

Multicast VPN User Guide for EX9200 Switches

Published
2024-12-11

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Multicast VPN User Guide for EX9200 Switches
Copyright © 2024 Juniper Networks, Inc. All rights reserved.

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

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | vi

1

Overview

Understanding Multicast VPNs | 2

MBGP Multicast VPN Sites | 2

Multicast VPN Terminology | 3

Inclusive tree | 4

Selective tree | 4

Understanding Layer 3 VPNs | 5

Introduction to Configuring Layer 3 VPNs | 5

Layer 3 VPN Platform Support | 8

Supported Standards | 9

Supported Multicast VPN Standards | 9

2

Configuring Multicast on Layer 3 VPNs

Creating Next Generation MVPN VRF Import and Export Policies | 12

Limiting Routes to Be Advertised by an MVPN VRF Instance | 12

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN | 13

Signaling Provider Tunnels in Next Generation MVPNs | 18

PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs | 18

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN | 19

Requirements | 19

Overview and Topology | 20

Configuration | 23

Verification | 29

Example: Configuring MBGP Multicast VPNs | 31

Requirements | 31

Overview and Topology | 32

| Configuration | 33

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs | 55

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs | 56

| Requirements | 57

| Overview | 57

| Configuration | 57

| Verification | 68

Distributing Next Generation MVPN Routes | 71

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs | 71

Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs | 73

Configuring Internet Multicast Using Ingress Replication Provider Tunnels | 75

Example: Configuring PIM State Limits | 79

Controlling PIM Resources for Multicast VPNs Overview | 80

Example: Configuring PIM State Limits | 83

| Requirements | 83

| Overview | 83

| Configuration | 84

| Verification | 95

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN | 97

Configuring a Selective Provider Tunnel Using Wildcards | 103

Example: Configuring Selective Provider Tunnels Using Wildcards | 104

Configuring NLRI Parameters for an MBGP MVPN | 105

Configuring Routing Instances for an MBGP MVPN | 107

Configuring Point-to-Multipoint LSPs for an MBGP MVPN | 108

Configuring PIM Provider Tunnels for an MBGP MVPN | 115

Configuring PIM-SSM GRE Selective Provider Tunnels | 115

Configuring Draft Rosen VPNs | 117

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN | 117

| Requirements | 117

Overview and Topology | 118

Configuration | 122

Verification | 126

Configuring GRE Tunnel Interfaces for Layer 3 VPNs | 129

Configuring GRE Tunnels for Layer 3 VPNs | 129

3

Troubleshooting

Tracing Operations | 136

Tracing MBGP MVPN Traffic and Operations | 136

Knowledge Base | 139

4

Configuration Statements and Operational Commands

Operational-Mode Commands | 141

CLI Operational Mode Command Overview | 141

Junos CLI Reference Overview | 144

About This Guide

The Junos operating system (Junos OS) supports multicast VPN on the EX9200 switches. Use the topics on this page to configure MBGP MVPN.

1

PART

Overview

[Understanding Multicast VPNs | 2](#)

[Understanding Layer 3 VPNs | 5](#)

[Supported Standards | 9](#)

Understanding Multicast VPNs

IN THIS CHAPTER

- MBGP Multicast VPN Sites | 2
- Multicast VPN Terminology | 3

MBGP Multicast VPN Sites

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.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
11.1R2	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.

RELATED DOCUMENTATION

[Example: Allowing MBGP MVPN Remote Sources](#)

[Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN](#)

Multicast VPN Terminology

IN THIS SECTION

- [Inclusive tree](#) | 4
- [Selective tree](#) | 4

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.

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.

Understanding Layer 3 VPNs

IN THIS CHAPTER

- [Introduction to Configuring Layer 3 VPNs | 5](#)
- [Layer 3 VPN Platform Support | 8](#)

Introduction to Configuring Layer 3 VPNs

To configure Layer 3 virtual private network (VPN) functionality, you must enable VPN support on the provider edge (PE) router. You must also configure any provider (P) routers that service the VPN, and you must configure the customer edge (CE) routers so that their routes are distributed into the VPN.

To configure Layer 3 VPNs, you include the following statements:

```
description text;  
instance-type vrf;  
interface interface-name;  
protocols {  
    bgp {  
        group group-name {  
            peer-as as-number;  
            neighbor ip-address;  
        }  
        multihop ttd-value;  
    }  
    (ospf | ospf3) {  
        area area {  
            interface interface-name;  
        }  
        domain-id domain-id;  
        domain-vpn-tag number;  
        sham-link {  
            local address;
```

```

    }
    sham-link-remote address <metric number>;
  }
  rip {
    rip-configuration;
  }
}
route-distinguisher (as-number:id | ip-address:id);
router-id address;
routing-options {
  autonomous-system autonomous-system {
    independent-domain;
    loops number;
  }
  forwarding-table {
    export [ policy-names ];
  }
  interface-routes {
    rib-group group-name;
  }
  martians {
    destination-prefix match-type <allow>;
  }
  maximum-paths {
    path-limit;
    log-interval interval;
    log-only;
    threshold percentage;
  }
  maximum-prefixes {
    prefix-limit;
    log-interval interval;
    log-only;
    threshold percentage;
  }
  multipath {
    vpn-unequal-cost;
  }
  options {
    syslog (level level | upto level);
  }
  rib routing-table-name {
    martians {

```

```

        destination-prefix match-type <allow>;
    }
    multipath {
        vpn-unequal-cost;
    }
    static {
        defaults {
            static-options;
        }
        route destination-prefix {
            next-hop [next-hops];
            static-options;
        }
    }
}
static {
    defaults {
        static-options;
    }
    route destination-prefix {
        policy [ policy-names ];
        static-options;
    }
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-target (community | export community-name | import community-name);
vrf-table-label;

```

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*]



NOTE: The [edit logical-systems] hierarchy level is not applicable in ACX Series routers. The sham-link, sham-link-remote, and vrf-advertise-selective statements are not applicable in ACX Series routers.

For Layer 3 VPNs, only some of the statements in the [edit routing-instances] hierarchy are valid. For the full hierarchy, see [Junos OS Routing Protocols Library](#).

In addition to these statements, you must enable a signaling protocol, IBGP sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and P routers.

By default, Layer 3 VPNs are disabled.

Many of the configuration procedures for Layer 3 VPNs are common to all types of VPNs.

Layer 3 VPN Platform Support

Layer 3 VPNs are supported on most combinations of Juniper Networks routing and switching platforms and PICs capable of running the JUNOS Software.

MX Series routers configured to be in Ethernet services mode can support some of the Junos OS Layer 3 VPN features. For Layer 3 VPNs, Ethernet services mode supports configuring a loopback interface for a VPN routing and forwarding (VRF) instance. You can configure up to two VRF instances in Ethernet services mode. Each VRF instance can handle up to 10,000 routes. The `ping mpls l3vpn operational mode command` is also supported.

Supported Standards

IN THIS CHAPTER

- [Supported Multicast VPN Standards | 9](#)

Supported Multicast VPN Standards

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

- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPN*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discovery Routes*
- Internet draft draft-morin-l3vpn-mvpn-fast-failover-06.txt, *Multicast VPN Fast Upstream Failover*
- Internet draft draft-raggarwa-l3vpn-bgp-mvpn-extranet-08.txt, *Extranet in BGP Multicast VPN (MVPN)*
- RFC 7900, *Extranet Multicast in BGP/IP MPLS VPNs (partial support)*
- RFC 8534, *Explicit Tracking with Wildcard Routes in Multicast VPN (partial support)*
- RFC 9081, *Interoperation between Multicast Virtual Private Network (MVPN) and Multicast Source Directory Protocol (MSDP) Source-Active Routes*

RELATED DOCUMENTATION

Supported Carrier-of-Carriers and Interprovider VPN Standards

Supported VPWS Standards

Supported Layer 2 VPN Standards

Supported Layer 3 VPN Standards

Supported VPLS Standards

[Supported MPLS Standards](#)

[Supported Standards for BGP](#)

[Accessing Standards Documents on the Internet](#)

2

PART

Configuring Multicast on Layer 3 VPNs

[Creating Next Generation MVPN VRF Import and Export Policies | 12](#)

[Signaling Provider Tunnels in Next Generation MVPNs | 18](#)

[Distributing Next Generation MVPN Routes | 71](#)

[Configuring Draft Rosen VPNs | 117](#)

[Configuring GRE Tunnel Interfaces for Layer 3 VPNs | 129](#)

Creating Next Generation MVPN VRF Import and Export Policies

IN THIS CHAPTER

- [Limiting Routes to Be Advertised by an MVPN VRF Instance | 12](#)
- [Configuring VRF Route Targets for Routing Instances for an MBGP MVPN | 13](#)

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. However, if you configure the `vrf-advertise-selective` statement without any of its options, the router has the same behavior as if you configured the `no-vrf-advertise` statement. VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers.

