Junos® OS

Multicast VPNs Configuration Guide

Release
12.1

Published: 2012-03-13

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS Multicast VPNs Configuration Guide
12.1

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

   ```
   [edit]
   user@host# edit system scripts
   [edit system scripts]
   ```


3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

   [edit system scripts]
   user@host# **load merge relative /var/tmp/ex-script-snippet.conf**
   load complete

For more information about the **load** command, see the *Junos OS CLI User Guide*.

Documentation Conventions

Table 1 on page xiii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page xiii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*

Table 2: Text and Syntax Conventions *(continued)*

<table>
<tr><th>Convention</th><th>Description</th><th>Examples</th></tr>
<tr><td>*Italic text like this*</td><td>Represents variables (options for which you substitute a value) in commands or configuration statements.</td><td>Configure the machine's domain name:

[edit]
root@# **set system domain-name** *domain-name*</td></tr>
<tr><td>Text like this</td><td>Represents names of configuration statements, commands, files, and directories; interface names; configuration hierarchy levels; or labels on routing platform components.</td><td>• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level.
• The console port is labeled **CONSOLE**.</td></tr>
<tr><td>< > (angle brackets)</td><td>Enclose optional keywords or variables.</td><td>**stub** <**default-metric** *metric*>;</td></tr>
<tr><td>| (pipe symbol)</td><td>Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.</td><td>**broadcast | multicast**

(*string1* | *string2* | *string3*)</td></tr>
<tr><td># (pound sign)</td><td>Indicates a comment specified on the same line as the configuration statement to which it applies.</td><td>**rsvp { # Required for dynamic MPLS only**</td></tr>
<tr><td>[ ] (square brackets)</td><td>Enclose a variable for which you can substitute one or more values.</td><td>**community name members [** *community-ids* **]**</td></tr>
<tr><td>Indention and braces ({ })</td><td>Identify a level in the configuration hierarchy.</td><td rowspan="2">[edit]
routing-options {
  static {
    route default {
      nexthop *address*;
      retain;
    }
  }
}</td></tr>
<tr><td>; (semicolon)</td><td>Identifies a leaf statement at a configuration hierarchy level.</td></tr>
<tr><td colspan="3">J-Web GUI Conventions</td></tr>
<tr><td>**Bold text like this**</td><td>Represents J-Web graphical user interface (GUI) items you click or select.</td><td>• In the Logical Interfaces box, select **All Interfaces**.
• To cancel the configuration, click **Cancel**.</td></tr>
<tr><td>**>** (bold right angle bracket)</td><td>Separates levels in a hierarchy of J-Web selections.</td><td>In the configuration editor hierarchy, select **Protocols>Ospf**.</td></tr>
</table>

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to **techpubs-comments@juniper.net**, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf .
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/ .
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/ .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html .

PART 1

Overview

CHAPTER 1

Introduction to Multicast VPNs

MBGP Multicast VPN Sites

There are two ways to implement Layer 3 VPNs using PIM, or multicast VPNs (MVPNs). There are no official names for the two methods: using dual PIM MVPNs (also known informally as “draft-rosen”) and multiprotocol BGP (MBGP)-based MVPNs (the “next generation” method of MVPN configuration). Both methods are supported and equally effective, but the MBGP-based MVPN method does not require multicast configuration on the service provider backbone. In other words, the PIM state information is maintained between the PE routers using the same architecture that is used for unicast VPNs. The main advantage of deploying MVPNs with MBGP is simplicity of configuration and operation because multicast is not needed on the service provider VPN backbone connecting the PE routers.

Multiprotocol BGP multicast VPNs employ the intra-autonomous system (AS) next-generation BGP control plane and PIM sparse mode as the data plane.

There are several multicast applications driving the deployment of next-generation Layer 3 MVPNs. Some of the key emerging applications include the following:

- Layer 3 VPN multicast service offered by service providers to enterprise customers
- Video transport applications for wholesale IPTV and multiple content providers attached to the same network
- Distribution of media-rich financial services or enterprise multicast services
- Multicast backhaul over a metro network

The main characteristics of MBGP MVPNs are:

- They extend Layer 3 VPN service (RFC 4364) to support IP multicast for Layer 3 VPN service providers.
- They follow the same architecture as specified by RFC 4364 for unicast VPNs. Specifically, BGP is used as the provider edge (PE) router-to-PE router control plane for multicast VPN.

- They eliminate the requirement for the virtual router (VR) model (as specified in Internet draft draft-rosen-vpn-mcast, *Multicast in MPLS/BGP VPNs*) for multicast VPNs and the RFC 4364 model for unicast VPNs.
- They rely on RFC 4364-based unicast with extensions for intra-AS and inter-AS communication.

An MBGP MVPN defines two types of site sets, a sender site set and a receiver site set. These sites have the following properties:

- Hosts within the sender site set can originate multicast traffic for receivers in the receiver site set.
- Receivers outside the receiver site set should not be able to receive this traffic.
- Hosts within the receiver site set can receive multicast traffic originated by any host in the sender site set.
- Hosts within the receiver site set should not be able to receive multicast traffic originated by any host that is not in the sender site set.

A site can be in both the sender site set and the receiver site set, so hosts within such a site can both originate and receive multicast traffic. For example, the sender site set could be the same as the receiver site set, in which case all sites could both originate and receive multicast traffic from one another.

Sites within a given MBGP MVPN might be within the same organization or in different organizations, which means that an MBGP MVPN can be either an intranet or an extranet. A given site can be in more than one MBGP MVPN, so MBGP MVPNs might overlap. Not all sites of a given MBGP MVPN have to be connected to the same service provider, meaning that an MBGP MVPN can span multiple service providers. Feature parity for the MVPN extranet functionality or overlapping MVPNs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

Another way to look at an MBGP MVPN is to say that an MBGP MVPN is defined by a set of administrative policies. These policies determine both the sender site set and the receiver site set. These policies are established by MBGP MVPN customers, but implemented by service providers using the existing BGP and MPLS VPN infrastructure.

Related Documentation

- Example: Allowing MBGP MVPN Remote Sources
- Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN
- Example: Configuring MBGP Multicast VPN Extranets

PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs

You can configure PIM sparse mode, PIM dense mode, auto-RP, and bootstrap router (BSR) for MBGP MVPN networks:

- PIM sparse mode—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode includes an explicit join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from the receivers to the rendezvous point (RP).
- PIM dense mode—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. Packets are forwarded to all interfaces except the incoming interface. Unlike PIM sparse mode, where explicit joins are required for packets to be transmitted downstream, packets are flooded to all routers in the routing instance in PIM dense mode.
- Auto-RP—Uses PIM dense mode to propagate control messages and establish RP mapping. You can configure an auto-RP node in one of three different modes: discovery mode, announce mode, and mapping mode.
- BSR—Establishes RPs. A selected router in a network acts as a BSR, which selects a unique RP for different group ranges. BSR messages are flooded using a data tunnel between PE routers.

Related Documentation

- Example: Allowing MBGP MVPN Remote Sources
- Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN
- Example: Configuring MBGP Multicast VPN Extranets

CHAPTER 2

Introduction to Configuring Multicast VPNs

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Selective LSPs are also referred to as selective provider tunnels. Selective provider tunnels carry traffic from some multicast groups in a VPN and extend only to the PE routers that have receivers for these groups. You can configure a selective provider tunnel for group prefixes and source prefixes, or you can use wildcards for the group and source, as described in the Internet draft draft-rekhter-mvpn-wildcard-spmsi-01.txt, *Use of Wildcard in S-PMSI Auto-Discovery Routes*.

The following sections describe the scenarios and special considerations when you use wildcards for selective provider tunnels.

About S-PMSI

The provider multicast service interface (PMSI) is a BGP tunnel attribute that contains the tunnel ID used by the PE router for transmitting traffic through the core of the provider network. A selective PMSI (S-PMSI) autodiscovery route advertises binding of a given MVPN customer multicast flow to a particular provider tunnel. The S-PMSI autodiscovery route advertised by the ingress PE router contains /32 IPv4 or /128 IPv6 addresses for the customer source and the customer group derived from the source-tree customer multicast route.

Figure 1 on page 8 shows a simple MVPN topology. The ingress router, PE1, originates the S-PMSI autodiscovery route. The egress routers, PE2 and PE3, have join state as a result of receiving join messages from CE devices that are not shown in the topology. In response to the S-PMSI autodiscovery route advertisement sent by PE1, PE2, and PE3, elect whether or not to join the tunnel based on the join state. The selective provider tunnel configuration is configured in a VRF instance on PE1.

NOTE: The MVPN mode configuration (RPT-SPT or SPT-only) is configured on all three PE routers for all VRFs that make up the VPN. If you omit the MVPN mode configuration, the default mode is SPT-only.

Figure 1: Simple MVPN Topology

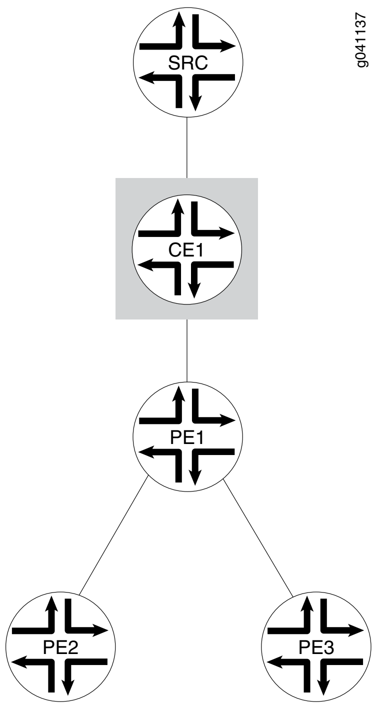

Scenarios for Using Wildcard S-PMSI

A wildcard S-PMSI has the source or the group (or both the source and the group) field set to the wildcard value of 0.0.0.0/0 and advertises binding of multiple customer multicast flows to a single provider tunnel in a single S-PMSI autodiscovery route.

The scenarios under which you might configure a wildcard S-PMSI are as follows:

- When the customer multicast flows are PIM-SM in ASM-mode flows. In this case, a PE router connected to an MVPN customer's site that contains the customer's RP (C-RP) could bind all the customer multicast flows traveling along a customer's RPT tree to a single provider tunnel.
- When a PE router is connected to an MVPN customer’s site that contains multiple sources, all sending to the same group.

- When the customer multicast flows are PIM-bidirectional flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows for the same group that have been originated within the sites of a given MVPN connected to that PE, and advertise such binding in a single S-PMSI autodiscovery route.
- When the customer multicast flows are PIM-SM in SSM-mode flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows coming from a given source located in a site connected to that PE router.
- When you want to carry in the provider tunnel all the customer multicast flows originated within the sites of a given MVPN connected to a given PE router.

Types of Wildcard S-PMSI

The following types of wildcard S-PMSI are supported:

- A (*,G) S-PMSI matches all customer multicast routes that have the group address. The customer source address in the customer multicast route can be any address, including 0.0.0.0/0 for shared-tree customer multicast routes. A (*, C-G) S-PMSI autodiscovery route is advertised with the source field set to 0 and the source address length set to 0. The multicast group address for the S-PMSI autodiscovery route is derived from the customer multicast joins.
- A (*,*) S-PMSI matches all customer multicast routes. Any customer source address and any customer group address in a customer multicast route can be bound to the (*,*) S-PMSI. The S-PMSI autodiscovery route is advertised with the source address and length set to 0 and the group address and length set 0. The remaining fields in the S-PMSI autodiscovery route follow the same rule as (C-S, C-G) S-PMSI, as described in section 12.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

Differences Between Wildcard S-PMSI and (S,G) S-PMSI

For dynamic provider tunnels, each customer multicast stream is bound to a separate provider tunnel, and each tunnel is advertised by a separate S-PMSI autodiscovery route. For static LSPs, multiple customer multicast flows are bound to a single provider tunnel by having multiple S-PMSI autodiscovery routes advertise the same provider tunnel.

When you configure a wildcard (*,G) or (*,*) S-PMSI, one or more matching customer multicast routes share a single S-PMSI. All customer multicast routes that have a matching source and group address are bound to the same (*,G) or (*,*) S-PMSI and share the same tunnel. The (*,G) or (*,*) S-PMSI is established when the first matching remote customer multicast join message is received in the ingress PE router, and deleted when the last remote customer multicast join is withdrawn from the ingress PE router. Sharing a single S-PMSI autodiscovery route improves control plane scalability.

Wildcard (*,*) S-PMSI and PIM Dense Mode

For (S,G) and (*,G) S-PMSI autodiscovery routes in PIM dense mode (PIM-DM), all downstream PE routers receive PIM-DM traffic. If a downstream PE router does not have receivers that are interested in the group address, the PE router instantiates prune state and stops receiving traffic from the tunnel.

Now consider what happens for (*,*) S-PMSI autodiscovery routes. If the PIM-DM traffic is not bound by a longer matching (S,G) or (*,G) S-PMSI, it is bound to the (*,*) S-PMSI. As is always true for dense mode, PIM-DM traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. Because there is no group information in the (*,*) S-PMSI autodiscovery route, egress PE routers join a (*,*) S-PMSI tunnel if there is any configuration on the egress PE router indicating interest in PIM-DM traffic.

Interest in PIM-DM traffic is indicated if the egress PE router has one of the following configurations in the VRF instance that corresponds to the instance that imports the S-PMSI autodiscovery route:

- At least one interface is configured in dense mode at the **[edit routing-instances *instance-name* protocols pim interface]** hierarchy level.
- At least one group is configured as a dense-mode group at the **[edit routing-instances *instance-name* protocols pim dense-groups *group-address*]** hierarchy level.

Wildcard (*,*) S-PMSI and PIM-BSR

For (S,G) and (*,G) S-PMSI autodiscovery routes in PIM bootstrap router (PIM-BSR) mode, an ingress PE router floods the PIM bootstrap message (BSM) packets over the provider tunnel to all egress PE routers. An egress PE router does not join the tunnel unless the message has the ALL-PIM-ROUTERS group. If the message has this group, the egress PE router joins the tunnel, regardless of the join state. The group field in the message determines the presence or absence of the ALL-PIM-ROUTERS address.

Now consider what would happen for (*,*) S-PMSI autodiscovery routes used with PIM-BSR mode. If the PIM BSM packets are not bound by a longer matching (S,G) or (*,G) S-PMSI, they are bound to the (*,*) S-PMSI. As is always true for PIM-BSR, BSM packets are flooded to downstream PE routers over the provider tunnel to the ALL-PIM-ROUTERS destination group. Because there is no group information in the (*,*) S-PMSI autodiscovery route, egress PE routers always join a (*,*) S-PMSI tunnel. Unlike PIM-DM, the egress PE routers might have no configuration suggesting use of PIM-BSR as the RP discovery mechanism in the VRF instance. To prevent all egress PE routers from always joining the (*,*) S-PMSI tunnel, the (*,*) wildcard group configuration must be ignored.

