JUNIPER NETWORKS

# Technology Overview

## Understanding Junos OS Next-Generation Multicast VPNs

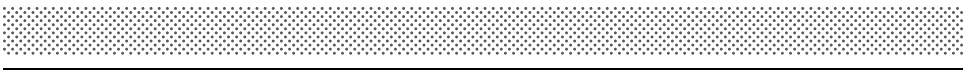

Published: 2014-01-10

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



*Technology Overview  Understanding Junos OS Next-Generation Multicast VPNs*
NCE0090

# Table of Contents

## Introduction

This document provides an overview of next-generation multicast virtual private networks (MVPNs) and describes how next-generation MVPN control and data plane protocols work together in the Juniper Networks® Junos® operating system (Junos OS). The target audience of this document is network architects, engineers, and operators.

The document includes the following:

- Overview of next-generation MVPNs—These sections include background material of how next-generation MVPNs work in general: concepts, terminology, control plane, and data plane.
- Junos OS next-generation MVPNs—These sections detail how Juniper Networks routers operate and interact with each other to set up next-generation MVPN routing and forwarding state in the network.

The scope of this document includes the following Junos OS features:

- Intra-AS MVPN membership discovery via BGP MCAST-VPN address family
- BGP customer-multicast (C-multicast) route exchange when the PE-CE protocol is Protocol Independent Multicast (PIM) sparse mode source-specific multicast (SSM), PIM sparse mode [any-source multicast (ASM)], PIM dense mode, or Internet Group Management Protocol (IGMP - source-tree-only mode)
- Generic routing encapsulation (GRE)-based inclusive provider tunnels signaled by PIM sparse mode (ASM)
- MPLS inclusive provider tunnels signaled by RSVP-Traffic Engineering (RSVP-TE) point-to-multipoint label-switched paths (LSPs)
- MPLS selective provider tunnels signaled by RSVP-TE point-to-multipoint LSPs

In addition to features identified in this document, Junos OS also supports the following features:

- Next-generation MVPN applications: extranet with point-to-multipoint traffic engineering
- GRE-based inclusive provider tunnels signaled by PIM sparse mode (SSM)
- Next-generation MVPN applications: hub and spoke

It is assumed that you are familiar with the following RFCs and IETF drafts:

- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4601, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*

- Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-mvpn-considerations-06.txt, *Mandatory Features in a Layer 3 Multicast BGP/MPLS VPN Solution*

**Related Documentation**

- Next-Generation MVPN Network Topology on page 2
- Next-Generation MVPN Concepts and Terminology on page 3

## Next-Generation MVPN Network Topology

Layer 3 BGP-MPLS virtual private networks (VPNs) are widely deployed in today's networks worldwide. Multicast applications, such as IPTV, are rapidly gaining popularity as is the number of networks with multiple, media-rich services merging over a shared Multiprotocol Label Switching (MPLS) infrastructure. The demand for delivering multicast service across a BGP-MPLS infrastructure in a scalable and reliable way is also increasing.

RFC 4364 describes protocols and procedures for building unicast BGP-MPLS VPNs. However, there is no framework specified in the RFC for provisioning multicast VPN (MVPN) services. In the past, Multiprotocol Label Switching Virtual Private Network (MVPN) traffic was overlaid on top of a BGP-MPLS network using a virtual LAN model based on Draft Rosen. Using the Draft Rosen approach, service providers were faced with control and data plane scaling issues of an overlay model and the maintenance of two routing/forwarding mechanisms: one for VPN unicast service and one for VPN multicast service. For more information about the limitations of Draft Rosen, see draft-rekhter-mboned-mvpn-deploy.

As a result, the IETF Layer 3 VPN working group published an Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*, that outlines a different architecture for next-generation MVPNs, as well as an accompanying RFC 2547 that proposes a BGP control plane for MVPNs. In turn, Juniper Networks delivered the industry's first implementation of BGP next-generation MVPNs in 2007.

All examples in this document refer to the network topology shown in Figure 1 on page 3:

- The service provider in this example offers VPN unicast and multicast services to Customer A (vpna).
- The VPN multicast source is connected to Site 1 and transmits data to groups 232.1.1.1 and 224.1.1.1.
- VPN multicast receivers are connected to Site 2 and Site 3.
- The provider edge router 1 (Router PE1) VRF table acts as the C-RP (using address 10.12.53.1) for C-PIM-SM ASM groups.
- The service provider uses RSVP-TE point-to-multipoint LSPs for transmitting VPN multicast data across the network.