RELATED DOCUMENTATION

family

inet-mvpn

inet6-mvpn

no-vrf-advertise

vrf-advertise-selective

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

IN THIS SECTION

- [Configuring the Export Target for an MBGP MVPN | 15](#)
- [Configuring the Import Target for an MBGP MVPN | 15](#)

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

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

To configure the import target community, include the target statement and specify the target community. The target community must be 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.

Signaling Provider Tunnels in Next Generation MVPNs

IN THIS CHAPTER

- [PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs | 18](#)
- [Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN | 19](#)
- [Example: Configuring MBGP Multicast VPNs | 31](#)
- [Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs | 55](#)
- [Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs | 56](#)

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 PIM Join Load Balancing on Next-Generation Multicast VPN

IN THIS SECTION

- [Requirements | 19](#)
- [Overview and Topology | 20](#)
- [Configuration | 23](#)
- [Verification | 29](#)

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 Byte-wise-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 bitwise 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 EBGP 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 1 on page 23](#), 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 1 on page 23](#), 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 Bitwise-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 1: PIM Join Load Balancing on Next-Generation MVPN



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 24](#)
- [Procedure | 25](#)
- [Results | 27](#)

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

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 bitwise-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 bitwise-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 bitwise-xor-hash

```

Procedure

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 {
        bitwise-xor-hash;
    }
}
}
```

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

Verification

IN THIS SECTION

- [Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages | 29](#)

Confirm that the configuration is working properly.

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>
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
0.0.0.0/0:203.0.113.1/24    RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
192.0.2.2/24:203.0.113.1/24  RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
0.0.0.0/0:203.0.113.2/24    RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
192.0.2.2/24:203.0.113.2/24  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):
 - 192.0.2.2/24:203.0.113.1/24 (S,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
 - 192.0.2.2/24:203.0.113.2/24 (S,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).
- For shared join messages (*,G):
 - 0.0.0.0/0:203.0.113.1/24 (*,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
 - 0.0.0.0/0:203.0.113.2/24 (*,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 MBGP Multicast VPNs

IN THIS SECTION

- [Requirements | 31](#)
- [Overview and Topology | 32](#)
- [Configuration | 33](#)

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

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 system 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

IN THIS SECTION

- Topology | 32

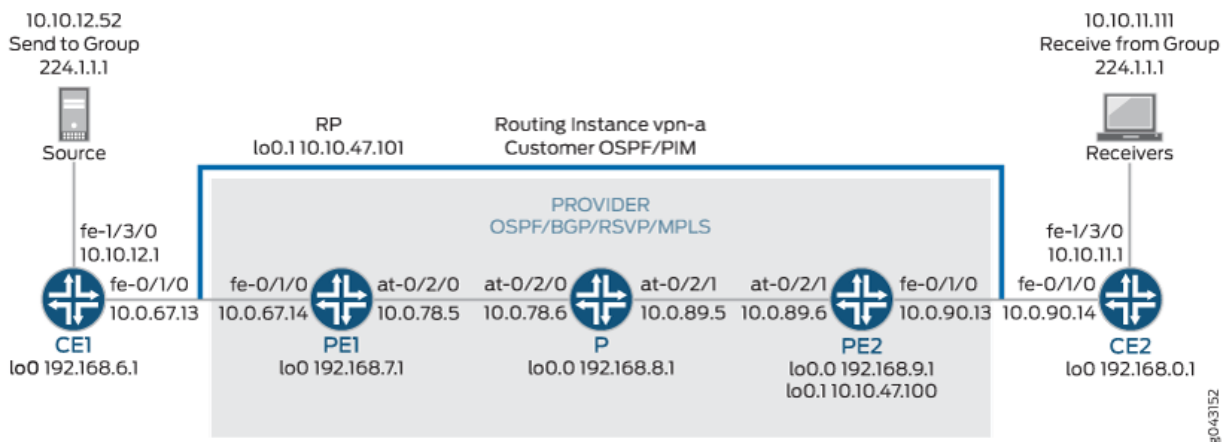
This example shows how to configure the following technologies:

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

Topology

The topology of the network is shown in [Figure 2 on page 32](#).

Figure 2: Multicast Over Layer 3 VPN Example Topology



8043152

Configuration

IN THIS SECTION

- [Configuring Interfaces | 34](#)
- [Configuring OSPF | 36](#)
- [Configuring BGP | 37](#)
- [Configuring RSVP | 38](#)
- [Configuring MPLS | 39](#)
- [Configuring the VRF Routing Instance | 40](#)
- [Configuring PIM | 42](#)
- [Configuring the Provider Tunnel | 43](#)
- [Configuring the Rendezvous Point | 43](#)
- [Results | 44](#)



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 2 on page 32](#), 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 [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 `100.0` address. The neighbor addresses are the PE routers' `100.0` 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 `100.0` address. The neighbor addresses are the `100.0` 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 routers, enable 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.

2. On the CE routers, enable 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.

```
user@PE1# set routing-instances vpn-a provider-tunnel rsvp-te label-switched-path-template default-template
```

Use the show configuration routing-instance vpn-a command to verify that the provider tunnel is configured to use the default LSP template.

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

```
user@PE2# set routing-instances vpn-a provider-tunnel rsvp-te label-switched-path-template default-template
```

Use the show configuration routing-instance vpn-a command to verify that the provider tunnel is configured to use the default LSP template.

Configuring the Rendezvous Point

Step-by-Step Procedure

1. Configure Router PE1 to be the rendezvous point. Specify the 10.10.47.1 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 RP.

2. On Router PE2, configure the static rendezvous point. Specify the 10.10.47.1 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 RP.

3. On the CE routers, configure the static rendezvous point. Specify the `100.1` 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 RP.

4. Use the `commit check` command to verify that the configuration can be successfully committed. If the configuration passes the check, commit the configuration.
5. Start the multicast sender device connected to CE1.
6. Start the multicast receiver device connected to CE2.
7. Verify that the receiver is receiving the multicast stream.
8. 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;
  }
}
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 {
    rsvp-te {
        label-switched-path-template {
            default-template;
        }
    }
}
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;
            }
        }
    }
}
```



```

    }
}

```

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;
    }
}
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 {
            rsvp-te {
                label-switched-path-template {
                    default-template;
                }
            }
        }
        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;
    }
}
}
}

```

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

In multiprotocol BGP (MBGP) multicast VPNs (MVPNs), VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).

Starting in Junos OS Release 12.3, you can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance. When the active VT interface fails, the secondary one takes over, and you can continue managing multicast traffic with no duplication.

Redundant VT interfaces are supported with RSVP point-to-multipoint provider tunnels as well as multicast LDP provider tunnels. This feature also works for extranets.

You can configure one of the VT interfaces to be the primary interface. If a VT interface is configured as the primary, it becomes the next hop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance. When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.

If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the next hop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.

To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
12.3	Starting in Junos OS Release 12.3, you can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance.

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

IN THIS SECTION

- [Requirements | 57](#)
- [Overview | 57](#)
- [Configuration | 57](#)
- [Verification | 68](#)

This example shows how to configure redundant virtual tunnel (VT) interfaces in multiprotocol BGP (MBGP) multicast VPNs (MVPNs). To configure, include multiple VT interfaces in the routing instance and, optionally, apply the `primary` statement to one of the VT interfaces.

Requirements

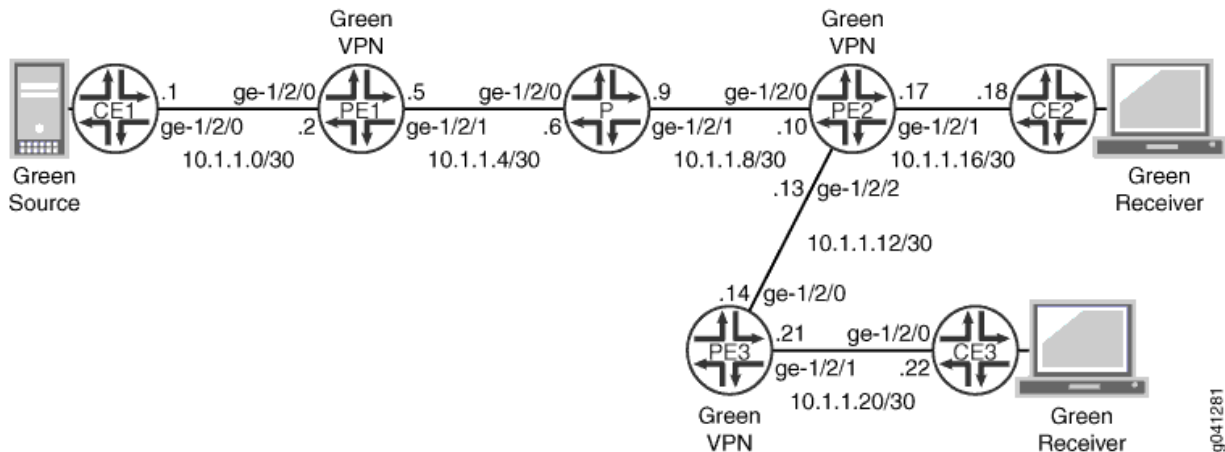
The routing device that has redundant VT interfaces configured must be running Junos OS Release 12.3 or later.