This means that if you configure PIM-BSR, a wildcard-group S-PMSI can be configured for all other group addresses. The (*,*) S-PMSI is not used for PIM-BSR traffic. Either a matching (*,G) or (S,G) S-PMSI (where the group address is the ALL-PIM-ROUTERS group) or an inclusive provider tunnel is needed to transmit data over the provider core. For PIM-BSR, the longest-match lookup is (S,G), (*,G), and the inclusive provider tunnel, in that order. If you do not configure an inclusive tunnel for the routing instance, you must configure a (*,G) or (S,G) selective tunnel. Otherwise, the data is dropped. This is because PIM-BSR functions like PIM-DM, in that traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. However, unlike PIM-DM, the egress PE routers might have no configuration to indicate interest or noninterest in PIM-BSR traffic.

Wildcard Source and the 0.0.0.0/0 Source Prefix

You can configure a 0.0.0.0/0 source prefix and a wildcard source under the same group prefix in a selective provider tunnel. For example, the configuration might look as follows:

```
routing-instances {
  vpna {
    provider-tunnel {
      selective {
        group 224.1.1.0/24 {
          source 0.0.0.0/0 {
            rsvp-te {
              label-switched-path-template {
                sptnl3;
              }
            }
          }
          wildcard-source {
            rsvp-te {
              label-switched-path-template {
                sptnl2;
              }
              static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbps;
          }
        }
      }
    }
  }
}
```

The functions of the **source 0.0.0.0/0** and **wildcard-source** configuration statements are different. The 0.0.0.0/0 source prefix only matches (C-S, C-G) customer multicast join messages and triggers (C-S, C-G) S-PMSI autodiscovery routes derived from the customer multicast address. Because all (C-S, C-G) join messages are matched by the 0.0.0.0/0 source prefix in the matching group, the wildcard source S-PMSI is used only for (*,C-G) customer multicast join messages. In the absence of a configured 0.0.0.0/0 source prefix, the wildcard source matches (C-S, C-G) and (*,C-G) customer multicast join messages. In the example, a join message for (10.0.1.0/24, 224.1.1.0/24) is bound to **sptnl3**. A join message for (*, 224.1.1.0/24) is bound to **sptnl2**.

Related Documentation

- Configuring a Selective Provider Tunnel Using Wildcards on page 34
- Example: Configuring Selective Provider Tunnels Using Wildcards on page 35
- Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 19
- Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 21

PART 2

Configuration

CHAPTER 3

Configuring Multicast VPNs

Introduction to Configuring MBGP MVPNs

You configure multiprotocol BGP-based (MBGP) multicast VPNs (MVPNs) at a number of different hierarchy levels within the Junos OS. However, a majority of MBGP MVPN statements are configured within a routing instance as follows:

```
description text;
instance-type vrf;
interface interface-name;
route-distinguisher (as-number:number | ip-address:number);
vrf-export [policy-names];
vrf-import [policy-names];
vrf-target (community | export community-name | import community-name);
protocols {
  mvpn {
    mvpn-mode (rpt-spt | spt-only);
    receiver-site;
    sender-site;
```

```
        route-target {
          export-target {
            target target-community;
            unicast;
          }
          import-target {
            target {
              target-value;
              receiver target-value;
              sender target-value;
            }
            unicast {
              receiver;
              sender;
            }
          }
        }
      }
    }
    provider-tunnel {
      ldp-p2mp;
      pim-asm {
        group-address address;
      }
      pim-ssm {
        group-address address;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
      }
      selective {
        group multicast--prefix/prefix-length {
          source ip--prefix/prefix-length {
            ldp-p2mp;
            pim-ssm {
              group-range multicast-prefix;
            }
            rsvp-te {
              label-switched-path-template {
                (default-template | lsp-template-name);
              }
              static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbps;
          }
          wildcard-source {
            ldp-p2mp;
            pim-ssm {
              group-range multicast-prefix;
            }
            rsvp-te {
              label-switched-path-template {
                (default-template | lsp-template-name);
```



```
                }
                static-lsp point-to-multipoint-lsp-name;
            }
            threshold-rate kbps;
        }
    }
    tunnel-limit number;
    wildcard-group-inet {
        wildcard-source {
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
    wildcard-group-inet6 {
        wildcard-source {
            ldp-p2mp;
            pim-ssm {
                group-range multicast-prefix;
            }
            rsvp-te {
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
                static-lsp lsp-name;
            }
            threshold-rate number;
        }
    }
}
threshold-rate number;
}
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Related Documentation

- *Junos OS Multicast over Layer 3 VPNs Feature Guide*

Configuring Routing Instances for an MBGP MVPN

To configure MBGP MVPNs, include the **mvpn** statement:

```
mvpn {
```

```
    mvpn-mode (rpt-spt | spt-only);
    receiver-site;
    route-target {
      export-target {
        target target-community;
        unicast;
      }
      import-target {
        target {
          target-value;
          receiver target-value;
          sender target-value;
        }
        unicast {
          receiver;
          sender;
        }
      }
    }
    sender-site;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

By default an MBGP MVPN routing instance is associated with both the multicast sender and the receiver sites. If you configure the **receiver-site** option, the routing instance is associated with only multicast receiver sites. Configuring the **sender-site** option associates the routing instance with only multicast sender sites.

NOTE: When you configure the routing instance for the MBGP MVPN, you must configure MPLS LSPs (either RSVP-signaled or LDP-signaled) between the PE routers of the routing instance to ensure VPN unicast connectivity. Point-to-multipoint LSPs are used for multicast data forwarding only.

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation is shortest path tree only (SPT-only) mode. In SPT-only mode, the active multicast sources are learned through multicast VPN source-active routes. This mode of operation is described in section 14 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

In contrast to SPT-only mode, rendezvous point tree (RPT)-SPT mode (also known as shared-tree data distribution) supports the native PIM model of transmitting (*,G) messages from the receiver to the RP for intersite shared-tree join messages.

In SPT-only mode, when a PE router receives a (*, C-G) join message, the router looks for an active source transmitting data to the customer group. If the PE router has a source-active route for the customer group, the router creates a source tree customer multicast route and sends the route to the PE router connected to the VPN site with the source. The source is determined by MVPN's single-forwarder election. When a receiver sends a (*,G) join message in a VPN site, the (*,G) join message only travels as far as the PE router. After the join message is converted to a type 6 multicast route, which is equivalent to a (S,G) join message, the route is installed with the no-advertise community setting.

NOTE: The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

Single-forwarder election guarantees selection of a unique forwarder for a given customer source (C-S). The upstream PE router might differ for the source tree and the shared tree because the election is based on the customer source and C-RP, respectively. Although the single-forwarder election is sufficient for SPT-only mode, the alternative RPT-SPT mode involves procedures to prevent duplicate traffic from being sent on the shared tree and the source tree. These procedures might require administrator-configured parameters to reduce duplicate traffic and reduce blackholes during RPT to SPT switch and the reverse.

In SPT-only mode, when a source is active, PIM creates a register state for the source both on the DR and on the C-RP (or on a PE router that is running Multicast Source Discovery Protocol [MSDP] between itself and the C-RP). After the register states are created, MVPN creates a source-active route. These type 5 source-active routes are installed on all PE routers. When the egress PE router with the (*,G) join message receives the source-active route, it has two routes that it can combine to produce the (S,G) multicast route. The type 6 route informs the PE router that a receiver is interested in group G. The source active route informs the PE router that a source S is transmitting data to group G. MVPN combines this information to produce a multicast join message and advertises this to the ingress PE router, as determined by the single-forwarder election.

NOTE: If you are using a Trio FPC, the MVPN sender PE router can not have a local receiver.

For some service providers, the SPT-only implementation is not ideal because it creates a restriction on C-RP configuration. For a PE router to create customer multicast routes from (*, C-G) join messages, the router must learn about active sources through MVPN type 5 source-active routes. These source-active routes can be originated only by a PE router. This means that a PE router in the MVPN must learn about all PIM register messages sent to the RP, which is possible only in the following cases:

- The C-RP is colocated on one of the PEs in the MVPN.
- MSDP is run between the C-RP and the VRF instance on one of the PE routers in the MVPN.

If this restriction is not acceptable, providers can use RPT-SPT mode instead of the default SPT-only mode. However, because SPT-only mode does not transmit (*,G) routes between VPN sites, SPT-only mode has the following advantages over RPT-SPT mode:

- Simplified operations by exchanging and processing only source-tree customer multicast routes among PE routers
- Simplified operations by eliminating the need for the service provider to suppress MVPN transient duplicates during the switch from RPT to SPT
- Less control plane overhead in the service provider space by limiting the type of customer multicast routes exchanged, which results in more scalable deployments
- More stable traffic patterns in the backbone without the traffic shifts involved in the RPT-SPT mode
- Easier maintenance in the service provider space due to less state information

To configure SPT-only mode:

1. Explicitly configure SPT-only mode:

   [edit routing-instances *routing-instance-name* protocols mvpn mvpn-mode]
   user@router# set **spt-only**

2. Include the **spt-only** statement for all VRFs that make up the VPN.

Related Documentation

- Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 21

- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7

Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation supports only intersite shortest-path trees (SPTs) for customer PIM (C-PIM) join messages. It does not support rendezvous-point trees (RPTs) for C-PIM join messages. The default mode of operation provides advantages, but it requires either that the customer rendezvous point (C-RP) be located on a PE router or that the Multicast Source Discovery Protocol (MSDP) be used between the C-RP and a PE router so that the PE router can learn about active sources advertised by other PE routers.

If the default mode is not suitable for your environment, you can configure RPT-SPT mode (also known as *shared-tree data distribution*), as documented in section 13 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). RPT-SPT mode supports the native PIM model of transmitting (*,G) messages from the receiver to the RP for intersite shared-tree join messages. This means that the type 6 (*,G) routes get transmitted from one PE router to another. In RPT-SPT mode, the shared-tree multicast routes are advertised from an egress PE router to the upstream router connected to the VPN site with the C-RP. The single-forwarder election is performed for the C-RP rather than for the source. The egress PE router takes the upstream hop to advertise the (*,G) and sends the type 6 route toward the upstream PE router. To send the data on the RPT, either inclusive or selective provider tunnels can be used. After the data starts flowing on the RPT, the last-hop router switches to SPT mode, unless you include the **spt-threshold infinity** statements in the configuration.

NOTE: The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

The switch to SPT mode is performed by PIM and not by MVPN type 5 and type 6 routes. After the last-hop router switches to SPT mode, the SPT (S,G) join messages follow the same rules as the SPT-only default mode.

The advantage of RPT-SPT mode is that it provides a method for PE routers to discover sources in the multicast VPN when the C-RP is located on the customer site instead of on a PE router. Because the shared C-tree is established between VPN sites, there is no need to run MSDP between the C-RP and the PE routers. RPT-SPT mode also enables egress PE routers to switch to receiving data from the PE connected to the source after the source information is learned, instead of receiving data from the RP.

CAUTION: When you configure RPT-SPT mode, receivers or sources directly attached to the PE router are not supported. As a workaround, place a CE router between any receiver or source and the PE router.

To configure RPT-SPT mode:

1. Enable shared-tree data distribution:

   [edit routing-instances *routing-instance-name* protocols mvpn mvpn-mode]
   user@router# **set rpt-spt**

2. Include the **rpt-spt** statement for all VRFs that make up the VPN.

Related Documentation

- Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 19
- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7

Limiting Routes to Be Advertised by an MVPN VRF Instance

If a hub-and-spoke deployment uses one VPN routing and forwarding (VRF) routing instance for unicast routing and a separate VRF for MVPN routing, you need to limit the PE routers at the hub site to advertise only IPv4 MVPN routes, only IPv6 MVPN routes, or both. This is necessary to prevent the multicast VRF instance from advertising unicast VPN routes to other PE routers.

NOTE: This configuration does not prevent the exportation of VPN routes to other VRF instances on the same router if the **auto-export** statement is included in the **[edit routing-options]** hierarchy.

To configure a VRF routing instance with the name **green** to advertise MVPN routes from both the **inet** and **inet6** address families, perform the following steps:

1. Configure the VRF routing instance to advertise IPv4 routes.

   user@host# **set routing-instances green vrf-advertise-selective family inet-mvpn**

2. Configure the VRF routing instance to advertise IPv6 routes.

   user@host# **set routing-instances green vrf-advertise-selective family inet6-mvpn**

After the configuration is committed, only the MVPN routes for the specified address families are advertised from the VRF instance to remote PE routers. To remove the restriction on routes being advertised, delete the **vrf-advertise-selective** statement.

NOTE: You cannot include the **vrf-advertise-selective** statement and the **no-vrf-advertise** statement in the same VRF configuration.

Related Documentation

- family on page 88
- inet-mvpn on page 92
- inet6-mvpn on page 94
- vrf-advertise-selective on page 115

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

By default, the VPN routing and forwarding (VRF) import and export route targets (configured either using VRF import and export policies or using the **vrf-target** statement) are used for importing and exporting routes with the MBGP MVPN network layer reachability information (NLRI).

You can use the **export-target** and **import-target** statements to override the default VRF import and export route targets. Export and import targets can also be specified specifically for sender sites or receiver sites, or can be borrowed from a configured unicast route target. Note that a sender site export route target is always advertised when security association routes are exported.

NOTE: When you configure an MBGP MVPN routing instance, you should not configure a target value for an MBGP MVPN specific route target that is identical to a target value for a unicast route target configured in another routing instance.

Specifying route targets in the MBGP MVPN NLRI for sender and receiver sites is useful when there is a mix of sender only, receiver only, and sender and receiver sites. A sender site route target is used for exporting automatic discovery routes by a sender site and for importing automatic discovery routes by a receiver site. A receiver site route target is used for exporting routes by a receiver site and importing routes by a sender site. A sender and receiver site exports and imports routes with both route targets.

A provider edge (PE) router with sites in a specific MBGP MVPN must determine whether a received automatic discovery route is from a sender site or receiver site based on the following:

- If the PE router is configured to be only in a sender site, route targets are imported only from receiver sites. Imported automatic discovery routes must be from a receiver site.
- If the PE router is configured to be only in a receiver site, route targets are imported only from sender sites. Imported automatic discovery routes must be from a sender site.

- If a PE router is configured to be in both sender sites and receiver sites, these guidelines apply:
  - Along with an import route target, you can optionally configure whether the route target is from a receiver or a sender site.
  - If a configuration is not provided, an imported automatic discovery route is treated as belonging to both the sender site set and the receiver site set.

To configure a route target for the MBGP MVPN routing instance, include the **route-target** statement:

```
route-target {
  export-target {
    target target-community;
    unicast;
  }
  import-target {
    target {
      target-value;
      receiver target-value;
      sender target-value;
    }
    unicast {
      receiver;
      sender;
  }
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols mvpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]**

The following sections describes how to configure the export target and the import target for an MBGP MVPN:

- Configuring the Export Target for an MBGP MVPN on page 24
- Configuring the Import Target for an MBGP MVPN on page 25

Configuring the Export Target for an MBGP MVPN

To configure an export target, include the **export-target** statement:

```
export-target {
  target target-community;
  unicast;
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols mvpn route-target]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route target]**

Configure the **target** option to specify the export target community. Configure the **unicast** option to use the same target community that has been specified for unicast.