Figure 1: Next-Generation MVPN Topology

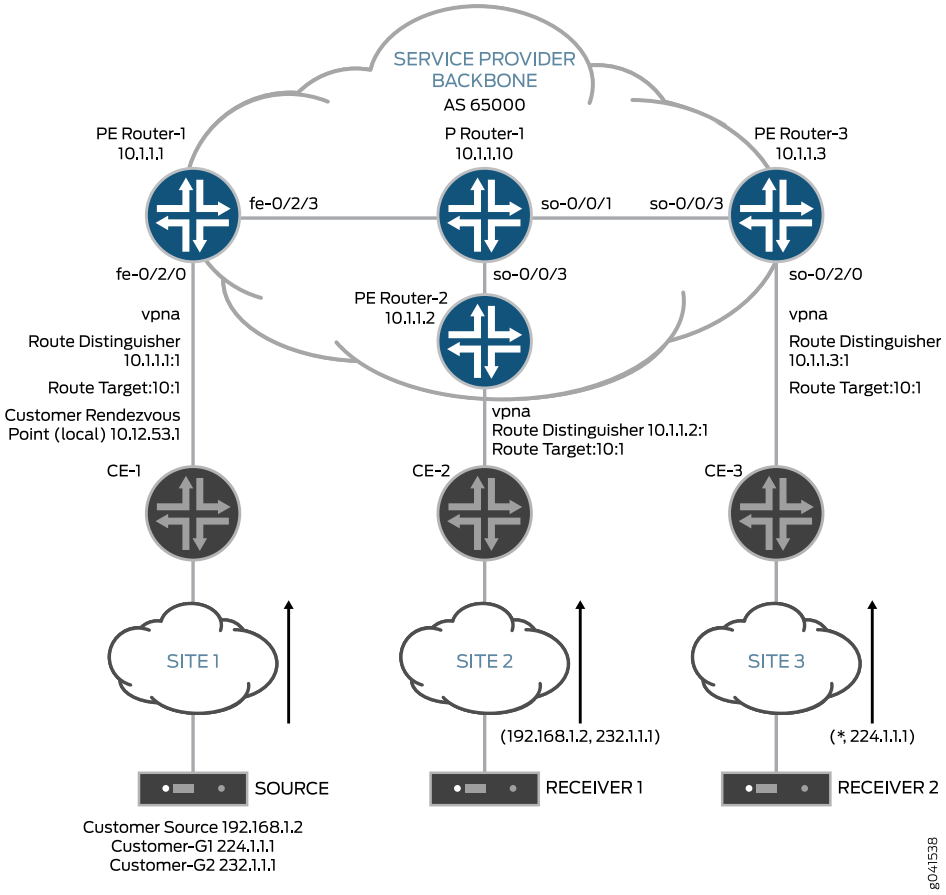


**Related Documentation**

- Next-Generation MVPN Concepts and Terminology on page 3
- Next-Generation MVPN Control Plane on page 6
- Next-Generation MVPN Data Plane on page 15
- *Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview*

## Next-Generation MVPN Concepts and Terminology

This section includes background material about how next-generation MVPNs work.

### Route Distinguisher and VRF Route Target Extended Community

Route distinguisher and VPN routing and forwarding (VRF) route target extended communities are an integral part of unicast BGP-MPLS virtual private networks (VPNs). Route distinguisher and route target are often confused in terms of their purpose in BGP-MPLS networks. As they play an important role in BGP next-generation MVPNs, it is important to understand what they are and how they are used as described in RFC 4364.

RFC 4364 describes the purpose of route distinguisher as the following:

*"A VPN-IPv4 address is a 12-byte quantity, beginning with an 8-byte Route Distinguisher (RD) and ending with a 4-byte IPv4 address. If several VPNs use the same IPv4 address prefix, the PEs translate these into unique VPN-IPv4 address prefixes. This ensures that if the same address is used in several different VPNs, it is possible for BGP to carry several completely different routes to that address, one for each VPN."*

Typically, each VRF table on a provider edge (PE) router is configured with a unique route distinguisher. Depending on the routing design, the route distinguisher can be unique or the same for a given VRF on other PE routers. A route distinguisher is an 8-byte number with two fields. The first field can be either an AS number (2 or 4 bytes) or an IP address (4 bytes). The second field is assigned by the user.

RFC 4364 describes the purpose of a VRF route target extended community as the following:

*"Every VRF is associated with one or more Route Target (RT) attributes.*

*When a VPN-IPv4 route is created (from an IPv4 route that the PE router has learned from a CE) by a PE router, it is associated with one or more route target attributes. These are carried in BGP as attributes of the route.*

*Any route associated with Route Target T must be distributed to every PE router that has a VRF associated with Route Target T. When such a route is received by a PE router, it is eligible to be installed in those of the PE's VRFs that are associated with Route Target T."*

The route target also contains two fields and is structured similar to a route distinguisher. The first field of the route target is either an AS number (2 or 4 bytes) or an IP address (4 bytes), and the second field is assigned by the user. Each PE router advertises its VPN-IPv4 routes with the route target (as one of the BGP path attributes) configured for the VRF table. The route target attached to the advertised route is referred to as the export route target. On the receiving PE router, the route target attached to the route is compared to the route target configured for the local VRF tables. The locally configured route target that is used in deciding whether a VPN-IPv4 route should be installed in a VRF table is referred to as the import route target.