Overview

In this example, Device PE2 has redundant VT interfaces configured in a multicast LDP routing instance, and one of the VT interfaces is assigned to be the primary interface.

Figure 3 on page 57 shows the topology used in this example.

Figure 3: Multiple VT Interfaces in MBGP MVPN Topology



The following example shows the configuration for the customer edge (CE), provider (P), and provider edge (PE) devices in Figure 3 on page 57. The section "Step-by-Step Procedure" on page 62 describes the steps on Device PE2.

Configuration

IN THIS SECTION

- Procedure | 58

Procedure

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.

Device CE1

```
set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.1/24
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols pim rp static address 198.51.100.0
set protocols pim interface all
set routing-options router-id 192.0.2.1
```

Device CE2

```
set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.6/24
set protocols sap listen 192.168.0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols pim rp static address 198.51.100.0
set protocols pim interface all
set routing-options router-id 192.0.2.6
```

Device CE3

```
set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.7/24
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols pim rp static address 198.51.100.0
```



```

set protocols pim interface all
set routing-options router-id 192.0.2.7

```

Device P

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.0.2.3/24
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/1.0
set protocols ldp p2mp
set routing-options router-id 192.0.2.3

```

Device PE1

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.2/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 2 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.2/24
set interfaces lo0 unit 1 family inet address 198.51.100.0/24
set protocols mpls interface ge-1/2/1.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ldp interface ge-1/2/1.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp

```

```

set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-1/2/0.0
set routing-instances vpn-1 interface vt-1/2/0.2 multicast
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/0.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 1001

```

Device PE2

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/1/0 unit 0 family inet
set interfaces vt-1/2/1 unit 0 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.4/24
set interfaces lo0 unit 1 family inet address 203.0.113.4/24
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/2.0

```

```

set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/1/0.0 multicast
set routing-instances vpn-1 interface vt-1/1/0.0 primary
set routing-instances vpn-1 interface vt-1/2/1.0 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.4
set routing-options autonomous-system 1001

```

Device PE3

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 0 family inet address 192.0.2.5/24
set interfaces lo0 unit 1 family inet address 203.0.113.5/24
set protocols mpls interface ge-1/2/0.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept

```

```

set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set routing-instances vpn-1 protocols pim rp static address 198.51.100.0
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 1001

```

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 [CLI User Guide](#).

To configure redundant VT interfaces in an MBGP MVPN:

1. Configure the physical interfaces and loopback interfaces.

```

[edit interfaces]
user@PE2# set ge-1/2/0 unit 0 family inet address 10.1.1.10/30
user@PE2# set ge-1/2/0 unit 0 family mpls
user@PE2# set ge-1/2/2 unit 0 family inet address 10.1.1.13/30
user@PE2# set ge-1/2/2 unit 0 family mpls
user@PE2# set ge-1/2/1 unit 0 family inet address 10.1.1.17/30
user@PE2# set ge-1/2/1 unit 0 family mpls
user@PE2# set lo0 unit 0 family inet address 192.0.2.4/24
user@PE2# set lo0 unit 1 family inet address 203.0.113.4/24

```

2. Configure the VT interfaces.

Each VT interface is configurable under one routing instance.

```
[edit interfaces]
user@PE2# set vt-1/1/0 unit 0 family inet
user@PE2# set vt-1/2/1 unit 0 family inet
```

3. Configure MPLS on the physical interfaces.

```
[edit protocols mpls]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```

4. Configure BGP.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 192.0.2.4
user@PE2# set family inet-vpn any
user@PE2# set family inet-mvpn signaling
user@PE2# set neighbor 192.0.2.2
user@PE2# set neighbor 192.0.2.5
```

5. Configure an interior gateway protocol.

```
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface lo0.0 passive
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```

6. Configure LDP.

```
[edit protocols ldp]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
user@PE2# set p2mp
```

7. Configure the routing policy.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE2# set from protocol bgp
user@PE2# set then accept
```

8. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface ge-1/2/1.0
user@PE2# set interface lo0.1
user@PE2# set route-distinguisher 100:100
user@PE2# set vrf-target target:1:1
user@PE2# set protocols ospf export parent_vpn_routes
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.1 passive
user@PE2# set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
user@PE2# set protocols pim rp static address 198.51.100.0
user@PE2# set protocols pim interface ge-1/2/1.0 mode sparse
user@PE2# set protocols mvpn
```

9. Configure redundant VT interfaces in the routing instance.

Make vt-1/1/0.0 the primary interface.

```
[edit routing-instances vpn-1]
user@PE2# set interface vt-1/1/0.0 multicast primary
user@PE2# set interface vt-1/2/1.0 multicast
```

10. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@PE2# set router-id 192.0.2.4
user@PE2# set autonomous-system 1001
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@PE2# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/2 {
  unit 0 {
    family inet {
      address 10.1.1.13/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 0 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/1/0 {
  unit 0 {
    family inet;
  }
}
vt-1/2/1 {
  unit 0 {
    family inet;
  }
}
lo0 {
```

```
unit 0 {
    family inet {
        address 192.0.2.4/24;
    }
}
unit 1 {
    family inet {
        address 203.0.113.4/24;
    }
}
}
```

```
user@PE2# show protocols
mpls {
    interface ge-1/2/0.0;
    interface ge-1/2/2.0;
}
bgp {
    group ibgp {
        type internal;
        local-address 192.0.2.4;
        family inet-vpn {
            any;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 192.0.2.2;
        neighbor 192.0.2.5;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-1/2/0.0;
        interface ge-1/2/2.0;
    }
}
ldp {
```



```
interface ge-1/2/0.0;
interface ge-1/2/2.0;
p2mp;
}
```

```
user@PE2# show policy-options
policy-statement parent_vpn_routes {
  from protocol bgp;
  then accept;
}
```

```
user@PE2# show routing-instances
vpn-1 {
  instance-type vrf;
  interface vt-1/1/0.0 {
    multicast;
    primary;
  }
  interface vt-1/2/1.0 {
    multicast;
  }
  interface ge-1/2/1.0;
  interface lo0.1;
  route-distinguisher 100:100;
  vrf-target target:1:1;
  protocols {
    ospf {
      export parent_vpn_routes;
      area 0.0.0.0 {
        interface lo0.1 {
          passive;
        }
        interface ge-1/2/1.0;
      }
    }
    pim {
      rp {
        static {
          address 198.51.100.0;
        }
      }
    }
  }
}
```

```

    }
    interface ge-1/2/1.0 {
        mode sparse;
    }
}
mvpn;
}
}

```

```

user@PE2# show routing-options
router-id 192.0.2.4;
autonomous-system 1001;

```

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

Verification

IN THIS SECTION

- [Checking the LSP Route | 68](#)

Confirm that the configuration is working properly.



NOTE: The `show multicast route` extensive instance `instance-name` command also displays the VT interface in the multicast forwarding table when multicast traffic is transmitted across the VPN.

Checking the LSP Route

Purpose

Verify that the expected LT interface is assigned to the LDP-learned route.