Configuring the Import Target for an MBGP MVPN

To configure an import target, include the **import-target** statement:

```
import-target {
  target target-value {
    receiver;
    sender;
  }
  unicast {
    receiver;
    sender;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols mvpn route-target]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target]**

The following sections describe how to configure the import target and unicast parameters:

- Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 25
- Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 26

Configuring the Import Target Receiver and Sender for an MBGP MVPN

To configure the import target community, include the **target** statement and specify the target community. The target community must by in the format **target:x:y**. The **x** value is either an IP address or an AS number followed by an optional **L** to indicate a 4 byte AS number, and **y** is a number (for example, **target:123456L:100**)

```
target target-value {
  receiver;
  sender;
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]**

You can specify the target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the target community used when importing receiver site sets.
- **sender**—Specify the target community used when importing sender site sets.

Configuring the Import Target Unicast Parameters for an MBGP MVPN

To configure a unicast target community as the import target, include the **unicast** statement:

```
unicast {
  receiver;
  sender;
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]**

You can specify the unicast target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the unicast target community used when importing receiver site sets.
- **sender**—Specify the unicast target community used when importing sender site sets.

Configuring NLRI Parameters for an MBGP MVPN

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv4 address family, include the **family inet-mvpn** statement:

```
inet-mvpn {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
    loops number;
    prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
  }
}
```

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv6 address family, include the **family inet6-mvpn** statement:

```
inet6-mvpn {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown percentage {
```

```
                idle-timeout (forever | minutes);
            }
        }
        loops number
        prefix-limit {
            maximum number;
            teardown percentage {
                idle-timeout (forever | minutes);
            }
        }
    }
}
```

Configuring PIM Provider Tunnels for an MBGP MVPN

To configure a Protocol Independent Multicast (PIM) sparse mode provider tunnel for a multicast VPN, include the **pim-asm** statement:

```
pim-asm {
    group-address address;
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]**

To complete the PIM sparse mode provider tunnel configuration, you also need to specify the group address using the **group-address** option. The source address for a PIM sparse mode provider tunnel is configured to be the loopback address of the loopback interface in the **inet.0** routing table.

Configuring PIM-SSM GRE Selective Provider Tunnels

This topic describes how to configure a PIM-SSM GRE selective provider tunnel for an MBGP MVPN. A selective provider tunnel uses a point-to-multipoint LSP.

Creating a selective provider tunnel enables you to move high-rate traffic off the inclusive tunnel and deliver the multicast traffic only to receivers that request it. This improves bandwidth utilization.

To configure a PIM-SSM GRE selective provider tunnel for the 224.1.1.1/32 customer multicast group address, the 10.2.2.2/32 customer source address, and a virtual routing instance named **green**:

1. Configure the multicast group address range to be used for creating selective tunnels. The address prefix can be any valid nonreserved IPv4 multicast address range. Whether you configure a range of addresses or a single address, make sure that you configure enough group addresses for all the selective tunnels needed.

   ```
   user@host# set routing-instances green provider-tunnel selective group 224.1.1.1/32
     source 10.2.2.2/32 pim-ssm group-range 232.1.1.0/24
   ```

2. Configure the threshold rate in kilobits per second (Kbps) for triggering the creation of the selective tunnel. If you set the threshold rate to zero Kbps, the selective tunnel is created immediately, and the multicast traffic does not use an inclusive tunnel at all. Optionally, you can leave the threshold rate unconfigured and the result is the same as setting the threshold to zero.

   user@host# **set routing-instances green provider-tunnel selective group 224.1.1.1/32 source 10.2.2.2/32 threshold-rate 0**

3. Configure the autonomous system number in the global routing options. This is required in MBGP MVPNs.

   user@host# **set routing-options autonomous-system 100**

When configuring PIM-SSM GRE selective provider tunnels, keep the following in mind:

- Aggregation of multiple customer multicast routes to a single PIM S-PMSI is not supported.
- Provider tunnel multicast group addresses must be IPv4 addresses, even in configurations in which the customer multicast group and source are IPv6 addresses.

Related Documentation

- Multicast VPN Terminology on page 83
- **pim-ssm on page 100**
- **group-range on page 90**
- **threshold-rate on page 111**

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

The Junos OS supports point-to-multipoint label-switched paths (LSPs) for MBGP MVPNs. Point-to-multipoint LSPs for multicast VPNs are supported for intra-autonomous system (AS) environments (within an AS), but are not supported for inter-AS environments (between autonomous systems). A point-to-multipoint LSP is an RSVP-signaled LSP with a single source and multiple destinations.

You can configure point-to-multipoint LSPs for MBGP MVPNs as follows:

- Static point-to-multipoint LSPs—Configure static point-to-multipoint LSPs using the standard MPLS LSP statements specified at the **[edit protocols mpls]** hierarchy level. You manually configure each of the leaf nodes for the point-to-multipoint LSP.
- Dynamic point-to-multipoint LSPs using the default template—Configuring dynamic point-to-multipoint LSPs using the **default-template** option causes the leaf nodes to be discovered automatically. The leaf nodes are discovered through BGP intra-AS automatic discovery. The **default-template** option allows you to minimize the amount of configuration needed. However, it does not allow you to configure any of the standard MPLS options.
- Dynamic point-to-multipoint LSPs using a user-configured template—Configuring dynamic point-to-multipoint LSPs using a user-configured template also causes the leaf nodes to be discovered automatically. By creating your own template for the

point-to-multipoint LSPs, all of the standard MPLS features (such as bandwidth allocation and traffic engineering) can be configured.

Be aware of the following properties for the egress PE router in a point-to-multipoint LSP configured for a multicast VPN:

- Penultimate hop-popping is not used by point-to-multipoint LSPs for multicast VPNs. Only ultimate hop-popping is used.
- You must configure either the **vrf-table-label** statement or a virtual loopback tunnel interface on the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends two copies of each packet over the link to the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is not a transit router for the point-to-multipoint LSP, the penultimate hop router can send just one copy of each packet over the link to the egress PE router.
- If you configure a virtual loopback tunnel interface on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends just one copy of each packet over the link to the egress PE router. A virtual loopback tunnel interface can perform two lookups on an incoming packet, one for the multicast MPLS lookup and one for the IP lookup.

NOTE: **Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.**

The following sections describe how to configure point-to-multipoint LSPs for MBGP MVPNs:

Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled inclusive point-to-multipoint LSPs for MBGP MVPNs. Aggregation is not supported, so you need to configure an inclusive point-to-multipoint LSP for each sender PE router in each multicast VPN routing instance. The sender PE router is in the sender site set of the MBGP MVPN.

To configure a static RSVP-signaled inclusive point-to-multipoint LSP, include the **static-lsp** statement:

static-lsp *lsp-name*;

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]**

- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]**

To configure dynamic inclusive point-to-multipoint LSPs, include the **label-switched-path-template** statement:

```
label-switched-path-template {
  (default-template | lsp-template-name);
}
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]**

You can configure either the **default-template** option or manually configure a point-to-multipoint LSP template and specify the template name.

Configuring Selective Provider Tunnels for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled selective point-to-multipoint LSPs (also referred to as selective provider tunnels) for MBGP MVPNs. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the multicast VPNs, helping to minimize flooding in the service provider's network.

As with inclusive point-to-multipoint LSPs, you can configure both dynamic and static selective tunnels for the multicast VPN.

To configure selective point-to-multipoint provider tunnels, include the **selective** statement:

```
selective {
  group multicast--prefix/prefix-length {
    source ip--prefix/prefix-length {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
      }
      threshold-rate kbps;
    }
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
```

```
                (default-template | lsp-template-name);
            }
            static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbps;
    }
  }
  tunnel-limit number;
  wildcard-group-inet {
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
      }
      threshold-rate number;
    }
  }
  wildcard-group-inet6 {
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
      }
      threshold-rate number;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]**

The following sections describe how to configure selective point-to-multipoint LSPs for MBGP MVPNs:

- Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 33
- Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 34

Configuring the Multicast Group Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, you need to specify a multicast group address by including the **group** statement:

group *address* { ... }

You can include this statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]**

The address must be a valid multicast group address. Multicast uses the Class D IP address range (**224.0.0.0** through **239.255.255.255**).

Configuring the Multicast Source Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, specify a multicast source address by including the **source** statement:

source *address* { ... }

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group *address*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address*]**

Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a static selective point-to-multipoint LSP for an MBGP MVPN. You need to configure a static LSP using the standard MPLS LSP statements at the **[edit protocols mpls]** hierarchy level. You then include the static LSP in your selective point-to-multipoint LSP configuration by using the **static-lsp** statement. Once this functionality is enabled on the source PE router, the static point-to-multipoint LSP is created based on your configuration.

To configure a static selective point-to-multipoint LSP, include the **rsvp-te** and the **static-lsp** statements:

rsvp-te static-lsp *lsp-name*;

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**

- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**

Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a dynamic selective point-to-multipoint LSP for an MBGP MVPN. The leaf nodes for a dynamic point-to-multipoint LSP can be automatically discovered using leaf automatic discovery routes. Selective provider multicast service interface (S-PMSI) automatic discovery routes are also supported.

To configure a dynamic selective point-to-multipoint provider tunnel, include the **rsvp-te** and **label-switched-path-template** statements:

```
rsvp-te label-switched-path-template {
  (default-template | lsp-template-name);
}
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**

The **label-switched-path-template** statement includes the following options:

- **default-template**—Specify that point-to-multipoint LSPs are generated dynamically based on the default template. No user configuration is required for the LSPs. However, the automatically generated LSPs include none of the common LSP features, such as bandwidth allocation and traffic engineering.
- ***lsp-template-name***—Specify the name of an LSP template to be used for the point-to-multipoint LSP. You need to configure the LSP template to be used as a basis for the point-to-multipoint LSPs. You can configure any of the common LSP features for this template.

Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

To configure a selective point-to-multipoint LSP dynamically, you need to specify the data threshold (in kilobits per second) required before a new tunnel is created using the **threshold-rate** statement:

```
threshold-rate number;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]**

Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

To configure a limit on the number of tunnels that can be generated for a dynamic point-to-multipoint LSP, include the **tunnel-limit** statement:

```
tunnel-limit number;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]**

Related Documentation

- Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

Configuring a Selective Provider Tunnel Using Wildcards

When you configure a selective provider tunnel for MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), you can use wildcards for the multicast group and source address prefixes. Using wildcards enables a PE router to use a single route to advertise the binding of multiple multicast streams of a given MVPN customer to a single provider's tunnel, as described in http://tools.ietf.org/html/draft-rekhter-mvpn-wildcard-spmsi-00 .

Sharing a single route improves control plane scalability because it reduces the number of S-PMSI autodiscovery routes.

To configure a selective provider tunnel using wildcards:

1. Configure a wildcard group matching any group IPv4 address and a wildcard source for (*,*) join messages.

   ```
   [edit routing-instances vpna provider-tunnel selective]
   user@router# set wildcard-group-inet wildcard-source
   ```

2. Configure a wildcard group matching any group IPv6 address and a wildcard source for (*,*) join messages.

   ```
   [edit routing-instances vpna provider-tunnel selective]
   user@router# set wildcard-group-inet6 wildcard-source
   ```

3. Configure an IP prefix of a multicast group and a wildcard source for (*,G) join messages.

   ```
   [edit routing-instances vpna provider-tunnel selective]
   user@router# set group 224.0.0/24 wildcard-source
   ```

4. Map the IPv4 join messages to a selective provider tunnel.

   ```
   [edit routing-instances vpna provider-tunnel selective wildcard-group-inet
     wildcard-source]
   user@router# set rsvp-te label-switched-path-template provider-tunnel1
   ```

5. Map the IPv6 join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective wildcard-group-inet6
  wildcard-source]
user@router# set rsvp-te label-switched-path-template provider-tunnel2
```

6. Map the (*,224.0.0/24) join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective group 224.0.0/24
  wildcard-source]
user@router# set rsvp-te label-switched-path-template provider-tunnel3
```

Related Documentation

- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7
- Example: Configuring Selective Provider Tunnels Using Wildcards on page 35
- Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 19
- Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 21

Example: Configuring Selective Provider Tunnels Using Wildcards

With the (*,G) and (*,*) S-PMSI, a customer multicast join message can match more than one S-PMSI. In this case, a customer multicast join message is bound to the longest matching S-PMSI. The longest match is a (S,G) S-PMSI, followed by a (*,G) S-PMSI and a (*,*) S-PMSI, in that order.

Consider the following configuration:

```
routing-instances {
  vpna {
    provider-tunnel {
      selective {
        wildcard-group-inet {
          wildcard-source {
            rsvp-te {
              label-switched-path-template {
                sptnl1;
              }
            }
          }
        }
        group 224.1.1.0/24 {
          wildcard-source {
            rsvp-te {
              label-switched-path-template {
                sptnl2;
              }
            }
          }
          source 10.1.1/24 {
            rsvp-te {
              label-switched-path-template {
```