## C-Multicast Routing

Customer multicast (C-multicast) routing information exchange refers to the distribution of customer PIM (C-PIM) join/prune messages received from local customer edge (CE) routers to other PE routers (toward the VPN multicast source).

## BGP MVPNs

BGP MVPNs use BGP as the control plane protocol between PE routers for MVPNs, including the exchange of C-multicast routing information. The support of BGP as a PE-PE protocol for exchanging C-multicast routes is mandated by Internet draft draft-ietf-l3vpn-mvpn-considerations-06.txt, *Mandatory Features in a Layer 3 Multicast BGP/MPLS VPN Solution*. The use of BGP for distributing C-multicast routing information is closely modeled after its highly successful counterpart of VPN unicast route distribution. Using BGP as the control plane protocol allows service providers to take advantage of

this widely deployed, feature-rich protocol. It also enables service providers to leverage their knowledge and investment in managing BGP-MPLS VPN unicast service to offer VPN multicast services.

## Sender and Receiver Site Sets

Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt describes an MVPN as a set of administrative policies that determine the PE routers that are in sender and receiver site sets.

A PE router can be a sender, a receiver, or both a sender and a receiver, depending on the configuration:

- A sender site set includes PE routers with local VPN multicast sources (VPN customer multicast sources either directly connected or connected via a CE router). A PE router that is in the sender site set is the sender PE router.
- A receiver site set includes PE routers that have local VPN multicast receivers. A PE router that is in the receiver site set is the receiver PE router.

## Provider Tunnels

Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt defines provider tunnels as the transport mechanisms used for forwarding VPN multicast traffic across service provider networks. Different tunneling technologies, such as generic routing encapsulation (GRE) and MPLS, can be used to create provider tunnels. Provider tunnels can be signaled by a variety of signaling protocols. This topic describes only PIM-SM (ASM) signaled IP GRE provider tunnels and RSVP-Traffic Engineering (RSVP-TE) signaled MPLS provider tunnels.

In BGP MVPNs, the sender PE router distributes information about the provider tunnel in a BGP attribute called provider multicast service interface (PMSI). By default, all receiver PE routers join and become the leaves of the provider tunnel rooted at the sender PE router.

Provider tunnels can be inclusive or selective:

- An inclusive provider tunnel (I-PMSI provider tunnel) enables a PE router that is in the sender site set of an MVPN to transmit multicast data to all PE routers that are members of that MVPN.
- A selective provider tunnel (S-PMSI provider tunnel) enables a PE router that is in the sender site set of an MVPN to transmit multicast data to a subset of the PE routers.

**Related Documentation**

- Next-Generation MVPN Network Topology on page 2
- Generating Next-Generation MVPN VRF Import and Export Policies on page 21
- Exchanging C-Multicast Routes on page 42
- *Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview*

## Next-Generation MVPN Control Plane

The BGP next-generation multicast virtual private network (MVPN) control plane, as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt and Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, distributes all the necessary information to enable end-to-end C-multicast routing exchange via BGP. The main tasks of the control plane (Table 1 on page 6) include MVPN autodiscovery, distribution of provider tunnel information, and PE-PE C-multicast route exchange.

**Table 1: Next-Generation MVPN Control Plane Tasks**

| Control Plane Task | Description |
|---|---|
| MVPN autodiscovery | A provider edge (PE) router discovers the identity of the other PE routers that participate in the same MVPN. |
| Distribution of provider tunnel information | A sender PE router advertises the type and identifier of the provider tunnel that it will use to transmit VPN multicast packets. |
| PE-PE C-Multicast route exchange | A receiver PE router propagates C-multicast join messages (C-joins) received over its VPN interface toward the VPN multicast sources. |

## BGP MCAST-VPN Address Family and Route Types

Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt introduced a BGP address family called MCAST-VPN for supporting next-generation MVPN control plane operations. The new address family is assigned the subsequent address family identifier (SAFI) of 5 by the Internet Assigned Numbers Authority (IANA).

A PE router that participates in a BGP-based next-generation MVPN network is required to send a BGP update message that contains MCAST-VPN network layer reachability information (NLRI). An MCAST-VPN NLRI contains route type, length, and variable fields. The value of each variable field depends on the route type.

Seven types of next-generation MVPN BGP routes (also referred to as routes in this topic) are specified (Table 2 on page 7). The first five route types are called autodiscovery MVPN routes. This topic also refers to Type 1-5 routes as non-C-multicast MVPN routes. Type 6 and Type 7 routes are called C-multicast MVPN routes.