Action

1. From operational mode, enter the **show route table mpls** command.

```

user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 02:09:36, metric 1
           Receive
1          *[MPLS/0] 02:09:36, metric 1
           Receive
2          *[MPLS/0] 02:09:36, metric 1
           Receive
13         *[MPLS/0] 02:09:36, metric 1
           Receive
299776     *[LDP/9] 02:09:14, metric 1
           > via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
           > via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
           > via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
           > via ge-1/2/2.0, Pop
299808     *[LDP/9] 02:09:04, metric 1
           > via ge-1/2/0.0, Swap 299808
299824     *[VPN/170] 02:08:56
           > via ge-1/2/1.0, Pop
299840     *[VPN/170] 02:08:56
           > via ge-1/2/1.0, Pop
299856     *[VPN/170] 02:08:56
           receive table vpn-1.inet.0, Pop
299872     *[LDP/9] 02:08:54, metric 1
           > via vt-1/1/0.0, Pop
           via ge-1/2/2.0, Swap 299872

```

2. From configuration mode, change the primary VT interface by removing the primary statement from the vt-1/1/0.0 interface and adding it to the vt-1/2/1.0 interface.

```

[edit routing-instances vpn-1]
user@PE2# delete interface vt-1/1/0.0 primary

```

```
user@PE2# set interface vt-1/2/1.0 primary
user@PE2# commit
```

3. From operational mode, enter the **show route table mpls** command.

```
user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 02:09:36, metric 1
           Receive
1          *[MPLS/0] 02:09:36, metric 1
           Receive
2          *[MPLS/0] 02:09:36, metric 1
           Receive
13         *[MPLS/0] 02:09:36, metric 1
           Receive
299776     *[LDP/9] 02:09:14, metric 1
           > via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
           > via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
           > via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
           > via ge-1/2/2.0, Pop
299808     *[LDP/9] 02:09:04, metric 1
           > via ge-1/2/0.0, Swap 299808
299824     *[VPN/170] 02:08:56
           > via ge-1/2/1.0, Pop
299840     *[VPN/170] 02:08:56
           > via ge-1/2/1.0, Pop
299856     *[VPN/170] 02:08:56
           receive table vpn-1.inet.0, Pop
299872     *[LDP/9] 02:08:54, metric 1
           > via vt-1/2/1.0, Pop
           via ge-1/2/2.0, Swap 299872
```

Meaning

With the original configuration, the output shows the vt-1/1/0.0 interface. If you change the primary interface to vt-1/2/1.0, the output shows the vt-1/2/1.0 interface.

Distributing Next Generation MVPN Routes

IN THIS CHAPTER

- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs | 71](#)
- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs | 73](#)
- [Configuring Internet Multicast Using Ingress Replication Provider Tunnels | 75](#)
- [Example: Configuring PIM State Limits | 79](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN | 97](#)
- [Configuring a Selective Provider Tunnel Using Wildcards | 103](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards | 104](#)
- [Configuring NLRI Parameters for an MBGP MVPN | 105](#)
- [Configuring Routing Instances for an MBGP MVPN | 107](#)
- [Configuring Point-to-Multipoint LSPs for an MBGP MVPN | 108](#)
- [Configuring PIM Provider Tunnels for an MBGP MVPN | 115](#)
- [Configuring PIM-SSM GRE Selective Provider Tunnels | 115](#)

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 7 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 primary 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 null routes 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 7 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.

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.

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 primary 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.

In Junos OS Release 15.1 and later, in RPT-SPT mode, PIM SSG Joins are created on the egress PE even if no directly-connected receivers are present.



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.

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.

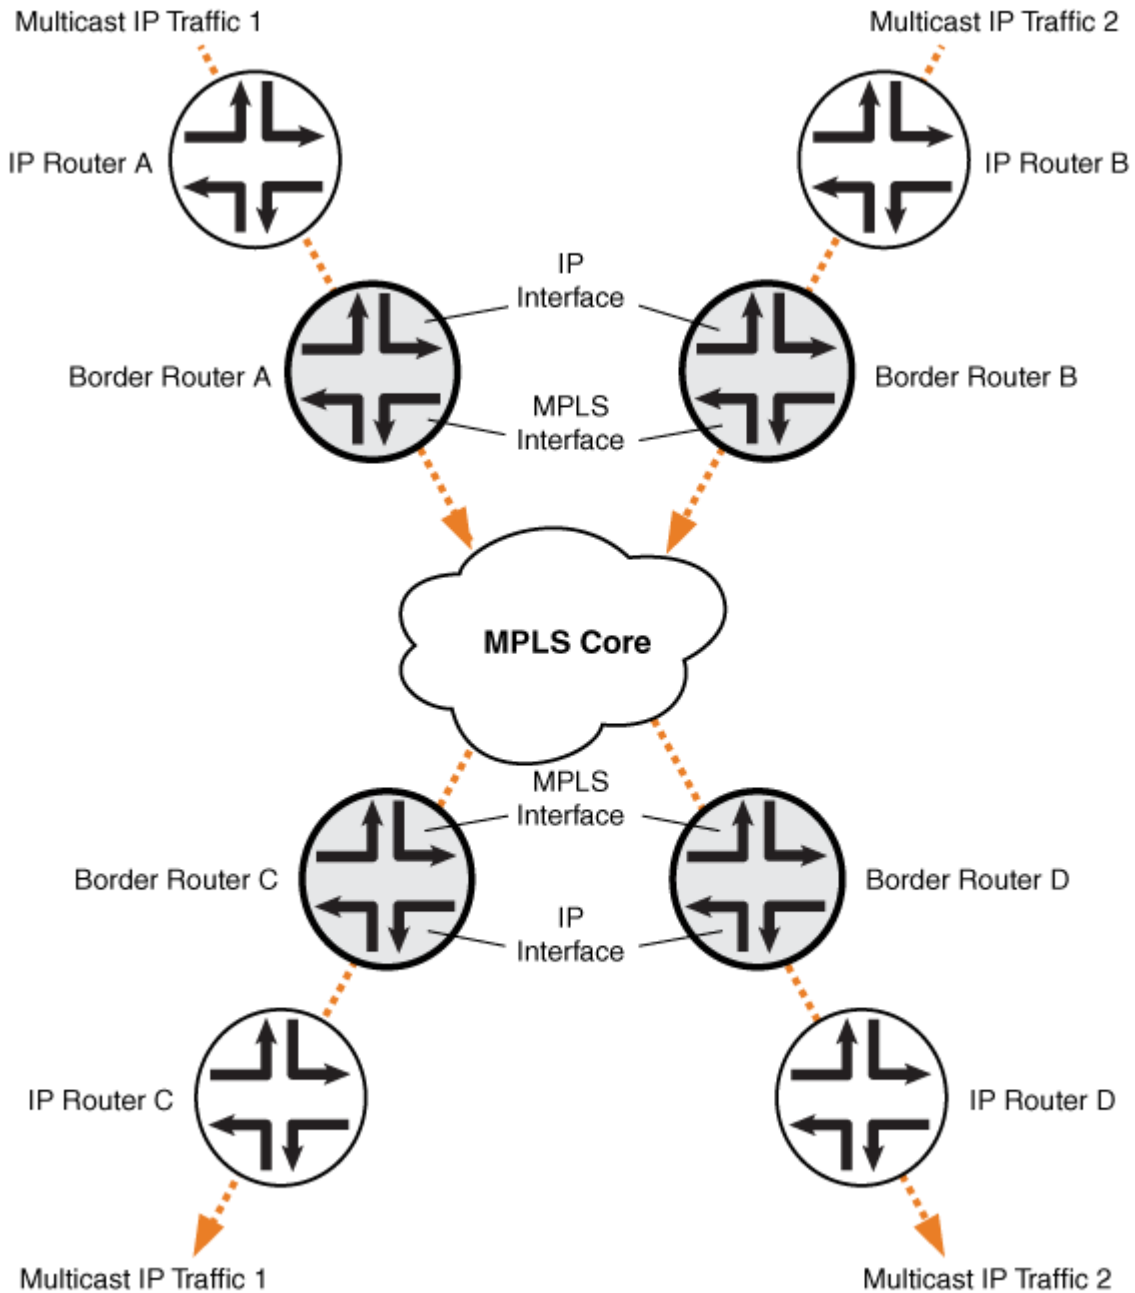
The topology consists of routers on the edge of the IP multicast domain that have a set of IP interfaces and a set of MPLS core-facing interfaces, see [Figure 4 on page 76](#). Internet multicast traffic is carried between the IP routers, through the MPLS cloud, using ingress replication tunnels for the data plane and a full-mesh IGBP session for the control plane.

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 a new destination needs to be added to the ingress replication provider tunnel, the resulting behavior differs depending on how the ingress replication provider tunnel is configured:

- `create-new-ucast-tunnel`—When this statement is configured, a new unicast tunnel to the destination is created, and is deleted when the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.
- `label-switched-path-template (Multicast)`—When this statement is configured, an LSP template is used for the for the point-to-multipoint LSP for ingress replication.