```
                sptnl3;
              }
            }
          }
        }
      }
    }
  }
}
```

For this configuration, the longest-match rule works as follows:

- A customer multicast (10.1.1.1, 224.1.1.1) join message is bound to the sptnl3 S-PMSI autodiscovery route.
- A customer multicast (10.2.1.1, 224.1.1.1) join message is bound to the sptnl2 S-PMSI autodiscovery route.
- A customer multicast (10.1.1.1, 224.2.1.1) join message is bound to the sptnl1 S-PMSI autodiscovery route.

When more than one customer multicast route is bound to the same wildcard S-PMSI, only one S-PMSI autodiscovery route is created. An egress PE router always uses the same matching rules as the ingress PE router that advertises the S-PMSI autodiscovery route. This ensures consistent customer multicast mapping on the ingress and the egress PE routers.

Related Documentation

- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7
- Configuring a Selective Provider Tunnel Using Wildcards on page 34

Tracing MBGP MVPN Traffic and Operations

To trace MBGP MVPN traffic, you can specify options with the **traceoptions** statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]**
- **[edit routing-instances *routing-instance-name* protocols mvpn]**

The following trace flags display the operations associated with multicast VPNs:

- **all**—All multicast VPN tracing options
- **error**—Error conditions
- **general**—General events

- **nlri**—Multicast VPN advertisements received or sent by means of BGP
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing
- **topology**—Multicast VPN topology changes caused by reconfiguration or advertisements received from other PE routers using BGP

Configuring Internet Multicast Using Ingress Replication Provider Tunnels

The routing instance type mpls-internet-multicast uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, enabling a faster path for multicast traffic between sender and receiver routers in large-scale implementations.

The mpls-internet-multicast routing instance is a non-forwarding instance used only for control plane procedures; it does not support any interface configurations. Only one mpls-internet-multicast routing instance can be defined for a logical system. All multicast and unicast routes used for Internet multicast are associated only with the master instance (**inet.0**), not with the routing instance.

Each router participating in Internet multicast must be configured with BGP MPLS-based Internet multicast for control plane procedures and with ingress replication for the data provider tunnel, which forms a full mesh of MPLS point-to-point LSPs. The ingress replication tunnel can be selective or inclusive, matching the configuration of the provider tunnel in the routing instance.

The topology consists of routers on the edge of the IP multicast domain that have a set of IP interfaces and a set of MPLS core-facing interfaces, see Figure 2 on page 39. Internet multicast traffic is carried between the IP routers, through the MPLS cloud, using ingress replication tunnels for the data plane and a full-mesh IGBP session for the control plane.

Figure 2: Internet Multicast Topology

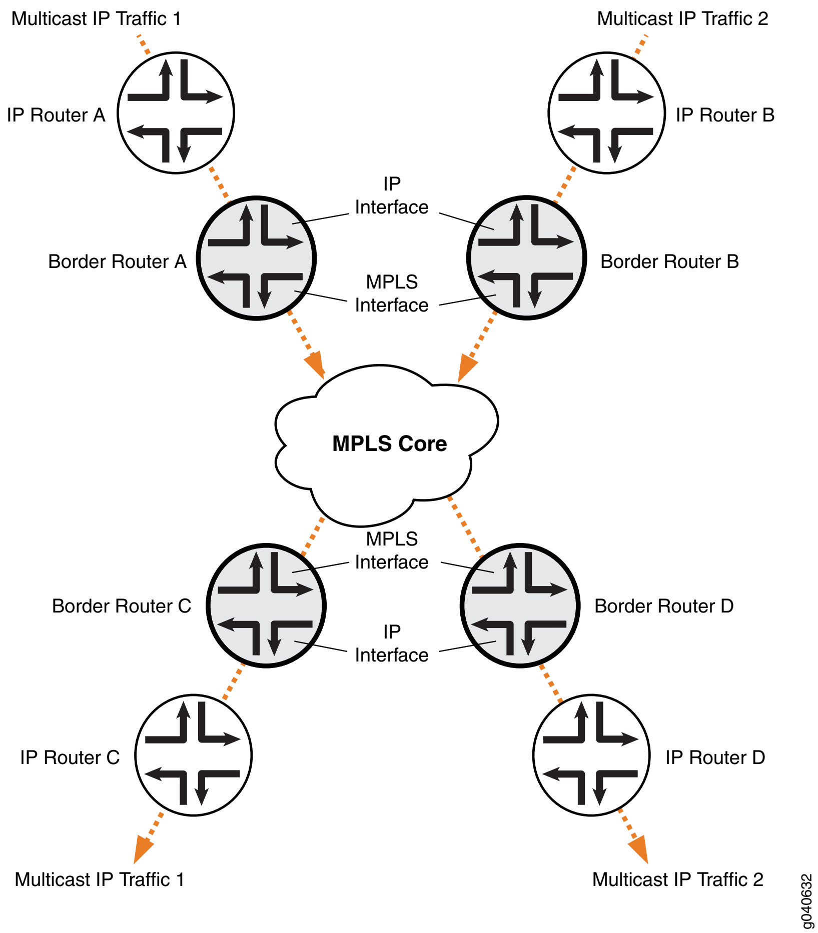

The mpls-internet-multicast routing instance type is configured for the default master instance on each router to support Internet multicast over MPLS. When using PIM as the multicast protocol, the mpls-internet-multicast configuration statement is also included at the **[edit protocols pim]** hierarchy level in the master instance. This creates a pseudo-interface that associates PIM with the mpls-internet-multicast routing instance.

When an application requests to add a destination to the ingress replication provider tunnel, the resulting behavior differs depending on which mode has been configured for the tunnel:

- existing-unicast-tunnel --- in this default mode, an existing unicast tunnel to the destination is used. If a unicast tunnel is not available, the destination is not added.
- create-new-ucast-tunnel --- when this mode is configured, a new unicast tunnel to the destination is added to the ingress replication provider tunnel, and is deleted if the application requests to delete the destination.

Example: Configure Internet Multicast Using Ingress Replication Tunnels

This example configures VPN-B with the instance type **mpls-internet-multicast**. This example also uses PIM for the multicast protocol.

1. Configure the routing instance type for VPN-B as **mpls-internet-multicast**:

   ```
   user@host# set routing-instances VPN-B instance-type
   mpls-internet-multicast
   ```

2. Configure the ingress replication provider tunnel to create a new unicast tunnel each time an application requests to add a destination:

   ```
   user@host# set routing-instances VPN-B provider-tunnel ingress-replication
   create-new-ucast-tunnel
   ```

3. Configure the point-to-point LSP to use the default template settings.

   ```
   user@host# set routing-instances VPN-B provider-tunnel ingress-replication
   label-switched-path label-switched-path-template default-template
   ```

4. Configure the ingress replication provider tunnel to be selective:

   ```
   user@host# set routing-instances VPN-B provider-tunnel selective group
   232.1.1.1/32 source 192.168.195.145/32 ingress-replication
   label-switched-path
   ```

5. Configure MVPN protocol in the routing instance:

   ```
   user@host# set routing-instances VPN-B protocols mvpn
   ```

6. Commit the configuration:

   ```
   user@host# commit
   ```

7. Use show command to verify the instance has been created:

   ```
   user@host# run show mvpn instance VPN-B

   MVPN instance:
   Legend for provider tunnel I-P-tnl -- inclusive provider tunnel S-P-tnl
   -- selective provider tunnel
   Legend for c-multicast routes properties (Pr)
   DS -- derived from (*, c-g)          RM -- remote VPN route
   Instance : VPN-B
      MVPN Mode : SPT-ONLY
     Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6

      Neighbor               I-P-tnl
      10.255.245.2           INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
      10.255.245.7           INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
      C-mcast IPv4 (S:G)     Ptnl                           St
      192.168.195.145/32:232.1.1.1/32 INGRESS-REPLICATION:MPLS Label
   18:10.255.245.6         RM
   ```

8. Add the **mpls-internet-multicast** configuration statement under the **[edit protocols pim]** hierarchy level in the master instance:

   ```
   user@host# set protocols pim mpls-internet-multicast
   ```

9. Commit the configuration:

   ```
   user@host# commit
   ```


10. Use **show ingress-replication mvpn** command to verify configuration settings:

```
user@host# run show ingress-replication mvpn

Ingress Tunnel: mvpn:1l
  Application: MVPN
  Unicast tunnels
    Leaf Address       Tunnel-type       Mode       State
    10.255.245.2       P2P LSP           New        Up
    10.255.245.4       P2P LSP           New        Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels
    Leaf Address       Tunnel-type       Mode       State
    10.255.245.2       P2P LSP           Existing   Up
```

11. Use this if you want to configure the ingress replication provider tunnel to be inclusive:

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
create-new-ucast-tunnel

user@host# set routing-instances VPN-B provider-tunnel ingress-replication
label-switched-path label-switched-path-template default-template
```

12. Use show mvpn instance command to verify tunnel is inclusive:

```
user@host# run show mypn instance VPN-B

MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance : VPN-A
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6

  Neighbor               I-P-tnl
  10.255.245.2           INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
  10.255.245.7           INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
  C-mcast IPv4 (S:G)     Ptnl                           St
  192.168.195.145/32:232.1.1.1/32 INGRESS-REPLICATION:MPLS Label
18:10.255.245.6         RM
```

Related Documentation

- create-new-ucast-tunnel on page 85
- existing-unicast-tunnel on page 86
- ingress-replication on page 95
- mpls-internet-multicast on page 97

CHAPTER 4

Additional Examples

Example: Configuring MBGP Multicast VPNs

This example provides a step-by-step procedure to configure multicast services across a multiprotocol BGP (MBGP) Layer 3 virtual private network.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.2 or later
- Five M Series, T Series, TX Series, or MX Series Juniper routers
- One host system capable of sending multicast traffic and supporting the Internet Group Management Protocol (IGMP)
- One host systems capable of receiving multicast traffic and supporting IGMP

Depending on the devices you are using, you might be required to configure static routes to:

- The multicast sender
- The Fast Ethernet interface to which the sender is connected on the multicast receiver
- The multicast receiver
- The Fast Ethernet interface to which the receiver is connected on the multicast sender

Overview and Topology

This example shows how to configure the following technologies:

- IPv4
- BGP
- OSPF
- RSVP
- MPLS
- PIM sparse mode
- Static RP

The topology of the network is shown in Figure 3 on page 44.

Figure 3: Multicast Over Layer 3 VPN Example Topology

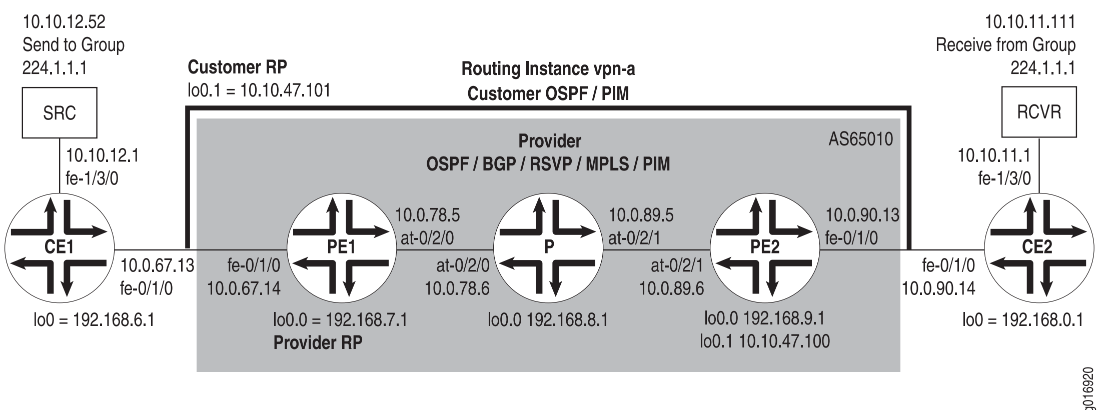

Configuration

NOTE: In any configuration session, it is a good practice to periodically verify that the configuration can be committed using the **commit check** command.

In this example, the router being configured is identified using the following command prompts:

- **CE1** identifies the customer edge 1 (CE1) router
- **PE1** identifies the provider edge 1 (PE1) router
- **P** identifies the provider core (P) router
- **CE2** identifies the customer edge 2 (CE2) router
- **PE2** identifies the provider edge 2 (PE2) router

To configure MBGP multicast VPNs for the network shown in Figure 3 on page 44, perform the following steps:

Configuring Interfaces

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode in the *Junos OS CLI User Guide*.

1. On each router, configure an IP address on the loopback logical interface 0 (**lo0.0**).

```
[edit interfaces]
user@CE1# set lo0 unit 0 family inet address 192.168.6.1/32 primary

user@PE1# set lo0 unit 0 family inet address 192.168.7.1/32 primary

user@P# set lo0 unit 0 family inet address 192.168.8.1/32 primary

user@PE2# set lo0 unit 0 family inet address 192.168.9.1/32 primary

user@CE2# set lo0 unit 0 family inet address 192.168.0.1/32 primary
```

   Use the **show interfaces terse** command to verify that the IP address is correct on the loopback logical interface.

2. On the PE and CE routers, configure the IP address and protocol family on the Fast Ethernet interfaces. Specify the **inet** protocol family type.

```
[edit interfaces]
user@CE1# set fe-1/3/0 unit 0 family inet address 10.10.12.1/24
user@CE1# set fe-0/1/0 unit 0 family inet address 10.0.67.13/30

[edit interfaces]
user@PE1# set fe-0/1/0 unit 0 family inet address 10.0.67.14/30

[edit interfaces]
user@PE2# set fe-0/1/0 unit 0 family inet address 10.0.90.13/30

[edit interfaces]
user@CE2# set fe-0/1/0 unit 0 family inet address 10.0.90.14/30
```



```
user@CE2# set fe-1/3/0 unit 0 family inet address 10.10.11.1/24
```

Use the **show interfaces terse** command to verify that the IP address is correct on the Fast Ethernet interfaces.

3. On the PE and P routers, configure the ATM interfaces' VPI and maximum virtual circuits. If the default PIC type is different on directly connected ATM interfaces, configure the PIC type to be the same. Configure the logical interface VCI, protocol family, local IP address, and destination IP address.

```
[edit interfaces]
user@PE1# set at-0/2/0 atm-options pic-type atm1
user@PE1# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@PE1# set at-0/2/0 unit 0 vci 0.128
user@PE1# set at-0/2/0 unit 0 family inet address 10.0.78.5/32 destination 10.0.78.6
```

```
[edit interfaces]
user@P# set at-0/2/0 atm-options pic-type atm1
user@P# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/0 unit 0 vci 0.128
user@P# set at-0/2/0 unit 0 family inet address 10.0.78.6/32 destination 10.0.78.5
user@P# set at-0/2/1 atm-options pic-type atm1
user@P# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/1 unit 0 vci 0.128
user@P# set at-0/2/1 unit 0 family inet address 10.0.89.5/32 destination 10.0.89.6
```

```
[edit interfaces]
user@PE2# set at-0/2/1 atm-options pic-type atm1
user@PE2# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@PE2# set at-0/2/1 unit 0 vci 0.128
user@PE2# set at-0/2/1 unit 0 family inet address 10.0.89.6/32 destination 10.0.89.5
```

Use the **show configuration interfaces** command to verify that the ATM interfaces' VPI and maximum VCs are correct and that the logical interface VCI, protocol family, local IP address, and destination IP address are correct.

Configuring OSPF

Step-by-Step Procedure

1. On the P and PE routers, configure the provider instance of OSPF. Specify the **lo0.0** and ATM core-facing logical interfaces. The provider instance of OSPF on the PE router forms adjacencies with the OSPF neighbors on the other PE router and Router P.

```
user@PE1# set protocols ospf area 0.0.0.0 interface at-0/2/0.0
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@P# set protocols ospf area 0.0.0.0 interface lo0.0
user@P# set protocols ospf area 0.0.0.0 interface all
user@P# set protocols ospf area 0.0.0.0 interface fxp0 disable
```

```
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.0
user@PE2# set protocols ospf area 0.0.0.0 interface at-0/2/1.0
```

Use the **show ospf interfaces** command to verify that the **lo0.0** and ATM core-facing logical interfaces are configured for OSPF.

2. On the CE routers, configure the customer instance of OSPF. Specify the loopback and Fast Ethernet logical interfaces. The customer instance of OSPF on the CE routers form adjacencies with the neighbors within the VPN routing instance of OSPF on the PE routers.

```
user@CE1# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@CE2# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface lo0.0
```

Use the **show ospf interfaces** command to verify that the correct loopback and Fast Ethernet logical interfaces have been added to the OSPF protocol.

3. On the P and PE routers, configure OSPF traffic engineering support for the provider instance of OSPF.

The **shortcuts** statement enables the master instance of OSPF to use a label-switched path as the next hop.

```
user@PE1# set protocols ospf traffic-engineering shortcuts
```

```
user@P# set protocols ospf traffic-engineering shortcuts
```

```
user@PE2# set protocols ospf traffic-engineering shortcuts
```

Use the **show ospf overview** or **show configuration protocols ospf** command to verify that traffic engineering support is enabled.

Configuring BGP

Step-by-Step Procedure

1. On Router P, configure BGP for the VPN. The local address is the local **lo0.0** address. The neighbor addresses are the PE routers' **lo0.0** addresses.

The **unicast** statement enables the router to use BGP to advertise network layer reachability information (NLRI). The **signaling** statement enables the router to use BGP as the signaling protocol for the VPN.

```
user@P# set protocols bgp group group-mvpn type internal
user@P# set protocols bgp group group-mvpn local-address 192.168.8.1
user@P# set protocols bgp group group-mvpn family inet unicast
user@P# set protocols bgp group group-mvpn family inet-mvpn signaling
user@P# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@P# set protocols bgp group group-mvpn neighbor 192.168.7.1
```

Use the **show configuration protocols bgp** command to verify that the router has been configured to use BGP to advertise NLRI.