Table 2: Next-generation MVPN BGP Route Types

| Usage | Type | Name | Description |
|---|---|---|---|
| Membership autodiscovery routes for inclusive provider tunnels | 1 | Inter autonomous system (inter-AS) I-PMSI autodiscovery route | • Originated by all next-generation MVPN PE routers.<br>• Used for advertising and learning intra autonomous system (intra-AS) MVPN membership information. |
| | 2 | Inter-AS I-PMSI AD route | • Originated by next-generation MVPN ASBR routers.<br>• Used for advertising and learning inter-AS MVPN membership information. |
| Autodiscovery routes for selective provider tunnels | 3 | S-PMSI AD route | • Originated by a sender router.<br>• Used for initiating a selective provider tunnel for a particular (C-S, C-G). |
| | 4 | Leaf AD route | • Originated by receiver PE routers in response to receiving a Type 3 route.<br>• Used by a sender PE router to discover the leaves of a selective provider tunnel.<br>• Also used for inter-AS operations that are not covered in this topic. |
| VPN multicast source discovery routes | 5 | Source active AD route | • Originated by the PE router that discovers an active VPN multicast source.<br>• Used by PE routers to learn the identity of active VPN multicast sources. |
| C-Multicast routes | 6 | Shared tree join route | • Originated by receiver PE routers.<br>• Originated when a PE router receives a shared tree C-join (C-*, C-G) through its PE-CE interface. |
| | 7 | Source tree join route | • Originated by receiver PE routers.<br>• Originated when a PE router receives a source tree C-join (C-S, C-G) or originated by the PE router that already has a Type 6 route and receives a Type 5 route. |

## Intra-AS MVPN Membership Discovery (Type 1 Routes)

All next-generation MVPN PE routers create and advertise a Type 1 intra-AS autodiscovery route (Figure 2 on page 8) for each MVPN to which they are connected. Table 3 on page 8 describes the format of each MVPN Type 1 intra-AS autodiscovery route.

Figure 2: Intra-AS I-PMSI AD Route Type MCAST-VPN NLRI Format

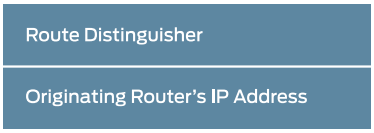


Table 3: Type 1 Intra-AS Autodiscovery Route MVPN Format Descriptions

| Field | Description |
|---|---|
| Route Distinguisher | Set to the route distinguisher configured for the VPN. |
| Originating Router's IP Address | Set to the IP address of the router originating this route. The address is typically the primary loopback address of the PE router. |

## Inter-AS MVPN Membership Discovery (Type 2 Routes)

Type 2 routes are used for membership discovery between PE routers that belong to different autonomous systems (ASs). Their use is not covered in this topic.

## Selective Provider Tunnels (Type 3 and Type 4 Routes)

A sender PE router that initiates a selective provider tunnel is required to originate a Type 3 intra-AS S-PMSI autodiscovery route with the appropriate PMSI attribute.

A receiver PE router responds to a Type 3 route by originating a Type 4 leaf autodiscovery route if it has local receivers interested in the traffic transmitted on the selective provider tunnel. Type 4 routes inform the sender PE router of the leaf PE routers.

## Source Active Autodiscovery Routes (Type 5 Routes)

Type 5 routes carry information about active VPN sources and the groups to which they are transmitting data. These routes can be generated by any PE router that becomes aware of an active source. Type 5 routes apply only for PIM-SM (ASM) when intersite source-tree-only mode is being used.

## C-Multicast Route Exchange (Type 6 and Type 7 Routes)

The C-multicast route exchange between PE routers refers to the propagation of C-joins from receiver PE routers to the sender PE routers.

In a next-generation MVPN, C-joins are translated into (or encoded as) BGP C-multicast MVPN routes and advertised via the BGP MCAST-VPN address family toward the sender PE routers.

Two types of C-multicast MVPN routes are specified:

- Type 6 C-multicast routes are used in representing information contained in a shared tree (C-*, C-G) join.
- Type 7 C-multicast routes are used in representing information contained in a source tree (C-S, C-G) join.

## PMSI Attribute

The provider multicast service interface (PMSI) attribute (Figure 3 on page 9) carries information about the provider tunnel. In a next-generation MVPN network, the sender PE router sets up the provider tunnel, and therefore is responsible for originating the PMSI attribute. The PMSI attribute can be attached to Type 1, Type 2, or Type 3 routes. Table 4 on page 9 describes each PMSI attribute format.

### Figure 3: PMSI Tunnel Attribute Format

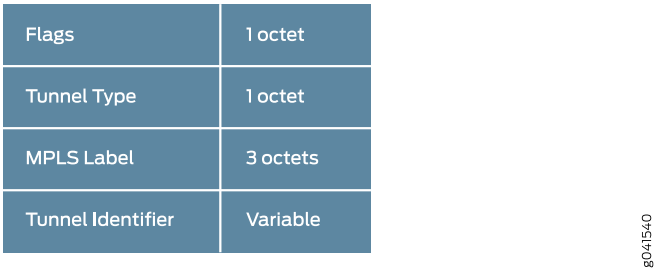


### Table 4: PMSI Tunnel Attribute Format Descriptions

| Field | Description |
|---|---|
| Flags | Currently has only one flag specified: Leaf Information Required. This flag is used for S-PMSI provider tunnel setup. |
| Tunnel Type | Identifies the tunnel technology used by the sender. Currently there are seven types of tunnels supported. |
| MPLS Label | Used when the sender PE router allocates the MPLS labels (also called upstream label allocation). This technique is described in RFC 5331 and is outside the scope of this topic. |
| Tunnel Identifier | Uniquely identifies the tunnel. Its value depends on the value set in the tunnel type field. |

For example, Router PE1 originates the following PMSI attribute:

**PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:6574:10.1.1.1]**

### VRF Route Import and Source AS Extended Communities

Two extended communities are specified to support next-generation MVPNs: source AS (**src-as**) and VRF route import extended communities.