Figure 4: Internet Multicast Topology



g040632

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 203.0.113.1/24 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
192.168.195.145/32:203.0.113.1/24 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:1
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 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-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
192.168.195.145/32:203.0.113.1/24 INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
RM

```

RELATED DOCUMENTATION

[*create-new-ucast-tunnel*](#)

[*ingress-replication*](#)

[*mpls-internet-multicast*](#)

Example: Configuring PIM State Limits

IN THIS SECTION

- [Controlling PIM Resources for Multicast VPNs Overview | 80](#)
- [Example: Configuring PIM State Limits | 83](#)

Controlling PIM Resources for Multicast VPNs Overview

IN THIS SECTION

- [System Log Messages for PIM Resources | 81](#)

A service provider network must protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances. Misbehaving CE devices can potentially advertise a large number of multicast routes toward a provider edge (PE) device, thereby consuming memory on the PE device and using other system resources in the network that are reserved for routes belonging to other VPNs.

To protect against potential misbehaving CE devices and VRF routing instances for specific multicast VPNs (MVPNs), you can control the following Protocol Independent Multicast (PIM) resources:

- Limit the number of accepted PIM join messages for any-source groups (*,G) and source-specific groups (S,G).

Note how the device counts the PIM join messages:

- Each (*,G) counts as one group toward the limit.
- Each (S,G) counts as one group toward the limit.
- Limit the number of PIM register messages received for a specific VRF routing instance. Use this configuration if the device is configured as a rendezvous point (RP) or has the potential to become an RP. When a source in a multicast network becomes active, the source's designated router (DR) encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

Note how the device counts PIM register messages:

- Each unique (S,G) join received by the RP counts as one group toward the configured register messages limit.
- Periodic register messages sent by the DR for existing or already known (S,G) entries do not count toward the configured register messages limit.
- Register messages are accepted until either the PIM register limit or the PIM join limit (if configured) is exceeded. Once either limit is reached, any new requests are dropped.
- Limit the number of group-to-RP mappings allowed in a specific VRF routing instance. Use this configuration if the device is configured as an RP or has the potential to become an RP. This configuration can apply to devices configured for automatic RP announce and discovery (Auto-RP) or

as a PIM bootstrap router. Every multicast device within a PIM domain must be able to map a particular multicast group address to the same RP. Both Auto-RP and the bootstrap router functionality are the mechanisms used to learn the set of group-to-RP mappings. Auto-RP is typically used in a PIM dense-mode deployment, and the bootstrap router is typically used in a PIM sparse-mode deployment.



NOTE: The group-to-RP mappings limit does not apply to static RP or embedded RP configurations.

Some important things to note about how the device counts group-to-RP mappings:

- One group prefix mapped to five RPs counts as five group-to-RP mappings.
- Five distinct group prefixes mapped to one RP count as five group-to-RP mappings.

Once the configured limits are reached, no new PIM join messages, PIM register messages, or group-to-RP mappings are accepted unless one of the following occurs:

- You clear the current PIM join states by using the `clear pim join` command. If you use this command on an RP configured for PIM register message limits, the register limit count is also restarted because the PIM join messages are unknown by the RP.



NOTE: On the RP, you can also use the `clear pim register` command to clear all of the PIM registers. This command is useful if the current PIM register count is greater than the newly configured PIM register limit. After you clear the PIM registers, new PIM register messages are received up to the configured limit.

- The traffic responsible for the excess PIM join messages and PIM register messages stops and is no longer present.



CAUTION: Never restart any of the software processes unless instructed to do so by a customer support engineer.

You restart the PIM routing process on the device. This restart clears all of the configured limits but disrupts routing and therefore requires a maintenance window for the change.

System Log Messages for PIM Resources

You can optionally configure a system log warning threshold for each of the PIM resources. With this configuration, you can generate and review system log messages to detect if an excessive number of PIM join messages, PIM register messages, or group-to-RP mappings have been received on the device.

The system log warning thresholds are configured per PIM resource and are a percentage of the configured maximum limits of the PIM join messages, PIM register messages, and group-to-RP mappings. You can further specify a log interval for each configured PIM resource, which is the amount of time (in seconds) between the log messages.

The log messages convey when the configured limits have been exceeded, when the configured warning thresholds have been exceeded, and when the configured limits drop below the configured warning threshold. [Table 1 on page 82](#) describes the different types of PIM system messages that you might see depending on your system log warning and log interval configurations.

Table 1: PIM System Log Messages

System Log Message	Definition
RPD_PIM_SG_THRESHOLD_EXCEED	Records when the (S,G)/(*,G) routes exceed the configured warning threshold.
RPD_PIM_REG_THRESH_EXCEED	Records when the PIM registers exceed the configured warning threshold.
RPD_PIM_GRP_RP_MAP_THRES_EXCEED	Records when the group-to-RP mappings exceed the configured warning threshold.
RPD_PIM_SG_LIMIT_EXCEED	Records when the (S,G)/(*,G) routes exceed the configured limit, or when the configured log interval has been met and the routes exceed the configured limit.
RPD_PIM_REGISTER_LIMIT_EXCEED	Records when the PIM registers exceed the configured limit, or when the configured log interval has been met and the registers exceed the configured limit.
RPD_PIM_GRP_RP_MAP_LIMIT_EXCEED	Records when the group-to-RP mappings exceed the configured limit, or when the configured log interval has been met and the mapping exceeds the configured limit.
RPD_PIM_SG_LIMIT_BELOW	Records when the (S,G)/(*,G) routes drop below the configured limit and the configured log interval.
RPD_PIM_REGISTER_LIMIT_BELOW	Records when the PIM registers drop below the configured limit and the configured log interval.

Table 1: PIM System Log Messages (Continued)

System Log Message	Definition
RPD_PIM_GRP_RP_MAP_LIMIT_BELOW	Records when the group-to-RP mappings drop below the configured limit and the configured log interval.

Example: Configuring PIM State Limits**IN THIS SECTION**

- [Requirements | 83](#)
- [Overview | 83](#)
- [Configuration | 84](#)
- [Verification | 95](#)

This example shows how to set limits on the Protocol Independent Multicast (PIM) state information so that a service provider network can protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, a multiprotocol BGP-based multicast VPN (next-generation MBGP MVPN) is configured with limits on the PIM state resources.

The `sglimit maximum` statement sets a limit for the number of accepted (*,G) and (S,G) PIM join states received for the `vpn-1` routing instance.

The `rp register-limit maximum` statement configures a limit for the number of PIM register messages received for the `vpn-1` routing instance. You configure this statement on the rendezvous point (RP) or on all the devices that might become the RP.


The `group-rp-mapping maximum` statement configures a limit for the number of group-to-RP mappings allowed in the `vpn-1` routing instance.

For each configured PIM resource, the `threshold` statement sets a percentage of the maximum limit at which to start generating warning messages in the PIM log file.

For each configured PIM resource, the `log-interval` statement is an amount of time (in seconds) between system log message generation.

Figure 5 on page 84 shows the topology used in this example.

Figure 5: PIM State Limits Topology



"CLI Quick Configuration" on page 85 shows the configuration for all of the devices in Figure 5 on page 84. The section **Device PE1** below describes the steps for Device PE1.

Configuration

IN THIS SECTION

- Procedure | 85

*Procedure***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.

Device CE1

```
set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 1 family mpls
set interfaces lo0 unit 1 family inet address 192.0.2.1/24
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.1
```

Device PE1

```
set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30
set interfaces ge-1/2/0 unit 2 family mpls
set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 5 family mpls
set interfaces vt-1/2/0 unit 2 family inet
set interfaces lo0 unit 2 family inet address 192.0.2.2/24
set interfaces lo0 unit 102 family inet address 203.0.113.1/24
set protocols mpls interface ge-1/2/1.5
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
set protocols ldp interface ge-1/2/1.5
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
```

```

set routing-instances vpn-1 interface ge-1/2/0.2
set routing-instances vpn-1 interface vt-1/2/0.2
set routing-instances vpn-1 interface lo0.102
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2
set routing-instances vpn-1 protocols pim sglimit family inet maximum 100
set routing-instances vpn-1 protocols pim sglimit family inet threshold 70
set routing-instances vpn-1 protocols pim sglimit family inet log-interval 10
set routing-instances vpn-1 protocols pim rp register-limit family inet maximum 100
set routing-instances vpn-1 protocols pim rp register-limit family inet threshold 80
set routing-instances vpn-1 protocols pim rp register-limit family inet log-interval 10
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval 10
set routing-instances vpn-1 protocols pim rp static address 203.0.113.1
set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.2
set routing-options autonomous-system 1001

```

Device P

```

set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 9 family mpls
set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 192.0.2.3/24
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9

```

```

set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 192.0.2.3

```

Device PE2

```

set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 192.0.2.4/24
set interfaces lo0 unit 104 family inet address 203.0.113.4/24
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval 10
set routing-instances vpn-1 protocols pim rp static address 203.0.113.1
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn

```

```

set routing-options router-id 192.0.2.4
set routing-options autonomous-system 1001

```

Device PE3

```

set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 192.0.2.5/24
set interfaces lo0 unit 105 family inet address 203.0.113.5/24
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.0.2.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 192.0.2.2
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 203.0.113.1
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 192.0.2.5
set routing-options autonomous-system 1001

```

Device CE2

```

set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls
set interfaces lo0 unit 6 family inet address 192.0.2.6/24
set protocols sap listen 192.168.0.0
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.6

```

Device CE3

```

set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 192.0.2.7/24
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 203.0.113.1
set protocols pim interface all
set routing-options router-id 192.0.2.7

```

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 [CLI User Guide](#).

To configure PIM state limits:

1. Configure the network interfaces.