2. On the PE and P routers, configure the BGP local autonomous system number.

```
user@PE1# set routing-options autonomous-system 0.65010
```



```
user@P# set routing-options autonomous-system 0.65010
```

```
user@PE2# set routing-options autonomous-system 0.65010
```

Use the **show configuration routing-options** command to verify that the BGP local autonomous system number is correct.

3. On the PE routers, configure BGP for the VPN. Configure the local address as the local **lo0.0** address. The neighbor addresses are the **lo0.0** addresses of Router P and the other PE router, PE2.

```
user@PE1# set protocols bgp group group-mvpn type internal
user@PE1# set protocols bgp group group-mvpn local-address 192.168.7.1
user@PE1# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE1# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

```
user@PE2# set protocols bgp group group-mvpn type internal
user@PE2# set protocols bgp group group-mvpn local-address 192.168.9.1
user@PE2# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE2# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.7.1
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

Use the **show bgp group** command to verify that the BGP configuration is correct.

4. On the PE routers, configure a policy to export the BGP routes into OSPF.

```
user@PE1# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE1# set policy-options policy-statement bgp-to-ospf then accept
```

```
user@PE2# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE2# set policy-options policy-statement bgp-to-ospf then accept
```

Use the **show policy bgp-to-ospf** command to verify that the policy is correct.

Configuring RSVP

Step-by-Step Procedure

1. On the PE routers, enable RSVP on the interfaces that participate in the LSP. Configure the Fast Ethernet and ATM logical interfaces.

```
user@PE1# set protocols rsvp interface fe-0/1/0.0
user@PE1# set protocols rsvp interface at-0/2/0.0
```

```
user@PE2# set protocols rsvp interface fe-0/1/0.0
user@PE2# set protocols rsvp interface at-0/2/1.0
```

2. On Router P, enable RSVP on the interfaces that participate in the LSP. Configure the ATM logical interfaces.

```
user@P# set protocols rsvp interface at-0/2/0.0
user@P# set protocols rsvp interface at-0/2/1.0
```

Use the **show configuration protocols rsvp** command to verify that the RSVP configuration is correct.

Configuring MPLS

Step-by-Step Procedure

1. On the PE routers, configure an MPLS LSP to the PE router that is the LSP egress point. Specify the IP address of the **lo0.0** interface on the router at the other end of the LSP. Configure MPLS on the ATM, Fast Ethernet, and **lo0.0** interfaces.

   To help identify each LSP when troubleshooting, configure a different LSP name on each PE router. In this example, we use the name **to-pe2** as the name for the LSP configured on PE1 and **to-pe1** as the name for the LSP configured on PE2.

   ```
   user@PE1# set protocols mpls label-switched-path to-pe2 to 192.168.9.1
   user@PE1# set protocols mpls interface fe-0/1/0.0
   user@PE1# set protocols mpls interface at-0/2/0.0
   user@PE1# set protocols mpls interface lo0.0
   ```

   ```
   user@PE2# set protocols mpls label-switched-path to-pe1 to 192.168.7.1
   user@PE2# set protocols mpls interface fe-0/1/0.0
   user@PE2# set protocols mpls interface at-0/2/1.0
   user@PE2# set protocols mpls interface lo0.0
   ```

   Use the **show configuration protocols mpls** and **show route label-switched-path to-pe1** commands to verify that the MPLS and LSP configuration is correct.

   After the configuration is committed, use the **show mpls lsp name to-pe1** and **show mpls lsp name to-pe2** commands to verify that the LSP is operational.

2. On Router P, enable MPLS. Specify the ATM interfaces connected to the PE routers.

   ```
   user@P# set protocols mpls interface at-0/2/0.0
   user@P# set protocols mpls interface at-0/2/1.0
   ```

   Use the **show mpls interface** command to verify that MPLS is enabled on the ATM interfaces.

3. On the PE and P routers, configure the protocol family on the ATM interfaces associated with the LSP. Specify the **mpls** protocol family type.

   ```
   user@PE1# set interfaces at-0/2/0 unit 0 family mpls
   ```

   ```
   user@P# set interfaces at-0/2/0 unit 0 family mpls
   user@P# set interfaces at-0/2/1 unit 0 family mpls
   ```

   ```
   user@PE2# set interfaces at-0/2/1 unit 0 family mpls
   ```

   Use the **show mpls interface** command to verify that the MPLS protocol family is enabled on the ATM interfaces associated with the LSP.

Configuring the VRF Routing Instance

Step-by-Step Procedure

1. On the PE routers, configure a routing instance for the VPN and specify the **vrf** instance type. Add the Fast Ethernet and **lo0.1** customer-facing interfaces. Configure the VPN instance of OSPF and include the BGP-to-OSPF export policy.

   ```
   user@PE1# set routing-instances vpn-a instance-type vrf
   user@PE1# set routing-instances vpn-a interface lo0.1
   user@PE1# set routing-instances vpn-a interface fe-0/1/0.0
   ```



```
user@PE1# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE1# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

```
user@PE2# set routing-instances vpn-a instance-type vrf
user@PE2# set routing-instances vpn-a interface lo0.1
user@PE2# set routing-instances vpn-a interface fe-0/1/0.0
user@PE2# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE2# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

Use the **show configuration routing-instances vpn-a** command to verify that the routing instance configuration is correct.

2. On the PE routers, configure a route distinguisher for the routing instance. A route distinguisher allows the router to distinguish between two identical IP prefixes used as VPN routes. Configure a different route distinguisher on each PE router. This example uses 65010:1 on PE1 and 65010:2 on PE2.

```
user@PE1# set routing-instances vpn-a route-distinguisher 65010:1
```

```
user@PE2# set routing-instances vpn-a route-distinguisher 65010:2
```

Use the **show configuration routing-instances vpn-a** command to verify that the route distinguisher is correct.

3. On the PE routers, configure default VRF import and export policies. Based on this configuration, BGP automatically generates local routes corresponding to the route target referenced in the VRF import policies. This example uses 2:1 as the route target.

NOTE: You must configure the same route target on each PE router for a given VPN routing instance.

```
user@PE1# set routing-instances vpn-a vrf-target target:2:1
```

```
user@PE2# set routing-instances vpn-a vrf-target target:2:1
```

Use the **show configuration routing-instances vpn-a** command to verify that the route target is correct.

4. On the PE routers, configure the VPN routing instance for multicast support.

```
user@PE1# set routing-instances vpn-a protocols mvpn
```

```
user@PE2# set routing-instances vpn-a protocols mvpn
```

Use the **show configuration routing-instance vpn-a** command to verify that the VPN routing instance has been configured for multicast support.

5. On the PE routers, configure an IP address on loopback logical interface 1 (**lo0.1**) used in the customer routing instance VPN.

```
user@PE1# set interfaces lo0 unit 1 family inet address 10.10.47.101/32
```



```
user@PE2# set interfaces lo0 unit 1 family inet address 10.10.47.100/32
```

Use the **show interfaces terse** command to verify that the IP address on the loopback interface is correct.

Configuring PIM

Step-by-Step Procedure

1. On the PE and P routers, enable the provider instance of PIM. Add the core-facing ATM interfaces. On the PE routers, also configure the **lo0.0** interface. Specify the mode as **sparse** and the version as **2**.

```
user@PE1# set protocols pim interface at-0/2/0.0 mode sparse
user@PE1# set protocols pim interface at-0/2/0.0 version 2
user@PE1# set protocols pim interface lo0.0 mode sparse
user@PE1# set protocols pim interface lo0.0 version 2

user@P# set protocols pim interface at-0/2/0.0 mode sparse
user@P# set protocols pim interface at-0/2/0.0 version 2
user@P# set protocols pim interface at-0/2/1.0 mode sparse
user@P# set protocols pim interface at-0/2/1.0 version 2

user@PE2# set protocols pim interface at-0/2/1.0 mode sparse
user@PE2# set protocols pim interface at-0/2/1.0 version 2
user@PE2# set protocols pim interface lo0.0 mode sparse
user@PE2# set protocols pim interface lo0.0 version 2
```

Use the **show pim interfaces** command to verify that PIM sparse-mode is enabled on the core-facing ATM interfaces.

2. On the PE routers, enable the VPN customer instance of PIM. Configure the **lo0.1** and the customer-facing Fast Ethernet interface. Specify the mode as **sparse** and the version as **2**.

```
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode
  sparse
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version
  2

user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode
  sparse
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version
  2
```

Use the **show pim interfaces instance vpn-a** command to verify that PIM sparse-mode is enabled on the **lo0.1** interface and the customer-facing Fast Ethernet interface.

3. On the CE routers, enable the customer instance of PIM. In this example, we configure all interfaces. Specify the mode as **sparse** and the version as **2**.

```
user@CE1# set protocols pim interface all
```

user@CE2# **set protocols pim interface all mode sparse**
user@CE2# **set protocols pim interface all version 2**

Use the **show pim interfaces** command to verify that PIM sparse mode is enabled on all interfaces.

Configuring the Provider Tunnel

Step-by-Step Procedure

1. On Router PE1, configure the provider tunnel. Specify the multicast address to be used.

   The **provider-tunnel** statement instructs the router to send multicast traffic across a tunnel. The **pim-asm** statement instructs the router to accept the multicast stream from any source.

   user@PE1# **set routing-instances vpn-a provider-tunnel pim-asm group-address 224.1.1.1**

   Use the **show configuration routing-instance vpn-a** command to verify that the multicast group address is correct on Router PE1.

2. On Router PE2, configure the provider tunnel. Specify the multicast address to be used.

   user@PE2# **set routing-instances vpn-a provider-tunnel pim-asm group-address 224.1.1.1**

   Use the **show configuration routing-instance vpn-a** command to verify that the multicast group address is correct on Router PE2.

Configuring the Rendezvous Point

Step-by-Step Procedure

1. Configure Router PE1 to be the rendezvous point for the provider instance of PIM. Specify the **lo0.0** address of Router PE1.

   user@PE1# **set protocols pim rp local address 192.168.7.1**

   Use the **show pim rps** command to verify that to correct local IP address is configured for the provider instance RP.

2. Configure the static rendezvous point on Router P and the PE2 router for the provider instance of PIM. Specify the **lo0.0** address of Router PE1. Specify the version as **2**.

   user@P# **set protocols pim rp static address 192.168.7.1 version 2**

   user@PE2# **set protocols pim rp static address 192.168.7.1 version 2**

   Use the **show pim rps** command to verify that the correct static IP address is configured for the provider instance RP.

3. Configure Router PE1 to be the rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1. Specify the multicast address to be used.

   user@PE1# **set routing-instances vpn-a protocols pim rp local address 10.10.47.101**
   user@PE1# **set routing-instances vpn-a protocols pim rp local group-ranges 224.1.1.1/32**

Use the **show pim rps instance vpn-a** command to verify that the correct local IP address is configured for the customer instance RP.

4. On Router PE2, configure the static rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1.

   user@PE2# **set routing-instances vpn-a protocols pim rp static address 10.10.47.101**

   Use the **show pim rps instance vpn-a** command to verify that the correct static IP address is configured for the customer instance RP.

5. On the CE routers, configure the static rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1.

   user@CE1# **set protocols pim rp static address 10.10.47.101 version 2**

   user@CE2# **set protocols pim rp static address 10.10.47.101 version 2**

   Use the **show pim rps** command to verify that the correct static IP address is configured for the customer instance RP.

6. Use the **commit check** command to verify that the configuration can be successfully committed. If the configuration passes the check, commit the configuration.
7. Start the multicast sender device connected to CE1.
8. Start the multicast receiver device connected to CE2.
9. Verify that the receiver is receiving the multicast stream.
10. Use **show** commands to verify the routing, VPN, and multicast operation.

Results The configuration and verification parts of this example have been completed. The following section is for your reference.

The relevant sample configuration for Router CE1 follows.

Router CE1

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.168.6.1/32 {
          primary;
        }
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.67.13/30;
      }
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.12.1/24;
```

```
            }
         }
      }
   }
   protocols {
      ospf {
         area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface lo0.0;
            interface fe-1/3/0.0;
         }
      }
      pim {
         rp {
            static {
               address 10.10.47.101 {
                  version 2;
               }
            }
         }
         interface all;
      }
   }
```

The relevant sample configuration for Router PE1 follows.

Router PE1

```
   interfaces {
      lo0 {
         unit 0 {
            family inet {
               address 192.168.7.1/32 {
                  primary;
               }
            }
         }
      }
      fe-0/1/0 {
         unit 0 {
            family inet {
               address 10.0.67.14/30;
            }
         }
      }
      at-0/2/0 {
         atm-options {
            pic-type atm1;
            vpi 0 {
               maximum-vcs 256;
            }
         }
         unit 0 {
            vci 0.128;
            family inet {
               address 10.0.78.5/32 {
                  destination 10.0.78.6;
               }
```



```
            }
            family mpls;
        }
    }
    lo0 {
        unit 1 {
            family inet {
                address 10.10.47.101/32;
            }
        }
    }
}
routing-options {
    autonomous-system 0.65010;
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface at-0/2/0.0;
    }
    mpls {
        label-switched-path to-pe2 {
            to 192.168.9.1;
        }
        interface fe-0/1/0.0;
        interface at-0/2/0.0;
        interface lo0.0;
    }
    bgp {
        group group-mvpn {
            type internal;
            local-address 192.168.7.1;
            family inet-vpn {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            neighbor 192.168.9.1;
            neighbor 192.168.8.1;
        }
    }
    ospf {
        traffic-engineering {
            shortcuts;
        }
        area 0.0.0.0 {
            interface at-0/2/0.0;
            interface lo0.0;
        }
    }
    pim {
        rp {
            local {
                address 192.168.7.1;
            }
```

```
        }
        interface at-0/2/0.0 {
            mode sparse;
            version 2;
        }
        interface lo0.0 {
            mode sparse;
            version 2;
        }
    }
}
policy-options {
    policy-statement bgp-to-ospf {
        from protocol bgp;
        then accept;
    }
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface lo0.1;
        interface fe-0/1/0.0;
        route-distinguisher 65010:1;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
        vrf-target target:2:1;
        protocols {
            ospf {
                export bgp-to-ospf;
                area 0.0.0.0 {
                    interface all;
                }
            }
            pim {
                rp {
                    local {
                        address 10.10.47.101;
                        group-ranges {
                            224.1.1.1/32;
                        }
                    }
                }
                interface lo0.1 {
                    mode sparse;
                    version 2;
                }
                interface fe-0/1/0.0 {
                    mode sparse;
                    version 2;
                }
            }
            mvpn;
        }
```



```
    }
  }
```

The relevant sample configuration for Router P follows.

Router P

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.168.8.1/32 {
          primary;
        }
      }
    }
  }
  at-0/2/0 {
    atm-options {
      pic-type atm1;
      vpi 0 {
        maximum-vcs 256;
      }
    }
    unit 0 {
      vci 0.128;
      family inet {
        address 10.0.78.6/32 {
          destination 10.0.78.5;
        }
      }
      family mpls;
    }
  }
  at-0/2/1 {
    atm-options {
      pic-type atm1;
      vpi 0 {
        maximum-vcs 256;
      }
    }
    unit 0 {
      vci 0.128;
      family inet {
        address 10.0.89.5/32 {
          destination 10.0.89.6;
        }
      }
      family mpls;
    }
  }
}
routing-options {
  autonomous-system 0.65010;
}
protocols {
  rsvp {
    interface at-0/2/0.0;
```

```
        interface at-0/2/1.0;
    }
    mpls {
        interface at-0/2/0.0;
        interface at-0/2/1.0;
    }
    bgp {
        group group-mvpn {
            type internal;
            local-address 192.168.8.1;
            family inet {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            neighbor 192.168.9.1;
            neighbor 192.168.7.1;
        }
    }
    ospf {
        traffic-engineering {
            shortcuts;
        }
        area 0.0.0.0 {
            interface lo0.0;
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        rp {
            static {
                address 192.168.7.1 {
                    version 2;
                }
            }
        }
        interface at-0/2/0.0 {
            mode sparse;
            version 2;
        }
        interface at-0/2/1.0 {
            mode sparse;
            version 2;
        }
    }
}
```

The relevant sample configuration for Router PE2 follows.