The source AS extended community is an AS-specific extended community that identifies the AS from which a route originates. This community is mostly used for inter-AS operations, which is not covered in this topic.

The VPN routing and forwarding (VRF) route import extended community is an IP-address-specific extended community that is used for importing C-multicast routes in the VRF table of the active sender PE router to which the source is attached.

Each PE router creates a unique route target import and src-as community for each VPN and attaches them to the VPN-IPv4 routes.

**Related Documentation**

- Next-Generation MVPN Data Plane on page 15
- Distributing C-Multicast Routes on page 10
- Enabling Next-Generation MVPN Services on page 18
- Signaling Provider Tunnels and Data Plane Setup on page 28
- Originating Type 1 Intra-AS Autodiscovery Routes on page 24
- Next-Generation MVPN Network Topology on page 2

## Distributing C-Multicast Routes

While non-C-multicast multicast virtual private network (MVPN) routes (Type 1 – Type 5) are generally used by all provider edge (PE) routers in the network, C-multicast MVPN routes (Type 6 and Type 7) are only useful to the PE router connected to the active C-S or candidate rendezvous point (RP). Therefore, C-multicast routes need to be installed only in the VPN routing and forwarding (VRF) table on the active sender PE router for a given C-G. To accomplish this, Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt specifies to attach a special and dynamic route target to C-multicast MVPN routes (Figure 4 on page 11).

Figure 4: Attaching a Special and Dynamic Route Target to C-Multicast MVPN Routes

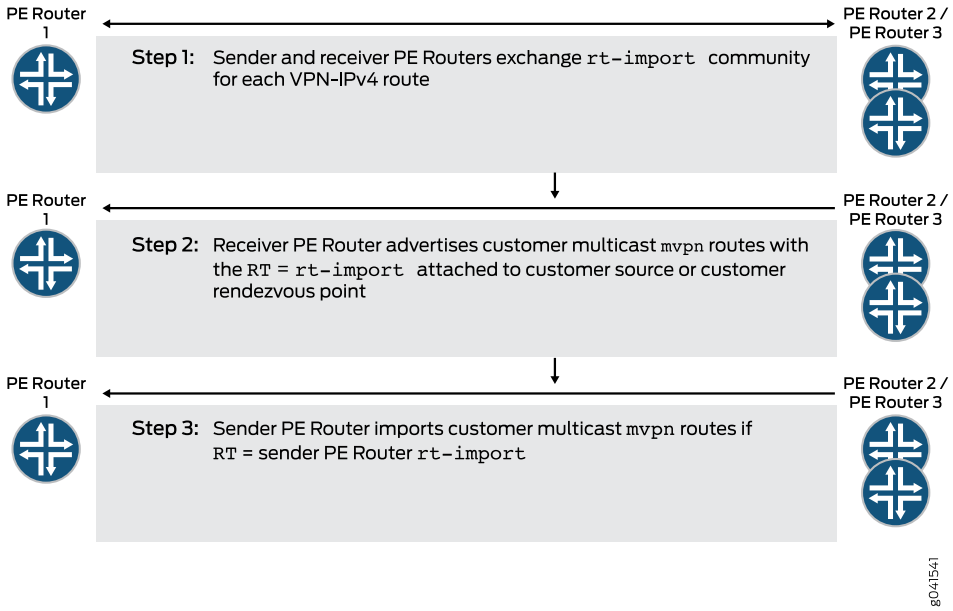


The route target attached to C-multicast routes is also referred to as the C-multicast import route target and should not to be confused with route target import (Table 5 on page 11). Note that C-multicast MVPN routes differ from other MVPN routes in one essential way: they carry a dynamic route target whose value depends on the identity of the active sender PE router at a given time and can change if the active PE router changes.

Table 5: Distinction Between Route Target Improt Attached to VPN-IPv4 Routes and Route Target Attached to C-Multicast MVPN Routes

| Route Target Import Attached to VPN-IPV4 Routes | Route Target Attached to C-Multicast MVPN Routes |
|---|---|
| Value generated by the originating PE router. Must be unique per VRF table. | Value depends on the identity of the active PE router. |
| Static. Created upon configuration to help identify to which PE router and to which VPN the VPN unicast routes belong. | Dynamic because if the active sender PE router changes, then the route target attached to the C-multicast routes must change to target the new sender PE router. For example, a new VPN source attached to a different PE router becomes active and preferred. |

A PE router that receives a local C-join determines the identity of the active sender PE router by performing a unicast route lookup for the C-S or candidate rendezvous point (router) [candidate RP] in the unicast VRF table. If there is more than one route, the receiver PE router chooses a single forwarder PE router. The procedures used for choosing a single forwarder are outlined in Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt and are not covered in this topic.