```

[edit interfaces]
user@PE1# set ge-1/2/0 unit 2 family inet address 10.1.1.2/30
user@PE1# set ge-1/2/0 unit 2 family mpls
user@PE1# set ge-1/2/1 unit 5 family inet address 10.1.1.5/30
user@PE1# set ge-1/2/1 unit 5 family mpls
user@PE1# set vt-1/2/0 unit 2 family inet

```

```

user@PE1# set lo0 unit 2 family inet address 192.0.2.2/24
user@PE1# set lo0 unit 102 family inet address 203.0.113.1/24

```

2. Configure MPLS on the core-facing interface.

```

[edit protocols mpls]
user@PE1# set interface ge-1/2/1.5

```

3. Configure internal BGP (IBGP) on the main router.

The IBGP neighbors are the other PE devices.

```

[edit protocols bgp group ibgp]
user@PE1# set type internal
user@PE1# set local-address 192.0.2.2
user@PE1# set family inet-vpn any
user@PE1# set family inet-mvpn signaling
user@PE1# set neighbor 192.0.2.4
user@PE1# set neighbor 192.0.2.5

```

4. Configure OSPF on the main router.

```

[edit protocols ospf area 0.0.0.0]
user@PE1# set interface lo0.2 passive
user@PE1# set interface ge-1/2/1.5

```

5. Configure a signaling protocol (RSVP or LDP) on the main router.

```

[edit protocols ldp]
user@PE1# set interface ge-1/2/1.5
user@PE1# set p2mp

```

6. Configure the BGP export policy.

```

[edit policy-options policy-statement parent_vpn_routes]
user@PE1# set from protocol bgp
user@PE1# set then accept

```


7. Configure the routing instance.

The customer-facing interfaces and the BGP export policy are referenced in the routing instance.

```
[edit routing-instances vpn-1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-1/2/0.2
user@PE1# set interface vt-1/2/0.2
user@PE1# set interface lo0.102
user@PE1# set route-distinguisher 100:100
user@PE1# set provider-tunnel ldp-p2mp
user@PE1# set vrf-target target:1:1
user@PE1# set protocols ospf export parent_vpn_routes
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.102 passive
user@PE1# set protocols ospf area 0.0.0.0 interface ge-1/2/0.2
user@PE1# set protocols pim rp static address 203.0.113.1
user@PE1# set protocols pim interface ge-1/2/0.2 mode sparse
user@PE1# set protocols mvpn
```

8. Configure the PIM state limits.

```
[edit routing-instances vpn-1 protocols pim]
user@PE1# set sglimit family inet maximum 100
user@PE1# set sglimit family inet threshold 70
user@PE1# set sglimit family inet log-interval 10
user@PE1# set rp register-limit family inet maximum 100
user@PE1# set rp register-limit family inet threshold 80
user@PE1# set rp register-limit family inet log-interval 10
user@PE1# set rp group-rp-mapping family inet maximum 100
user@PE1# set rp group-rp-mapping family inet threshold 80
user@PE1# set rp group-rp-mapping family inet log-interval 10
```

9. Configure the router ID and AS number.

```
[edit routing-options]
user@PE1# set router-id 192.0.2.2
user@PE1# set autonomous-system 1001
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show protocols`, `show policy-options`, `show routing-instances`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@PE1# show interfaces
ge-1/2/0 {
  unit 2 {
    family inet {
      address 10.1.1.2/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 2 {
    family inet;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.0.2.2/24;
    }
  }
  unit 102 {
    family inet {
      address 203.0.113.1/24;
    }
  }
}
```

```
}  
}
```

```
user@PE1# show protocols  
mpls {  
    interface ge-1/2/1.5;  
}  
bgp {  
    group ibgp {  
        type internal;  
        local-address 192.0.2.2;  
        family inet-vpn {  
            any;  
        }  
        family inet-mvpn {  
            signaling;  
        }  
        neighbor 192.0.2.4;  
        neighbor 192.0.2.5;  
    }  
}  
ospf {  
    area 0.0.0.0 {  
        interface lo0.2 {  
            passive;  
        }  
        interface ge-1/2/1.5;  
    }  
}  
ldp {  
    interface ge-1/2/1.5;  
    p2mp;  
}
```

```
user@PE1# show policy-options  
policy-statement parent_vpn_routes {  
    from protocol bgp;
```

```
    then accept;
}
```

```
user@PE1# show routing-instances
vpn-1 {
  instance-type vrf;
  interface ge-1/2/0.2;
  interface vt-1/2/0.2;
  interface lo0.102;
  route-distinguisher 100:100;
  provider-tunnel {
    ldp-p2mp;
  }
  vrf-target target:1:1;
  protocols {
    ospf {
      export parent_vpn_routes;
      area 0.0.0.0 {
        interface lo0.102 {
          passive;
        }
        interface ge-1/2/0.2;
      }
    }
    pim {
      sglimit {
        family inet {
          maximum 100;
          threshold 70;
          log-interval 10;
        }
      }
      rp {
        register-limit {
          family inet {
            maximum 100;
            threshold 80;
            log-interval 10;
          }
        }
        group-rp-mapping {
```

```
        family inet {
            maximum 100;
            threshold 80;
            log-interval 10;
        }
    }
    static {
        address 203.0.113.1;
    }
}
interface ge-1/2/0.2 {
    mode sparse;
}
}
mvpn;
}
```

```
user@PE1# show routing-options
router-id 192.0.2.2;
autonomous-system 1001;
```

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

Verification

IN THIS SECTION

- [Monitoring the PIM State Information | 95](#)

Confirm that the configuration is working properly.

Monitoring the PIM State Information

Purpose

Verify that the counters are set as expected and are not exceeding the configured limits.

Action

From operational mode, enter the `show pim statistics` command.

```

user@PE1> show pim statistics instance vpn-1
PIM Message type      Received      Sent  Rx errors
V2 Hello                393          390      0
...
V4 (S,G) Maximum                100
V4 (S,G) Accepted                0
V4 (S,G) Threshold              70
V4 (S,G) Log Interval           10
V4 (grp-prefix, RP) Maximum     100
V4 (grp-prefix, RP) Accepted     0
V4 (grp-prefix, RP) Threshold   80
V4 (grp-prefix, RP) Log Interval 10
V4 Register Maximum            100
V4 Register Accepted            0
V4 Register Threshold           80
V4 Register Log Interval        10

```

Meaning

The V4 (S,G) Maximum field shows the maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing instance. If this number is met, additional (S,G) entries are not accepted.

The V4 (S,G) Accepted field shows the number of accepted (S,G) IPv4 multicast routes.

The V4 (S,G) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).

The V4 (S,G) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 (grp-prefix, RP) Maximum field shows the maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.

The V4 (grp-prefix, RP) Accepted field shows the number of accepted group-to-RP IPv4 multicast mappings.

The V4 (grp-prefix, RP) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).

The V4 (grp-prefix, RP) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 Register Maximum field shows the maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.

The V4 Register Accepted field shows the number of accepted IPv4 PIM registers.

The V4 Register Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).

The V4 Register Log Interval field shows the time (in seconds) between consecutive log messages.

RELATED DOCUMENTATION

[Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces](#)

[Examples: Configuring the Multicast Forwarding Cache](#)

[Example: Configuring MSDP with Active Source Limits and Mesh Groups](#)

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

IN THIS SECTION

- [About S-PMSI | 98](#)
- [Scenarios for Using Wildcard S-PMSI | 99](#)
- [Types of Wildcard S-PMSI | 100](#)
- [Differences Between Wildcard S-PMSI and \(S,G\) S-PMSI | 100](#)
- [Wildcard \(*,*\) S-PMSI and PIM Dense Mode | 101](#)
- [Wildcard \(*,*\) S-PMSI and PIM-BSR | 101](#)
- [Wildcard Source and the 0.0.0.0/0 Source Prefix | 102](#)

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 6 on page 99](#) 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 6: 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 203.0.113.0/24 {
          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, 203.0.113.0/24) is bound to sptn13. A join message for (*, 203.0.113.0/24) is bound to sptn12.

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 <https://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 vjna 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 vjna 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 vjna provider-tunnel selective]
user@router# set group 203.0.113/24 wildcard-source
```