Router PE2

```
interfaces {
    lo0 {
        unit 0 {
```

```
            family inet {
                address 192.168.9.1/32 {
                    primary;
                }
            }
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.13/30;
            }
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm1;
            vpi 0 {
                maximum-vcs 256;
            }
        }
        unit 0 {
            vci 0.128;
            family inet {
                address 10.0.89.6/32 {
                    destination 10.0.89.5;
                }
            }
            family mpls;
        }
    }
    lo0 {
        unit 1 {
            family inet {
                address 10.10.47.100/32;
            }
        }
    }
}
routing-options {
    autonomous-system 0.65010;
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
    mpls {
        label-switched-path to-pe1 {
            to 192.168.7.1;
        }
        interface lo0.0;
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
    bgp {
```

```
        group group-mvpn {
            type internal;
            local-address 192.168.9.1;
            family inet-vpn {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            neighbor 192.168.7.1;
            neighbor 192.168.8.1;
        }
    }
    ospf {
        traffic-engineering {
            shortcuts;
        }
        area 0.0.0.0 {
            interface lo0.0;
            interface at-0/2/1.0;
        }
    }
    pim {
        rp {
            static {
                address 192.168.7.1 {
                    version 2;
                }
            }
        }
        interface lo0.0 {
            mode sparse;
            version 2;
        }
        interface at-0/2/1.0 {
            mode sparse;
            version 2;
        }
    }
}
policy-options {
    policy-statement bgp-to-ospf {
        from protocol bgp;
        then accept;
    }
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface fe-0/1/0.0;
        interface lo0.1;
        route-distinguisher 65010:2;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
```



```
        }
        vrf-target target:2:1;
        protocols {
            ospf {
                export bgp-to-ospf;
                area 0.0.0.0 {
                    interface all;
                }
            }
            pim {
                rp {
                    static {
                        address 10.10.47.101;
                    }
                }
                interface fe-0/1/0.0 {
                    mode sparse;
                    version 2;
                }
                interface lo0.1 {
                    mode sparse;
                    version 2;
                }
            }
            mvpn;
        }
    }
}
```

The relevant sample configuration for Router CE2 follows.

Router CE2

```
interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 192.168.0.1/32 {
                    primary;
                }
            }
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.14/30;
            }
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.11.1/24;
            }
            family inet6 {
                address fe80::205:85ff:fe88:ccdb/64;
            }
```

```
        }
      }
    }
    protocols {
      ospf {
        area 0.0.0.0 {
          interface fe-0/1/0.0;
          interface lo0.0;
          interface fe-1/3/0.0;
        }
      }
      pim {
        rp {
          static {
            address 10.10.47.101 {
              version 2;
            }
          }
        }
        interface all {
          mode sparse;
          version 2;
        }
      }
    }
```

Related Documentation

- Understanding Multiprotocol BGP-Based Multicast VPNs: Next-Generation

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics, and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running Draft-Rosen multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across external and internal BGP (EIBGP) upstream paths when the PE router has both external BGP (EBGP) and internal BGP (IBGP) paths toward the source or rendezvous point (RP).

- Requirements on page 62
- Overview and Topology on page 63
- Configuration on page 66
- Verification on page 69

Requirements

This example requires the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- Junos OS Release 12.1 or later running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
   - OSPF
   - MPLS
   - LDP
   - PIM
   - BGP
3. Configure a multicast VPN.

Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across unequal EIBGP routes, if a PE router has EBGP and IBGP paths toward the source (or RP). In previous releases, only the active EBGP path was used to send the join messages. This feature is applicable to IPv4 C-PIM join messages.

During load balancing, if a PE router loses one or more EBGP paths toward the source (or RP), the C-PIM join messages that were previously using the EBGP path are moved to a multicast tunnel interface, and the reverse path forwarding (RPF) neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGP path toward the source (or RP), only the new join messages get load-balanced across EIBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

Though the primary goal for multipath PIM join load balancing is to utilize unequal EIBGP paths for multicast traffic, potential join loops can be avoided if a PE router chooses only the EBGP path when there are one or more join messages for different groups from a remote PE router. If the remote PE router's join message arrives after the PE router has already chosen IBGP as the upstream path, then the potential loops can be broken by changing the selected upstream path to EBGP.

NOTE: During a graceful Routing Engine switchover (GRES), the EIBGP path selection for C-PIM join messages can vary, because the upstream interface selection is performed again for the new Routing Engine based on the join messages it receives from the CE and PE neighbors. This can lead to disruption of multicast traffic depending on the number of join messages received and the load on the network at the time of the graceful restart. However, the nonstop active routing feature is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.

In this example, PE1 and PE2 are the upstream PE routers for which the multipath PIM join load-balancing feature is configured. Routers PE1 and PE2 have one EBGP path and one IBGP path each toward the source. The Source and Receiver attached to customer edge (CE) routers are Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The existing join-count-based load balancing is performed such that the algorithm first selects the least loaded C-PIM interface. If there is equal or no load on all the C-PIM interfaces, the join messages get distributed equally across the available upstream interfaces.

   In Figure 4 on page 66, if the PE1 router receives PIM join messages from the CE2 router, and if there is equal or no load on both the EBGP and IBGP paths toward the source, the join messages get load-balanced on the EIBGP paths.

2. If the selected least loaded interface is a multicast tunnel interface, then there can be a potential join loop if the downstream list of the customer join (C-join) message already contains the multicast tunnel interface. In such a case, the least loaded interface among EBGP paths is selected as the upstream interface for the C-join message.

   Assuming that the IBGP path is the least loaded, the PE1 router sends the join messages to PE2 using the IBGP path. If PIM join messages from the PE3 router arrive on PE1, then the downstream list of the C-join messages for PE3 already contains a multicast tunnel interface, which can lead to a potential join loop, because both the upstream and downstream interfaces are multicast tunnel interfaces. In this case, PE1 uses only the EBGP path to send the join messages.

3. If the selected least loaded interface is a multicast tunnel interface and the multicast tunnel interface is not present in the downstream list of the C-join messages, the loop prevention mechanism is not necessary. If any PE router has already advertised data multicast distribution tree (MDT) type, length, and values (TLVs), that PE router is selected as the upstream neighbor.

   When the PE1 router sends the join messages to PE2 using the least loaded IBGP path, and if PE3 sends its join messages to PE2, no join loop is created.

4. If no data MDT TLV corresponds to the C-join message, the least loaded neighbor on a multicast tunnel interface is selected as the upstream interface.

On PE routers that have only IBGP paths toward the source (or RP), such as PE3, PIM join load balancing is performed as follows:

1. The PE router only finds a multicast tunnel interface as the RPF interface, and load balancing is done across the C-PIM neighbors on a multicast tunnel interface.

   Router PE3 load-balances PIM join messages received from the CE4 router across the IBGP paths to the PE1 and PE2 routers.

2. If any PE router has already advertised data MDT TLVs corresponding to the C-join messages, that PE router is selected as the RPF neighbor.

For a particular C-multicast flow, at least one of the PE routers having EIBGP paths toward the source (or RP) must use only the EBGP path to avoid or break join loops. As a result of the loop avoidance mechanism, a PE router is constrained to choose among EIBGP paths when a multicast tunnel interface is already present in the downstream list.

In Figure 4 on page 66, assuming that the CE2 host is interested in receiving traffic from the Source and CE2 initiates multiple PIM join messages for different groups (Group 1 with group address 225.1.1.1, and Group 2 with group address 225.1.1.2), the join messages for both groups arrive on the PE1 router.

Router PE1 then equally distributes the join messages between the EIBGP paths toward the Source. Assuming that Group 1 join messages are sent to the CE1 router directly using the EBGP path, and Group 2 join messages are sent to the PE2 router using the IBGP path, PE1 and PE2 become the RPF neighbors for Group 1 and Group 2 join messages, respectively.

When the CE3 router initiates Group 1 and Group 2 PIM join messages, the join messages for both groups arrive on the PE2 router. Router PE2 then equally distributes the join messages between the EIBGP paths toward the Source. Since PE2 is the RPF neighbor for Group 2 join messages, it sends the Group 2 join messages directly to the CE1 router using the EBGP path. Group 1 join messages are sent to the PE1 router using the IBGP path.

However, if the CE4 router initiates multiple Group 1 and Group 2 PIM join messages, there is no control over how these join messages received on the PE3 router get distributed to reach the Source. The selection of the RPF neighbor by PE3 can affect PIM join load balancing on EIBGP paths.

- If PE3 sends Group 1 join messages to PE1 and Group 2 join messages to PE2, there is no change in RPF neighbor. As a result, no join loops are created.
- If PE3 sends Group 1 join messages to PE2 and Group 2 join messages to PE1, there is a change in the RPF neighbor for the different groups resulting in the creation of join loops. To avoid potential join loops, PE1 and PE2 do not consider IBGP paths to send the join messages received from the PE3 router. Instead, the join messages are sent directly to the CE1 router using only the EBGP path.

The loop avoidance mechanism in a Draft-Rosen MVPN has the following limitations:

- Because the timing of arrival of join messages on remote PE routers determines the distribution of join messages, the distribution could be sub-optimal in terms of join count.
- Because join loops cannot be avoided and can occur due to the timing of join messages, the subsequent RPF interface change leads to loss of multicast traffic. This can be avoided by implementing the PIM make-before-break feature.

  The PIM make-before-break feature is an approach to detect and break C-PIM join loops in a Draft-Rosen MVPN. The C-PIM join messages are sent to the new RPF neighbor after establishing the PIM neighbor relationship, but before updating the related multicast forwarding entry. Though the upstream RPF neighbor would have updated its multicast forwarding entry and started sending the multicast traffic downstream, the downstream router does not forward the multicast traffic (because of RPF check failure) until the multicast forwarding entry is updated with the new RPF neighbor. This helps to ensure that the multicast traffic is available on the new path before switching the RPF interface of the multicast forwarding entry.

Figure 4: PIM Join Load Balancing on Draft-Rosen MVPN

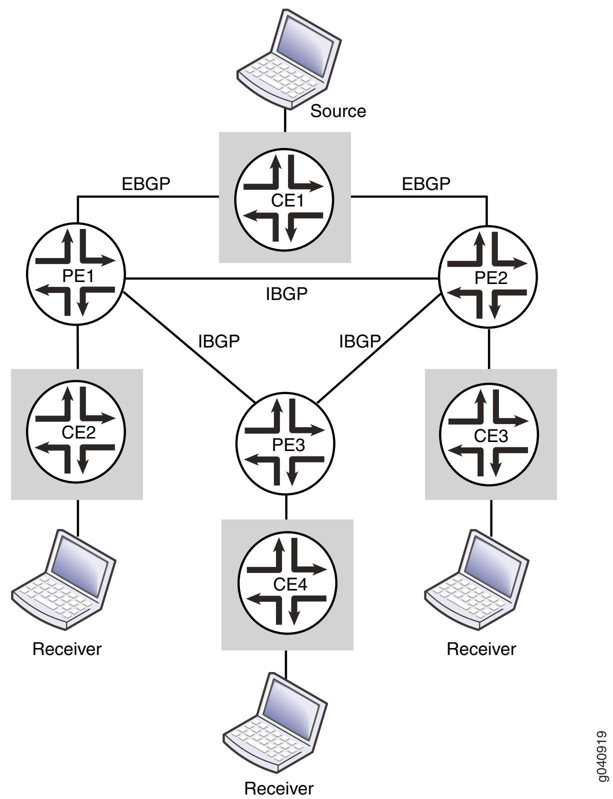

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

PE1

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-5/0/4.0
set routing-instances vpn1 interface ge-5/2/0.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 1:1
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
   equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 44.44.44.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 44.44.44.2 peer-as 3
```



```
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 11.11.11.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 11.11.11.2 peer-as 4
set routing-instances vpn1 protocols pim vpn-group-address 224.1.1.1
set routing-instances vpn1 protocols pim rp static address 10.255.8.168
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
```

PE2

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-2/0/3.0
set routing-instances vpn1 interface ge-4/0/5.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.90.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.90.10.2 peer-as 45
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 4
set routing-instances vpn1 protocols pim vpn-group-address 224.1.1.1
set routing-instances vpn1 protocols pim rp static address 10.255.8.168
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode. To configure the PE1 router:

NOTE: Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing and forwarding (VRF) instance.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-5/0/4.0
user@PE1# set interface ge-5/2/0.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set vrf-target target:1:1
```

2. Enable protocol-independent load balancing for the VRF instance.

```
[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 44.44.44.1
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 44.44.44.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 11.11.11.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 11.11.11.2 peer-as 4
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim vpn-group-address 224.1.1.1
user@PE1# set pim rp static address 10.255.8.168
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-5/0/4.0;
    interface ge-5/2/0.0;
    interface lo0.1;
    route-distinguisher 1:1;
    vrf-target target:1:1;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
    protocols {
      bgp {
        export direct;
        group bgp {
          type external;
          local-address 44.44.44.1;
          family inet {
            unicast;
          }
          neighbor 44.44.44.2 {
            peer-as 3;
          }
```



```
                }
                group bgp1 {
                    type external;
                    local-address 11.11.11.1;
                    family inet {
                        unicast;
                    }
                    neighbor 11.11.11.2 {
                        peer-as 4;
                    }
                }
            }
            pim {
                vpn-group-address 224.1.1.1;
                rp {
                    static {
                        address 10.255.8.168;
                    }
                }
                interface all;
                join-load-balance;
            }
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying PIM Join Load Balancing for Different Groups of Join Messages on page 69

Verifying PIM Join Load Balancing for Different Groups of Join Messages

Purpose Verify PIM join load balancing for the different groups of join messages received on the PE1 router.

Action From operational mode, run the **show pim join instance extensive** command.

user@PE1> show pim join instance extensive

```
Instance: PIM.vpn1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
    Source: *
    RP: 10.255.8.168
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-5/2/0.1
    Upstream neighbor: 10.10.10.2
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: ge-5/0/4.0
            10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 225.1.1.2
```