After the active sender (upstream) PE router is selected, the receiver PE router constructs the C-multicast MVPN route corresponding to the local C-join.

After the C-multicast route is constructed, the receiver PE router needs to attach the correct route target to this route targeting the active sender PE router. As mentioned, each PE router creates a unique VRF route target import community and attaches it to the VPN-IPv4 routes. When the receiver PE router does a route lookup for C-S or candidate RP, it can extract the value of the route target import associated with this route and set the value of the C-import route target to the value of the route target import.

On the active sender PE router, C-multicast routes are imported only if they carry the route target whose value is the same as the route target import that the sender PE router generated.

## Constructing C-Multicast Routes

A PE router originates a C-multicast MVPN route in response to receiving a C-join through its PE-CE interface. See Figure 5 on page 12 for the formats in the C-multicast route encoded in MCAST-VPN NLRI. Table 6 on page 12 describes each field.

### Figure 5: C-Multicast Route Type MCAST-VPN NLRI Format

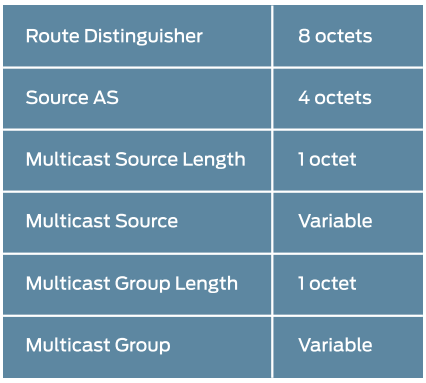


### Table 6: C-Multicast Route Type MCAST-VPN NLRI Format Descriptions

| Field | Description |
|---|---|
| Route Distinguisher | Set to the route distinguisher of the C-S or candidate RP (the route distinguisher associated with the upstream PE router). |
| Source AS | Set to the value found in the **src-as** community of the C-S or candidate RP. |
| Multicast Source Length | Set to 32 for IPv4 and to 128 for IPv6 C-S or candidate RP IP addresses. |
| Multicast Source | Set to the IP address of the C-S or candidate RP. |
| Multicast Group Length | Set to 32 for IPv4 and to 128 for IPv6 C-G addresses. |
| Multicast Group | Set to the C-G of the received C-join. |

This same structure is used for encoding both Type 6 and Type 7 routes with two differences:

- The first difference is the value used for the multicast source field. For Type 6 routes, this field is set to the IP address of the candidate RP configured. For Type 7 routes, this field is set to the IP address of the C-S contained in the (C-S, C-G) message.
- The second difference is the value used for the route distinguisher. For Type 6 routes, this field is set to the route distinguisher that is attached to the IP address of the candidate RP. For Type 7 routes, this field is set to the route distinguisher that is attached to the IP address of the C-S.

## Eliminating PE-PE Distribution of (C-*, C-G) State Using Source Active Autodiscovery Routes

PE routers must maintain additional state when the C-multicast routing protocol is Protocol Independent Multicast-Sparse Mode (PIM-SM) in any-source multicast (ASM). This is a requirement because with ASM, the receivers first join the shared tree rooted at the candidate RP (called a candidate RP tree or candidate RPT). However, as the VPN multicast sources become active, receivers learn the identity of the sources and join the tree rooted at the source (called a customer shortest-path tree or C-SPT). The receivers then send a prune message to the candidate RP to stop the traffic coming through the shared tree for the group that they joined to the C-SPT. The switch from candidate RPT to C-SPT is a complicated process requiring additional state.

Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt specifies optional procedures that completely eliminate the need for joining the candidate RPT. These procedures require PE routers to keep track of all active VPN sources using one of two options. The first option is to colocate the candidate RP on one of the PE routers. The second option is to use the Multicast Source Discovery Protocol (MSDP) between one of the PE routers and the customer candidate RP.

In this approach, a PE router that receives a local (C-*, C-G) join creates a Type 6 route, but does not advertise the route to the remote PE routers until it receives information about an active source. The PE router acting as the candidate RP (or that learns about active sources via MSDP) is responsible for originating a Type 5 route. A Type 5 route carries information about the active source and the group addresses. The information contained in a Type 5 route is enough for receiver PE routers to join the C-SPT by originating a Type 7 route toward the sender PE router, completely skipping the advertisement of the Type 6 route that is created when a C-join is received.
Figure 6 on page 14 shows the format of a source active (SA) autodiscovery route.
Table 7 on page 14 describes each format.