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

```
[edit routing-instances vjna provider-tunnel selective wildcard-group-inet wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective) 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 (Routing Instances Provider Tunnel Selective) label-switched-path-
template provider-tunnel2
```

6. Map the (*,203.0.113/24) join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective group 203.0.113/24 wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective) label-switched-path-
template provider-tunnel3
```

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 {
                sptn1;
              }
            }
          }
        }
      }
    }
    group 203.0.113.0/24 {
      wildcard-source {
        rsvp-te {
          label-switched-path-template {
            sptn2;
          }
        }
      }
    }
  }
}
```

```

    }
    }
}
source 10.1.1/24 {
    RSVP-TE {
        label-switched-path-template {
            sptn13;
        }
    }
}
}
}
}
}
}
}
}
}
}
}
}

```

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

- A customer multicast (10.1.1.1, 203.0.113.1) join message is bound to the sptn13 S-PMSI autodiscovery route.
- A customer multicast (10.2.1.1, 203.0.113.1) join message is bound to the sptn2 S-PMSI autodiscovery route.
- A customer multicast (10.1.1.1, 203.1.113.1) join message is bound to the sptn1 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.

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 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 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 Point-to-Multipoint LSPs for an MBGP MVPN

IN THIS SECTION

- [Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN | 109](#)
- [Configuring Selective Provider Tunnels for an MBGP MVPN | 110](#)

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 (Multicast) {
    (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 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 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.

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 224.0.113.1/24 customer multicast group address, the 10.2.2.2/32 customer source address, and a virtual routing instance named `green`. Since we are using addresses outside the reserved SSM address range of 232.0.0.0/8, we must also include the `group-range` address with the `pim-ssm` option.

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 224.0.113.1/24 source
10.2.2.2/32 pim-ssm group-range 224.0.113.1/24
```

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 224.0.113.1/24 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 | 3](#)

pim-ssm

group-range

threshold-rate

Configuring Draft Rosen VPNs

IN THIS CHAPTER

- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN | 117](#)

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

IN THIS SECTION

- [Requirements | 117](#)
- [Overview and Topology | 118](#)
- [Configuration | 122](#)
- [Verification | 126](#)

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

This example requires the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, 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 EBG and IBGP paths toward the source (or RP). In previous releases, only the active EBG 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 EBG paths toward the source (or RP), the C-PIM join messages that were previously using the EBG 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 EBG 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 EBG 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 EBG.



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 7 on page 121](#), 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 7 on page 121](#), 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 203.0.113.1, and Group 2 with group address 203.0.113.2), 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 7: PIM Join Load Balancing on Draft-Rosen MVPN



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 122](#)
- [Procedure | 123](#)

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 192.0.2.4
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 192.0.2.5 peer-as 3
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 192.0.2.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 192.0.2.2 peer-as 4
set routing-instances vpn1 protocols pim group-address 198.51.100.1
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 group-address 198.51.100.1
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

```

Procedure

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 192.0.2.4
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 192.0.2.5 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 192.0.2.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 192.0.2.2 peer-as 4
```

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

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim group-address 198.51.100.1
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 192.0.2.4;
          family inet {
            unicast;
          }
          neighbor 192.0.2.5 {
            peer-as 3;
          }
        }
        group bgp1 {
          type external;
          local-address 192.0.2.1;
          family inet {
```

```
        unicast;
    }
    neighbor 192.0.2.2 {
        peer-as 4;
    }
}
pim {
    group-address 198.51.100.1;
    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

IN THIS SECTION

- [Verifying PIM Join Load Balancing for Different Groups of Join Messages](#) | 126

Confirm that the configuration is working properly.

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: 203.0.113.1
  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: 203.0.113.2
  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: 203.0.113.3
  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: 203.0.113.4
```

```
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: 203.0.113.1) and Group 3 (group address: 203.0.113.3) join messages, the PE1 router has selected the EBGP path toward the CE1 router to send the join messages.
- For Group 2 (group address: 203.0.113.2) and Group 4 (group address: 203.0.113.4) 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

Configuring GRE Tunnel Interfaces for Layer 3 VPNs

IN THIS CHAPTER

- [Configuring GRE Tunnels for Layer 3 VPNs | 129](#)

Configuring GRE Tunnels for Layer 3 VPNs

IN THIS SECTION

- [Configuring GRE Tunnels Manually Between PE and CE Routers | 130](#)
- [Configuring GRE Tunnels Dynamically | 132](#)

Junos OS allows you to configure a generic routing encapsulation (GRE) tunnel between the PE and CE routers for a Layer 3 VPN. The GRE tunnel can have one or more hops. You can configure the tunnel from the PE router to a local CE router (as shown in [Figure 8 on page 130](#)) or to a remote CE router (as shown in [Figure 9 on page 130](#)).

Figure 8: GRE Tunnel Configured Between the Local CE Router and the PE Router

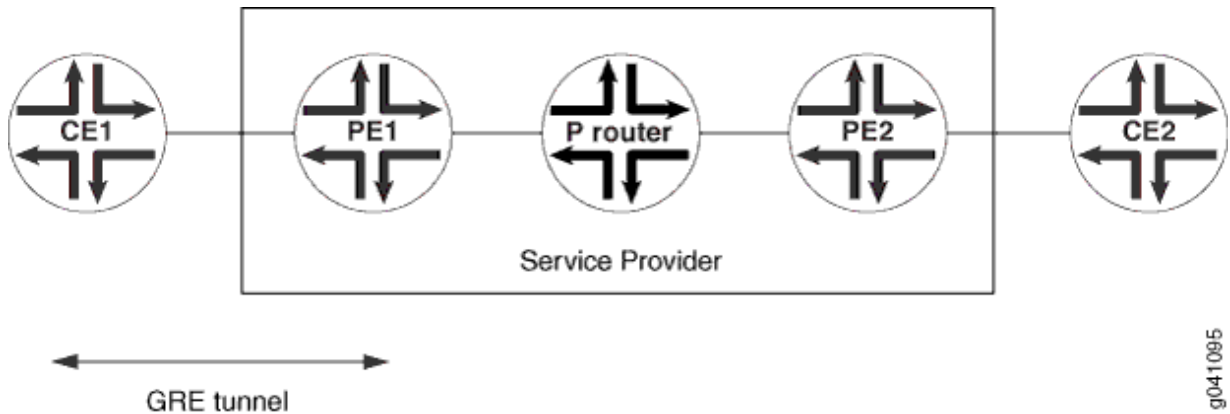
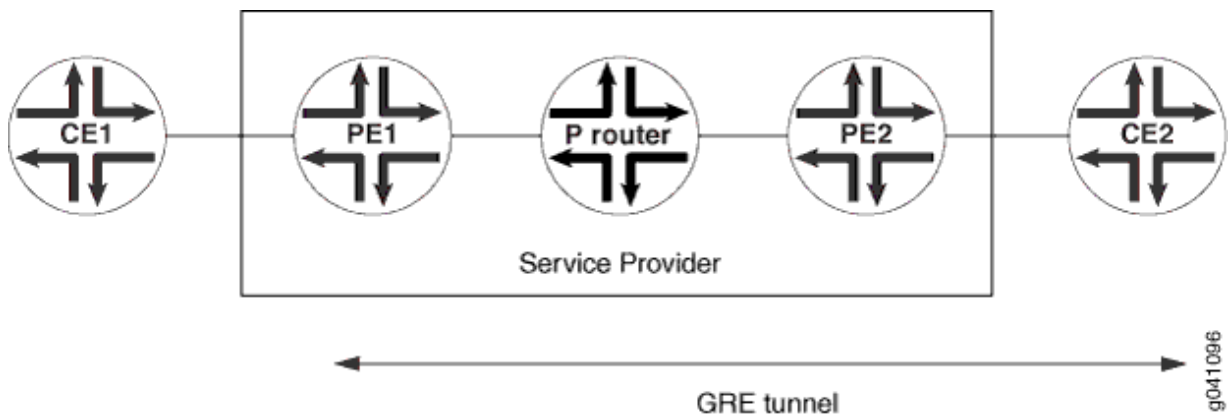


Figure 9: GRE Tunnel Configured Between the Remote CE Router and the PE Router



For more information about how to configure tunnel interfaces, see the [Junos OS Services Interfaces Library for Routing Devices](#).

You can configure the GRE tunnels manually or configure the Junos OS to instantiate GRE tunnels dynamically.

The following sections describe how to configure GRE tunnels manually and dynamically:

Configuring GRE Tunnels Manually Between PE and CE Routers

You can manually configure a GRE tunnel between a PE router and either a local CE router or a remote CE router for a Layer 3 VPN as explained in the following sections:

Configuring the GRE Tunnel Interface on the PE Router

You configure the GRE tunnel as a logical interface on the PE router. To configure the GRE tunnel interface, include the unit statement:

```
unit logical-unit-number {
    tunnel {
        source source-address;
        destination destination-address;
        routing-instance {
            destination routing-instance-name;
        }
    }
    family inet {
        address address;
    }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

As part of the GRE tunnel interface configuration, you need to include the following statements:

- *source source-address*—Specify the source or origin of the GRE tunnel, typically the PE router.
- *destination destination-address*—Specify the destination or end point of the GRE tunnel. The destination can be a Provider router, the local CE router, or the remote CE router.

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. If the tunnel destination address is not in inet.0, you need to specify which routing table to search for the tunnel destination address by configuring the *routing-instance* statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

- *destination routing-instance-name*—Specify the name of the routing instance when configuring the GRE tunnel interface on the PE router.

To complete the GRE tunnel interface configuration, include the *interface* statement for the GRE interface under the appropriate routing instance:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

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

Configuring the GRE Tunnel Interface on the CE Router

You can configure either the local or the remote CE router to act as the endpoint for the GRE tunnel.

To configure the GRE tunnel interface on the CE router, include the unit statement:

```
unit logical-unit-number {
  tunnel {
    source address;
    destination address;
  }
  family inet {
    address address;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

Configuring GRE Tunnels Dynamically

When the router receives a VPN route to a BGP next hop address, but no MPLS path is available, a GRE tunnel can be dynamically generated to carry the VPN traffic across the BGP network. The GRE tunnel is generated and then its routing information is copied into the inet.3 routing table. IPv4 routes are the only type of routes supported for dynamic GRE tunnels. Also, the routing platform must have a tunnel PIC.



NOTE: When configuring a dynamic GRE tunnel to a remote CE router, do not configure OSPF over the tunnel interface. It creates a routing loop forcing the router to take the GRE tunnel down. The router attempts to reestablish the GRE tunnel, but will be forced to take it down again when OSPF becomes active on the tunnel interface and discovers a route to the tunnel endpoint. This is not an issue when configuring static GRE tunnels to a remote CE router.

To generate GRE tunnels dynamically, include the `dynamic-tunnels` statement:

```
dynamic-tunnels tunnel-name {
    destination-networks prefix;
    source-address address;
}
```

You can include this statement at the following hierarchy levels:

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

Specify the IPv4 prefix range (for example, 10/8 or 11.1/16) for the destination network by including the `destination-networks` statement. Only tunnels within the specified IPv4 prefix range are allowed to be initiated.

```
destination-networks prefix;
```

You can include this statement at the following hierarchy levels:

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

Specify the source address for the GRE tunnels by including the `source-address` statement. The source address specifies the address used as the source for the local tunnel endpoint. This could be any local address on the router (typically the router ID or the loopback address).

```
source-address address;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit routing-instances *routing-instance-name* routing-options]

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

RELATED DOCUMENTATION

| [Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks](#)

3

PART

Troubleshooting

[Tracing Operations](#) | 136

[Knowledge Base](#) | 139

Tracing Operations

IN THIS CHAPTER

- [Tracing MBGP MVPN Traffic and Operations | 136](#)

Tracing MBGP MVPN Traffic and Operations

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

1. Specify the name of one or more MVPN trace files using the `file` option for the `traceoptions` at the `[edit routing-instances routing-instance-name protocols mvpn]` hierarchy level or at the `[edit protocols mvpn]` hierarchy level:

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

The `file` option includes the following sub-options:

- `filename`—Specify the name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.
- `files number`—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum `size`, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the specified maximum `number` of trace files specified is reached. Then the oldest trace file is overwritten.
- `size size`—(Optional) Maximum size of each trace file. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
- `world-readable | no-world-readable`—(Optional) Enable unrestricted file access or restrict file access to the user who created the file.

2. Specify the `flag` option for the `traceoptions` statement:

```
traceoptions {
  flag flag <flag-modifier> <disable>;
}
```

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

- `all`—All multicast VPN tracing options
- `cmcast-join`—Multicast VPN C-multicast join routes
- `error`—Error conditions
- `general`—General events
- `inter-as-ad`—Multicast VPN inter-AS automatic discovery routes
- `intra-as-ad`—Multicast VPN intra-AS automatic discovery routes
- `leaf-ad`—Multicast VPN leaf automatic discovery routes
- `mdt-safi-ad`—Multicast VPN *MDT SAFI* automatic discovery routes
- `nlri`—Multicast VPN advertisements received or sent by means of BGP
- `normal`—Normal events
- `policy`—Policy processing
- `route`—Routing information
- `source-active`—Multicast VPN source active routes
- `spmsi-ad`—Multicast VPN SPMSI auto discovery active routes
- `state`—State transitions
- `task`—Routing protocol task processing
- `timer`—Routing protocol timer processing
- `tunnel`—Provider tunnel events
- `umh`—Upstream multicast hop (UMH) events

RELATED DOCUMENTATION

| *traceoptions*

CHAPTER 10

Knowledge Base

4

PART

Configuration Statements and Operational Commands

[Operational-Mode Commands](#) | 141

[Junos CLI Reference Overview](#) | 144

Operational-Mode Commands

IN THIS SECTION

- [CLI Operational Mode Command Overview | 141](#)

CLI Operational Mode Command Overview

IN THIS SECTION

- [CLI Operational Mode Command Categories | 141](#)
- [Commonly Used Operational Mode Commands | 142](#)

You (the network administrator) can control all network operations using the Junos OS CLI operational mode commands described in this topic.

CLI Operational Mode Command Categories

CLI operational mode commands fall into the following broad categories:

- Operational mode commands for monitoring and troubleshooting—The following commands perform functions related to information and statistics about the software and to test network connectivity.
 - `clear`—Clear statistics and protocol database information.
 - `file`—Perform file operations.
 - `mtrace`—Trace a multicast path from source to receiver.
 - `monitor`—Perform real-time debugging of various software components, including the routing protocols and interfaces.
 - `ping`—Determine the reachability of a remote network host.

- `show`—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
- `test`—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
- `traceroute`—Trace the route to a remote network host.
- Commands for restarting software processes—The commands in the restart hierarchy restart the various system processes, including the routing protocol, interface, and SNMP.
- A command—`request`—Perform system-level operations, including stopping and rebooting the router or switch and loading operating system images.
- A command—`start`—Exit the CLI and start a UNIX shell.
- A command—`configure`—Enter configuration mode, which provides a series of commands that configure the system, including the routing protocols, interfaces, network management, and user access.

For more information about the CLI operational mode commands, see the [CLI Explorer](#). Alternatively, you can enter `?` at the operational mode command prompt to view a list of available commands.

Commonly Used Operational Mode Commands

The following table lists some operational commands you may find useful for monitoring router or switch operation.

Table 2: Commonly Used Operational Mode Commands

Items to Check	Description	Command
Software version	Versions of software running on the router or switch	<code>show version</code>
Log files	Contents of the log files	<code>monitor</code>
	Log files and their contents and recent user logins	<code>show log</code>
Remote systems	Host reachability and network connectivity	<code>ping</code>
	The route to a network system	<code>traceroute</code>

Table 2: Commonly Used Operational Mode Commands (Continued)

Items to Check	Description	Command
Configuration	Current system configuration	show configuration
File manipulation	List of files and directories on the router or switch	file list
	Contents of a file	file show
Interface information	Detailed information about interfaces	show interfaces
Chassis	Chassis alarm status	show chassis alarms
	Information currently on craft display	show chassis craft-interface
	Router or switch environment information	show chassis environment
	Hardware inventory	show chassis hardware
Routing table information	Information about entries in the routing tables	show route
Forwarding table information	Information about data in the kernel's forwarding table	show route forwarding-table

RELATED DOCUMENTATION

| [Understanding Logical Systems for Routers and Switches](#)

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)