```
        Source: *
        RP: 10.255.8.168
        Flags: sparse,rptree,wildcard
        Upstream interface:  mt-5/0/10.32768
        Upstream neighbor: 19.19.19.19
        Upstream state: Join to RP
        Downstream neighbors:
            Interface: ge-5/0/4.0
                10.40.10.2 State: Join Flags: SRW Timeout: 207

    Group: 225.1.1.3
        Source: *
        RP: 10.255.8.168
        Flags: sparse,rptree,wildcard
        Upstream interface: ge-5/2/0.1
        Upstream neighbor: 10.10.10.2
        Upstream state: Join to RP
        Downstream neighbors:
            Interface: ge-5/0/4.0
                10.40.10.2 State: Join Flags: SRW Timeout: 207

    Group: 225.1.1.4
        Source: *
        RP: 10.255.8.168
        Flags: sparse,rptree,wildcard
        Upstream interface:  mt-5/0/10.32768
        Upstream neighbor: 19.19.19.19
        Upstream state: Join to RP
        Downstream neighbors:
            Interface: ge-5/0/4.0
                10.40.10.2 State: Join Flags: SRW Timeout: 207
```

Meaning The output shows how the PE1 router has load-balanced the C-PIM join messages for four different groups.

- For Group 1 (group address: 225.1.1.1) and Group 3 (group address: 225.1.1.3) join messages, the PE1 router has selected the EBGP path toward the CE1 router to send the join messages.
- For Group 2 (group address: 225.1.1.2) and Group 4 (group address: 225.1.1.4) join messages, the PE1 router has selected the IBGP path toward the PE2 router to send the join messages.

Related Documentation

- PIM Join Load Balancing on Multipath MVPN Routes Overview
- Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN on page 70

Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running next-generation multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across available internal BGP (IBGP) upstream paths when there is no external BGP (EBGP) path present, and across available EBGP upstream paths when external and internal BGP (EIBGP) paths are present toward the source or rendezvous point (RP).

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series, MX Series, or T Series routers.
- Junos OS Release 12.1 running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
   - OSPF
   - MPLS
   - LDP
   - PIM
   - BGP
3. Configure a multicast VPN.

Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across all available IBGP paths when there are only IBGP paths present, and across all available upstream EBGP paths when EIBGP paths are present toward the source (or RP). Unlike Draft-Rosen MVPN, next-generation MVPN does not utilize unequal EIBGP paths to send C-PIM join messages. This feature is applicable to IPv4 C-PIM join messages.

By default, only one active IBGP path is used to send the C-PIM join messages for a PE router having only IBGP paths toward the source (or RP). When there are EIBGP upstream paths present, only one active EBGP path is used to send the join messages.

In a next-generation MVPN, C-PIM join messages are translated into (or encoded as) BGP customer multicast (C-multicast) MVPN routes and advertised with the BGP MCAST-VPN address family toward the sender PE routers. A PE router originates a C-multicast MVPN route in response to receiving a C-PIM join message through its PE router to customer edge (CE) router interface. The two types of C-multicast MVPN routes are:

- Shared tree join route (C-*, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a shared tree C-PIM join message through its PE-CE router interface.
- Source tree join route (C-S, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a source tree C-PIM join message (C-S, C-G), or originated by the PE router that already has a shared tree join route and receives a source active autodiscovery route.

  The upstream path in a next-generation MVPN is selected using the Bytewise-XOR hash algorithm as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast, *Multicast in MPLS/BGP IP VPNs*. The hash algorithm is performed as follows:

  1. The PE routers in the candidate set are numbered from lower to higher IP address, starting from **0**.
  2. A bytewise exclusive-or of all the bytes is performed on the C-root (source) and the C-G (group) address.
  3. The result is taken modulo *n*, where *n* is the number of PE routers in the candidate set. The result is **N**.
  4. **N** represents the IP address of the upstream PE router as numbered in Step 1.

During load balancing, if a PE router with one or more upstream IBGP paths toward the source (or RP) discovers a new IBGP path toward the same source (or RP), the C-PIM join messages distributed among previously existing IBGP paths get redistributed due to the change in the candidate PE router set.

In this example, PE1, PE2, and PE3 are the PE routers that have the multipath PIM join load-balancing feature configured. Router PE1 has two EBGP paths and one IBGP upstream path, PE2 has one EBGP path and one IBGP upstream path, and PE3 has two IBGP upstream paths toward the Source. Router CE4 is the customer edge (CE) router attached to PE3. Source and Receiver are the Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The C-PIM join messages are sent using EBGP paths only. IBGP paths are not used to propagate the join messages.

In Figure 5 on page 74, the PE1 router distributes the join messages between the two EBGP paths to the CE1 router, and PE2 uses the EBGP path to CE1 to send the join messages.

2. If a PE router loses one or more EBGP paths toward the source (or RP), the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

   On discovering the first EBGP path, only new join messages get load-balanced across available EBGP paths, whereas the existing join messages on the multicast tunnel interface are not redistributed.

   If the EBGP path from the PE2 router to the CE1 router goes down, PE2 sends the join messages to PE1 using the IBGP path. When the EBGP path to CE1 is restored, only new join messages that arrive on PE2 use the restored EBGP path, whereas join messages already sent on the IBGP path are not redistributed.

On PE routers that have only IBGP paths toward the source (or RP), such as the PE3 router, PIM join load balancing is performed as follows:

1. The C-PIM join messages from CE routers get load-balanced only as BGP C-multicast data messages among IBGP paths.

   In Figure 5 on page 74, assuming that the CE4 host is interested in receiving traffic from the Source, and CE4 initiates source join messages for different groups (Group 1 [C-S,C-G1] and Group 2 [C-S,C-G2]), the source join messages arrive on the PE3 router.

   Router PE3 then uses the Bytewise-XOR hash algorithm to select the upstream PE router to send the C-multicast data for each group. The algorithm first numbers the upstream PE routers from lower to higher IP address starting from **0**.

   Assuming that Router PE1 router is numbered **0** and Router PE2 is **1**, and the hash result for Group 1 and Group 2 join messages is **0** and **1**, respectively, the PE3 router selects PE1 as the upstream PE router to send Group 1 join messages, and PE2 as the upstream PE router to send the Group 2 join messages to the Source.

2. The shared join messages for different groups [C-*,C-G] are also treated in a similar way to reach the destination.

Figure 5: PIM Join Load Balancing on Next-Generation MVPN

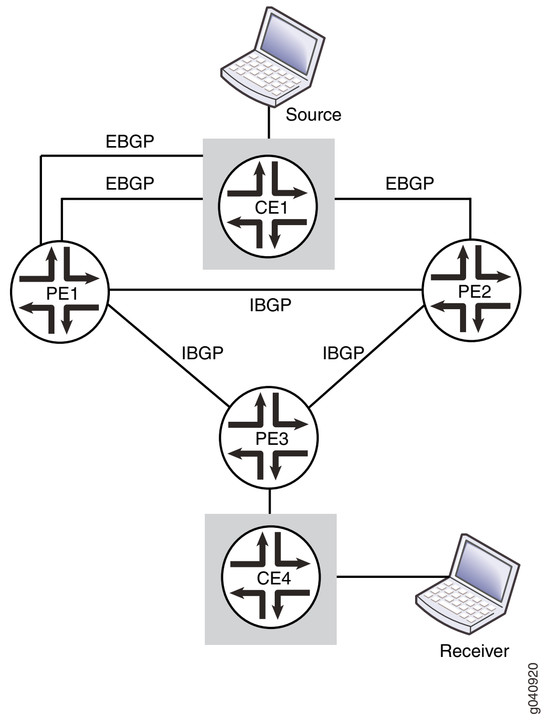

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

PE1

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-3/0/1.0
set routing-instances vpn1 interface ge-3/3/2.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 1:1
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
   default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
   equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.40.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.40.10.2 peer-as 3
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.10.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
```



```
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.10.10.2 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

PE2

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-1/0/9.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

PE3

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/8.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 3:3
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 routing-options autonomous-system 1
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.80.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.80.10.2 peer-as 2
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode. To configure the PE1 router:

NOTE: Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing forwarding (VRF) routing instance.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-3/0/1.0
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set provider-tunnel rsvp-te label-switched-path-template
  default-template
user@PE1# set vrf-target target:1:1
user@PE1# set vrf-table-label
```

2. Enable protocol-independent load balancing for the VRF instance.

```
[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 10.40.10.1
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 10.40.10.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 10.10.10.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 10.10.10.2 peer-as 3
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim rp static address 10.255.10.119
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

7. Configure the mode for C-PIM join messages to use rendezvous-point trees, and switch to the shortest-path tree after the source is known.


```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-mode rpt-spt
```

8. Configure the VRF instance to use the Bytewise-XOR hash algorithm.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-join-load-balance bytewise-xor-hash
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show routing-instances
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-3/0/1.0;
    interface ge-3/3/2.0;
    interface lo0.1;
    route-distinguisher 1:1;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          default-template;
        }
      }
    }
    vrf-target target:1:1;
    vrf-table-label;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
    protocols {
      bgp {
        export direct;
        group bgp {
          type external;
          local-address 10.40.10.1;
          family inet {
            unicast;
          }
          neighbor 10.40.10.2 {
            peer-as 3;
          }
        }
        group bgp1 {
          type external;
          local-address 10.10.10.1;
          family inet {
            unicast;
          }
          neighbor 10.10.10.2 {
            peer-as 3;
          }
```

```
                }
            }
            pim {
                rp {
                    static {
                        address 10.255.10.119;
                    }
                }
                interface all;
                join-load-balance;
            }
            mvpn {
                mvpn-mode {
                    rpt-spt;
                }
                mvpn-join-load-balance {
                    bytewise-xor-hash;
                }
            }
        }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages on page 78

Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages

Purpose Verify MVPN C-multicast route information for different groups of join messages received on the PE3 router.

Action From operational mode, run the **show mvpn c-multicast** command.

```
user@PE3> show mvpn c-multicast
MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : vpn1
  MVPN Mode : RPT-SPT
  C-mcast IPv4 (S:G)          Ptnl                           St
  0.0.0.0/0:225.1.1.1/32      RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
  4.4.4.2/32:225.1.1.1/32     RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
  0.0.0.0/0:225.1.1.2/32      RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
  4.4.4.2/32:225.1.1.2/32     RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
```


Meaning The output shows how the PE3 router has load-balanced the C-multicast data for the different groups.

- For source join messages (S,G):
  - 4.4.4.2/32:225.1.1.1/32 (S,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 4.4.4.2/32:225.1.1.2/32 (S,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).
- For shared join messages (*,G):
  - 0.0.0.0/0:225.1.1.1/32 (*,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 0.0.0.0/0:225.1.1.2/32 (*,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).

Related Documentation

- PIM Join Load Balancing on Multipath MVPN Routes Overview
- Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN on page 62

PART 3

Administration

CHAPTER 5

Multicast VPNs Reference

Multicast VPN Terminology

I

Inclusive tree A single multicast distribution tree in the backbone that carries all the multicast traffic from a specified set of one or more multicast VPNs. An inclusive tree that carries the traffic of more than one multicast VPN is an aggregate inclusive tree. An inclusive tree contains as its members all the PE routers that attach to the receiver sites of any of the multicast VPNs using the tree.

S

Selective tree A single multicast distribution tree in the backbone that carries traffic belonging only to a specified set of one or more multicast groups, from one or more multicast VPNs. An aggregate selective tree carries traffic for multicast groups that belong to different multicast VPNs. By default, traffic from most multicast groups could be carried by an inclusive tree, whereas traffic from high-bandwidth groups should be carried by a selective tree.

Supported Multicast VPN Standards

The Junos OS substantially supports the following Internet drafts, which define standards for multicast virtual private networks (VPNs).

- Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*

Related Documentation

- Supported Carrier-of-Carriers and Interprovider VPN Standards
- Supported Layer 2 Circuit Standards
- Supported Layer 2 VPN Standard
- Supported Layer 3 VPN Standards
- Supported VPLS Standards

- Supported MPLS Standards
- Supported BGP Standards
- Accessing Standards Documents on the Internet

CHAPTER 6

Summary of Multicast VPN Configuration Statements

create-new-ucast-tunnel

Syntax create-new-ucast-tunnel;

Hierarchy Level [edit routing-instances *routing-instance-name* provider-tunnel ingress-replication], [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* ingress-replication]

Release Information Statement introduced in Junos OS Release 10.4.

Description One of two modes for building unicast tunnels when ingress replication is configured for the provider tunnel. When this statement is configured, each time a new destination is added to the multicast distribution tree, a new unicast tunnel to the destination is created in the ingress replication tunnel. The new tunnel is deleted if the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs
- Configuring Routing Instances for an MBGP MVPN on page 17
- mpls-internet-multicast on page 97
- ingress-replication on page 95
- existing-unicast-tunnel on page 86

existing-unicast-tunnel

Syntax existing-unicast-tunnel;

Hierarchy Level [edit routing-instances *routing-instance-name* provider-tunnel ingress-replication],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* ingress-replication]

Release Information Statement introduced in Junos OS Release 10.4.

Description The default of two modes for selecting unicast tunnels when ingress replication is configured for the provider tunnel. When this is configured (or implied as default), each time a new destination is added to the multicast distribution tree, an existing unicast tunnel to the destination is used. If an existing tunnel is not available, the destination is not added. This is the only mode available when using LDP LSPs and ingress replication.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs
- Configuring Routing Instances for an MBGP MVPN on page 17
- mpls-internet-multicast on page 97
- ingress-replication on page 95
- existing-unicast-tunnel on page 86

existing-unicast-tunnel

Syntax existing-unicast-tunnel;

Hierarchy Level [edit routing-instances *routing-instance-name* provider-tunnel ingress-replication],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* ingress-replication]

Release Information Statement introduced in Junos OS Release 10.4.

Description The default of two modes for selecting unicast tunnels when ingress replication is configured for the provider tunnel. When this is configured (or implied as default), each time a new destination is added to the multicast distribution tree, an existing unicast tunnel to the destination is used. If an existing tunnel is not available, the destination is not added. This is the only mode available when using LDP LSPs and ingress replication.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs
- Configuring Routing Instances for an MBGP MVPN on page 17
- **mpls-internet-multicast on page 97**
- **ingress-replication on page 95**
- **existing-unicast-tunnel on page 86**

export-target

Syntax

```
export-target {
  target target-community;
  unicast;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target],
[edit routing-instances *routing-instance-name* protocols mvpn route-target]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN network layer reachability information (NLRI).

Options

target *target-community*—Specify the export target community.

unicast—Use the same target community as specified for unicast.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Export Target for an MBGP MVPN on page 24

family (VRF Advertisement)

Syntax

```
family {
  inet-mvpn;
  inet6-mvpn;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vrf-advertise-selective],
[edit routing-instances *routing-instance-name* vrf-advertise-selective],

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27
- inet-mvpn on page 92
- inet6-mvpn on page 94

group

Syntax

```
group address {
  source source-address {
    pim-ssm {
      group-range multicast-prefix;
    }
    ldp-p2mp;
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
    threshold-rate number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective],
[edit routing-instances *routing-instance-name* provider-tunnel selective]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP address for the multicast group configured for point-to-multipoint label-switched paths (LSPs) and PIM-SSM GRE selective provider tunnels.