Figure 6: Source Active Autodiscovery Route Type MCAST-VPN NLRI Format

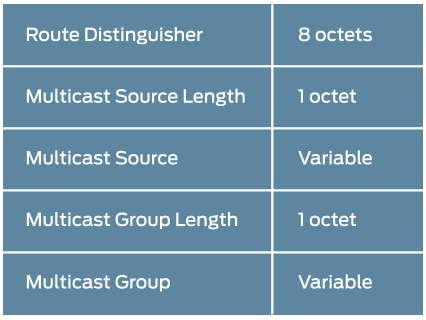


Table 7: Source Active Autodiscovery Route Type MCAST-VPN NLRI Format Descriptions

| Field | Description |
|---|---|
| Route Distinguisher | Set to the route distinguisher configured on the router originating the SA autodiscovery route. |
| Multicast Source Length | Set to 32 for IPv4 and to 128 for IPv6 C-S IP addresses. |
| Multicast Source | Set to the IP address of the C-S that is actively transmitting data to C-G. |
| Multicast Group Length | Set to 32 for IPv4 and to 128 for IPv6 C-G addresses. |
| Multicast Group | Set to the IP address of the C-G to which C-S is transmitting data. |

## Receiving C-Multicast Routes

The sender PE router imports C-multicast routes into the VRF table based on the route target of the route. If the route target attached to the C-multicast MVPN route matches the route target import community originated by this router, the C-multicast MVPN route is imported into the VRF table. If not, it is discarded.

Once the C-multicast MVPN routes are imported, they are translated back to C-joins and passed on to the VRF C-PIM protocol for further processing per normal PIM procedures.

**Related Documentation**

- Enabling Next-Generation MVPN Services on page 18
- Exchanging C-Multicast Routes on page 42
- Distributing C-Multicast Routes on page 10
- Next-Generation MVPN Network Topology on page 2

## Next-Generation MVPN Data Plane

A next-generation multicast virtual private network (MVPN) data plane is composed of provider tunnels originated by and rooted at the sender provider edge (PE) routers and the receiver PE routers as the leaves of the provider tunnel.

A provider tunnel can carry data for one or more VPNs. Those provider tunnels that carry data for more than one VPN are called aggregate provider tunnels and are outside the scope of this topic. Here, we assume that a provider tunnel carries data for only one VPN.

This topic covers two types of tunnel technologies: IP generic routing encapsulation (GRE) provider tunnels signaled by Protocol Independent Multicast-Sparse Mode (PIM-SM) any-source multicast (ASM) and MPLS provider tunnels signaled by RSVP-Traffic Engineering (RSVP-TE).

When a provider tunnel is signaled by PIM, the sender PE router runs another instance of the PIM protocol on the provider's network (P-PIM) that signals a provider tunnel for that VPN. When a provider tunnel is signaled by RSVP-TE, the sender PE router initiates a point-to-multipoint label-switched path (LSP) toward receiver PE routers by using point-to-multipoint RSVP-TE protocol messages. In either case, the sender PE router advertises the tunnel signaling protocol and the tunnel ID to other PE routers via BGP by attaching the provider multicast service interface (PMSI) attribute to either the Type 1 intra-AS autodiscovery routes (inclusive provider tunnels) or Type 3 S-PMSI autodiscovery routes (selective provider tunnels).

NOTE: The sender PE router goes through two steps when setting up the data plane. First, using the PMSI attribute, it advertises the provider tunnel it is using via BGP. Second, it actually signals the tunnel using whatever tunnel signaling protocol is configured for that VPN. This allows receiver PE routers to bind the tunnel that is being signaled to the VPN that imported the Type 1 intra-AS autodiscovery route. Binding a provider tunnel to a VRF table enables a receiver PE router to map the incoming traffic from the core network on the provider tunnel to the local target VRF table.

The PMSI attribute contains the provider tunnel type and an identifier. The value of the provider tunnel identifier depends on the tunnel type. Table 8 on page 15 identifies the tunnel types specified in Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt.

### Table 8: Tunnel Types Supported by PMSI Tunnel Attribute

| Tunnel Type | Description |
|---|---|
| 0 | No tunnel information present |
| 1 | RSVP-TE point-to-multipoint LSP |
| 2 | Multicast LDP point-to-multipoint LSP |

Table 8: Tunnel Types Supported by PMSI Tunnel Attribute *(continued)*

| Tunnel Type | Description |
|---|---|
| 3 | PIM-SSM tree |
| 4 | PIM-SM tree |
| 5 | PIM-Bidir tree |
| 6 | Ingress replication |
| 7 | Multicast LDP multipoint-to-multipoint LSP |

## Inclusive Provider Tunnels

This section describes various types of provider tunnels and attributes of provider tunnels.

### PMSI Attribute of Inclusive Provider Tunnels Signaled by PIM-SM

When the Tunnel Type field of the PMSI attribute is set to 4 (PIM-SM Tree), the tunnel identifier field contains **<Sender Address, P-Multicast Group Address>**. The **Sender Address** field is set to the router ID of the sender PE router. The P-multicast group address is set to a multicast group address from the service provider's P-multicast address space and uniquely identifies the VPN. A receiver PE router that receives an intra-AS autodiscovery route with a PMSI attribute whose tunnel type is PIM-SM is required to join the provider tunnel.