Options

address—Specify the IP address for the multicast group. This address must be a valid multicast group address.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Multicast Group Address for an MBGP MVPN on page 32
- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27

group-range (MBGP MVPN Tunnel)

Syntax group-range *multicast-prefix*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-address* source *source-address* pim-ssm],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-address* wildcard-source pim-ssm],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *group-address* source *source-address* pim-ssm],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *group-address* wildcard-source pim-ssm],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm]

Release Information Statement introduced in Junos OS Release 10.1.

Description Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.

Options *multicast-prefix*—Multicast group address range to be used to create MBGP MVPN source-specific multicast selective PMSI tunnels.

Range: Any valid, nonreserved IPv4 multicast address range

Default: None

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27

import-target

Syntax

```
import-target {
  target {
    target-value;
    receiver target-value;
    sender target-value;
  }
  unicast {
    receiver;
    sender;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target],
[edit routing-instances *routing-instance-name* protocols mvpn route-target]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.

Options

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Import Target for an MBGP MVPN on page 25

inet-mvpn (BGP)

Syntax

```
inet-mvpn {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
    loops number;
    prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp family],
[edit protocols bgp family],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* family],
[edit protocols bgp group *group-name* family]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable the **inet-mvpn** address family in BGP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring NLRI Parameters for an MBGP MVPN on page 26

inet-mvpn (VRF Advertisement)

Syntax inet-mvpn;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vrf-advertise-selective family],
[edit routing-instances *routing-instance-name* vrf-advertise-selective family]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable IPv4 MVPN routes to be advertised from the VRF instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Limiting Routes to Be Advertised by an MVPN VRF Instance on page 22

inet6-mvpn (BGP)

Syntax

```
inet6-mvpn {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
    loops number
    prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp family],
[edit protocols bgp family],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* family],
[edit protocols bgp group *group-name* family]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Enable the **inet6-mvpn** address family in BGP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring NLRI Parameters for an MBGP MVPN on page 26
- BGP Configuration Guidelines chapter in the *Routing Protocols Configuration Guide*

inet6-mvpn (VRF Advertisement)

Syntax inet6-mvpn;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vrf-advertise-selective family],
[edit routing-instances *routing-instance-name* vrf-advertise-selective family],

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable IPv6 MVPN routes to be advertised from the VRF instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting Routes to Be Advertised by an MVPN VRF Instance on page 22

ingress-replication

Syntax

```
ingress-replication {
  (create-new-ucast-tunnel | existing-unicast-tunnel);
    label-switched-path-template {
    }
  }
```

Hierarchy Level [edit routing-instances *routing-instance-name* provider-tunnel], [edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]

Release Information Statement introduced in Junos OS Release 10.4.

Description A provider tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.

Options

existing-unicast-tunnel—An existing tunnel to the destination is used for ingress replication. If an existing tunnel is not available, the destination is not added. Default mode if no option is specified.

create-new-ucast-tunnel—When specified, a new unicast tunnel to the destination is created and used for ingress replication. The unicast tunnel is deleted later if the destination is no longer included in the multicast distribution tree.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs
- Configuring Routing Instances for an MBGP MVPN on page 17
- **create-new-ucast-tunnel on page 85**
- **existing-unicast-tunnel on page 86**
- **mpls-internet-multicast on page 97**

label-switched-path-template

Syntax

```
label-switched-path-template {
   (default-template | lsp-template-name);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te],
[edit routing-instances *routing-instance-name* provider-tunnel rsvp-te],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the LSP template used for the point-to-multipoint LSP. There is no default setting for the **label-switched-path-template** statement, so you must configure either the default template using the **default-template** option or you must specify the name of your preconfigured point-to-multipoint LSP template.

Options

default-template—Specify that the default template be used for the point-to-multipoint LSP.

lsp-template-name—Specify the name of an LSP template to be used for the point-to-multipoint LSP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 29
- Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 33

mpls-internet-multicast

Syntax mpls-internet-multicast;

Hierarchy Level [edit routing-instances *routing-instance-name* instance-type]
[edit protocols pim]

Release Information Statement introduced in Junos OS Release 11.1.

Description A nonforwarding routing instance type that supports Internet multicast over an MPLS network for the default master instance. No interfaces can be configured for it. Only one **mpls-internet-multicast** instance can be configured for each logical system.

The **mpls-internet-multicast** configuration statement is also explicitly required under PIM in the master instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs
- ingress-replication on page 95

mvpn

Syntax

```
mvpn {
  mvpn-mode (rpt-spt | spt-only);
  receiver-site;
  sender-site;
  route-target {
    export-target {
      target target-community;
      unicast;
    }
    import-target {
      target {
        target-value;
        receiver target-value;
        sender target-value;
      }
      unicast {
        receiver;
        sender;
      }
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable next-generation multicast VPNs in a routing instance.

Options **receiver-site**—Allow sites with multicast receivers.

sender-site—Allow sites with multicast senders.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Routing Instances for an MBGP MVPN on page 17

mvpn-mode

Syntax mvpn-mode (rpt-spt | spt-only);

Hierarchy Level [edit logical-systems *profile-name* routing-instances *instance-name* protocols mvpn],
[edit routing-instances *instance-name* protocols mvpn]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the mode for customer PIM (C-PIM) join messages.

The remaining statements are explained separately.

Default spt-only

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 21
- Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 19

pim-asm

Syntax

```
pim-asm {
  group-address address;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],
[edit routing-instances *routing-instance-name* provider-tunnel]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify a Protocol Independent Multicast (PIM) sparse mode provider tunnel for an MBGP MVPN or for a draft-rosen MVPN.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Provider Tunnels for an MBGP MVPN on page 27

pim-ssm (Selective Tunnel)

Syntax

```
pim-ssm {
  group-range multicast-prefix;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-address* source *source-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-address* wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *group-address* wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27

provider-tunnel

Syntax

```
provider-tunnel {
  ingress-replication {
    (create-new-ucast-tunnel | existing-unicast-tunnel);
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
    }
  ldp-p2mp;
  pim-asm {
    group-address address;
  }
  pim-ssm {
    group-address address;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
  }
  selective {
    group multicast--prefix/prefix-length {
      source ip--prefix/prefix-length {
        ldp-p2mp;
        (create-new-ucast-tunnel | existing-unicast-tunnel);
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
        }
        pim-ssm {
          group-range multicast-prefix;
        }
        rsvp-te {
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
          static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbps;
      }
      wildcard-source {
        pim-ssm {
          group-range multicast-prefix;
        }
        rsvp-te {
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
          static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbps;
      }
```

```
        }
        tunnel-limit number;
        wildcard-group-inet {
            wildcard-source {
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp lsp-name;
                }
                threshold-rate number;
            }
        }
        wildcard-group-inet6 {
            wildcard-source {
                pim-ssm {
                    group-range multicast-prefix;
                }
                rsvp-te {
                    label-switched-path-template {
                        (default-template | lsp-template-name);
                    }
                    static-lsp lsp-name;
                }
                threshold-rate number;
            }
        }
    }
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
The **selective** statement and substatements added in Junos OS Release 8.5.
The **ingress-replication** statement and substatements added in Junos OS Release 10.4.

Description
Configure virtual private LAN service (VPLS) flooding of unknown unicast, broadcast, and multicast traffic using point-to-multipoint LSPs. Also configure point-to-multipoint LSPs for MBGP MVPNs.

Options
The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Flooding Unknown Traffic Using Point-to-Multipoint LSPs
- Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 29
- Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode

route-target

Syntax

```
route-target {
  export-target {
    target target-community;
    unicast;
  }
  import-target {
    target {
      target-value;
      receiver target-value;
      sender target-value;
    }
    unicast {
      receiver;
      sender;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn],
[edit routing-instances *routing-instance-name* protocols mvpn]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.

Default The multicast VPN routing instance uses the import and export route targets configured for the Layer 3 VPN.

Options The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring VRF Route Targets for Routing Instances for an MBGP MVPN on page 23

rpt-spt

Syntax rpt-spt;

Hierarchy Level [edit logical-systems *profile-name* routing-instances *instance-name* protocols mvpn mvpn-mode],
[edit routing-instances *instance-name* protocols mvpn mvpn-mode]

Release Information Statement introduced in Junos OS Release 10.0.

Description Use rendezvous-point trees for customer PIM (C-PIM) join messages, and switch to the shortest-path tree after the source is known.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 21

rsvp-te

Syntax

```
rsvp-te {
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
  static-lsp lsp-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure the properties of the RSVP traffic-engineered point-to-multipoint LSP for MBGP MVPNs.

The remaining statements are explained separately.

NOTE: Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Selective Provider Tunnels for an MBGP MVPN on page 30

selective

Syntax

```
selective {
  group multicast-prefix/prefix-length {
    source ip-prefix/prefix-length {
      ingress-replication {
        (create-new-ucast-tunnel | existing-unicast-tunnel);
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
        }
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
      }
      threshold-rate kbps;
    }
    wildcard-source {
      ldp-p2mp;
        pim-ssm {
          group-range multicast-prefix;
        }
        rsvp-te {
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
        }
      static-lsp point-to-multipoint-lsp-name;
      }
      threshold-rate kbps;
    }
  }
  tunnel-limit number;
  wildcard-group-inet {
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
      }
      threshold-rate number;
    }
  }
```



```
  wildcard-group-inet6 {
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range multicast-prefix;
      }
      rsvp-te {
        label-switched-path-template {
          (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
      }
      threshold-rate number;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel],
[edit routing-instances *routing-instance-name* provider-tunnel]

Release Information

Statement introduced in Junos OS Release 8.5.
The **ingress-replication** statement and substatements added in Junos OS Release 10.4.

Description

Configure selective point-to-multipoint LSPs for an MBGP MVPN. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the MBGP MVPNs, helping to minimize flooding in the service provider's network.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Selective Provider Tunnels for an MBGP MVPN on page 30
- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27

source

Syntax

```
source source-address {
  ldp-p2mp;
  pim-ssm {
    group-range multicast-prefix;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
  }
  threshold-rate number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address*],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP address for the multicast source. This statement is a part of the point-to-multipoint LSP and PIM-SSM GRE selective provider tunnel configuration for MBGP MVPNs.

Options

source-address—IP address for the multicast source.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Multicast Source Address for an MBGP MVPN on page 32
- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27

spt-only

Syntax spt-only;

Hierarchy Level [edit logical-systems *profile-name* routing-instances *instance-name* protocols mvpn mvpn-mode],
[edit routing-instances *instance-name* protocols mvpn mvpn-mode]

Release Information Statement introduced in Junos OS Release 10.0.

Description Set the MVPN mode to learn about active multicast sources using multicast VPN source-active routes. This is the default mode.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 19

static-lsp

Syntax static-lsp *lsp-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te],
[edit routing-instances *routing-instance-name* provider-tunnel rsvp-te],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address* rsvp-te]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the name of the static point-to-multipoint LSP used for an MBGP MVPN. Use this statement to specify the static LSP for both inclusive and selective point-to-multipoint LSPs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Selective Provider Tunnels for an MBGP MVPN on page 30

target

Syntax

```
target target-value {
   receiver target-value;
   sender target-value;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target],
[edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]

Release Information
Statement introduced in Junos OS Release 8.4.

Description
Specify the target value when importing sender and receiver site routes.

Options
target-value—Specify the target value when importing sender and receiver site routes.

receiver—Specify the target community used when importing receiver site routes.

sender—Specify the target community used when importing sender site routes.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 25

threshold-rate

Syntax

threshold-rate *kbps*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-address* wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* source *source-address*]
[edit routing-instances *routing-instance-name* provider-tunnel selective group *address* wildcard-source]
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet wildcard-source],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6 wildcard-source]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the data threshold required before a new tunnel is created for a dynamic selective point-to-multipoint LSP. This statement is part of the configuration for point-to-multipoint LSPs for MBGP MVPNs and PIM-SSM GRE or RSVP-TE selective provider tunnels.

Options

number—Specify the data threshold required before a new tunnel is created.

Range: 0 through 1,000,000 kilobits per second. Specifying 0 is equivalent to not including the statement.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 33
- Configuring PIM-SSM GRE Selective Provider Tunnels on page 27
- Configuring Intra-AS Selective Provider Tunnels

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn],
[edit routing-instances *routing-instance-name* protocols mvpn]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Trace traffic flowing through an MBGP MVPN.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can specify any of the following flags:

- **all**—All multicast VPN tracing options
- **error**—Error conditions
- **general**—General events
- **nlri**—Multicast VPN advertisements received or sent by means of the BGP
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

- **topology**—Multicast VPN topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP

flag-modifier—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing flag
- **receive**—Trace received packets
- **send**—Trace sent packets

no-world-readable—Do not allow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—Allow any user to read the log file.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing MBGP MVPN Traffic and Operations on page 36

tunnel-limit

Syntax tunnel-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective],
[edit routing-instances *routing-instance-name* provider-tunnel selective]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a limit on the number of tunnels that can be created for a point-to-multipoint LSP.

Options *number*—Specify the tunnel limit.

Range: 0 through 1024

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 34

unicast

Syntax

```
unicast {
  receiver;
  sender;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target],
[edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify the same target community configured for unicast.

Options **receiver**—Specify the unicast target community used when importing receiver site routes.

sender—Specify the unicast target community used when importing sender site routes.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 26

vrf-advertise-selective

Syntax

```
vrf-advertise-selective {
  family {
    inet-mvpn;
    inet6-mvpn;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.

The options are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting Routes to Be Advertised by an MVPN VRF Instance on page 22

wildcard-group-inet

Syntax

```
wildcard-group-inet {
  wildcard-source {
    ldp-p2mp;
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
    threshold-rate number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective],
[edit routing-instances *routing-instance-name* provider-tunnel selective]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure a wildcard group matching any group IPv4 address.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **wildcard-group-inet6 on page 117**
- Example: Configuring Selective Provider Tunnels Using Wildcards on page 35
- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7
- Configuring a Selective Provider Tunnel Using Wildcards on page 34

wildcard-group-inet6

Syntax

```
wildcard-group-inet6 {
  wildcard-source {
    ldp-p2mp;
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
    threshold-rate number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective],
[edit routing-instances *routing-instance-name* provider-tunnel selective]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure a wildcard group matching any group IPv6 address.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **wildcard-group-inet on page 116**
- Example: Configuring Selective Provider Tunnels Using Wildcards on page 35
- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7
- Configuring a Selective Provider Tunnel Using Wildcards on page 34

wildcard-source

Syntax

```
wildcard-source {
ldp-p2mp;
  pim-ssm {
    group-range multicast-prefix;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group *group-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6],
[edit routing-instances *routing-instance-name* provider-tunnel selective group *group-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet],
[edit routing-instances *routing-instance-name* provider-tunnel selective wildcard-group-inet6]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure a selective provider tunnel for a shared tree using a wildcard source.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **wildcard-group-inet on page 116**
- **wildcard-group-inet6 on page 117**
- Example: Configuring Selective Provider Tunnels Using Wildcards on page 35
- Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 7
- Configuring a Selective Provider Tunnel Using Wildcards on page 34

PART 4

Index

Index

Symbols

A

B

C

D

E

F

G

I

L

M

N

P

R

S

T

U

V

W