For example, if the service provider deploys PIM-SM provider tunnels (instead of RSVP-TE provider tunnels), Router PE1 advertises the following PMSI attribute:

**PMSI: 0:PIM-SM:label[0:0:0]:Sender10.1.1.1 Group 239.1.1.1**

### PMSI Attribute of Inclusive Provider Tunnels Signaled by RSVP-TE

When the tunnel type field of the PMSI attribute is set to 1 (RSVP-TE point-to-multipoint LSP), the tunnel identifier field contains an RSVP-TE point-to-multipoint session object as described in RFC 4875. The session object contains the **<Extended Tunnel ID, Reserved, Tunnel ID, P2MP ID>** associated with the point-to-multipoint LSPs.

The PE router that originates the PMSI attribute is required to signal an RSVP-TE point-to-multipoint LSP and the sub-LSPs. A PE router that receives this PMSI attribute must establish the appropriate state to properly handle the traffic received over the sub-LSP.

For example, Router PE1 advertises the following PMSI attribute:

**PMSI: Flags 0:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:6574:10.1.1.1]**

## Selective Provider Tunnels (S-PMSI Autodiscovery/Type 3 and Leaf Autodiscovery/Type 4 Routes)

A selective provider tunnel is used for mapping a specific C-multicast flow (a (C-S, C-G) pair) onto a specific provider tunnel. There are a variety of situations in which selective provider tunnels can be useful. For example, they can be used for putting high-bandwidth VPN multicast data traffic onto a separate provider tunnel rather than the default inclusive provider tunnel, thus restricting the distribution of traffic to only those PE routers with active receivers.

In BGP next-generation multicast virtual private networks (MVPNs), selective provider tunnels are signaled using Type 3 Selective-PMSI (S-PMSI) autodiscovery routes. See Figure 7 on page 17 and Table 9 on page 17 for details. The sender PE router sends a Type 3 route to signal that it is sending traffic for a particular (C-S, C-G) flow using an S-PMSI provider tunnel.

### Figure 7: S-PMSI Autodiscovery Route Type Multicast (MCAST)-VPN Network Layer Reachability Information (NLRI) Format

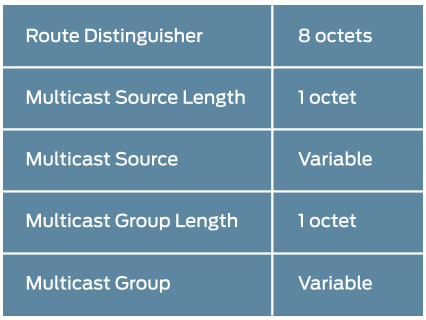


g041544

### Table 9: S-PMSI Autodiscovery Route Type Format Descriptions

| Field | Description |
|---|---|
| Route Distinguisher | Set to the route distinguisher configured on the router originating this route. |
| Multicast Source Length | Set to 32 for IPv4 and to 128 for IPv6 C-S IP addresses. |
| Multicast Source | Set to the C-S IP address. |
| Multicast Group Length | Set to 32 for IPv4 and to 128 for IPv6 C-G addresses. |
| Multicast Group | Set to the C-G address. |

The S-PMSI autodiscovery (Type 3) route carries a PMSI attribute similar to the PMSI attribute carried with intra-AS autodiscovery (Type 1) routes. The **Flags** field of the PMSI attribute carried by the S-PMSI autodiscovery route is set to the leaf information required. This flag signals receiver PE routers to originate a Type 4 leaf autodiscovery route (Figure 8 on page 18) to join the selective provider tunnel if they have active receivers.

See Table 10 on page 18 for details of leaf autodiscovery route type MCAST-VPN NLRI format descriptions.

Figure 8: Leaf Autodiscovery Route Type MCAST-VPN NLRI Format

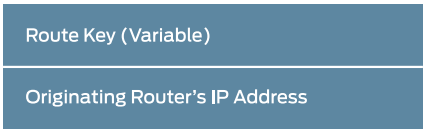


Table 10: Leaf Autodiscovery Route Type MCAST-VPN NLRI Format Descriptions

| Field | Description |
|---|---|
| Route Key | Contains the original Type 3 route received. |
| Originating Router's IP Address | Set to the IP address of the PE router originating the leaf autodiscovery route This is typically the primary loopback address. |

**Related Documentation**

- Next-Generation MVPN Control Plane on page 6
- Enabling Next-Generation MVPN Services on page 18
- Signaling Provider Tunnels and Data Plane Setup on page 28
- Next-Generation MVPN Network Topology on page 2

## Enabling Next-Generation MVPN Services

Juniper Networks introduced the industry's first implementation of BGP next-generation multicast virtual private networks (MVPNs). See Figure 9 on page 19 for a summary of a Junos OS next-generation MVPN routing flow.

Figure 9: Junos OS Next-Generation MVPN Routing Flow

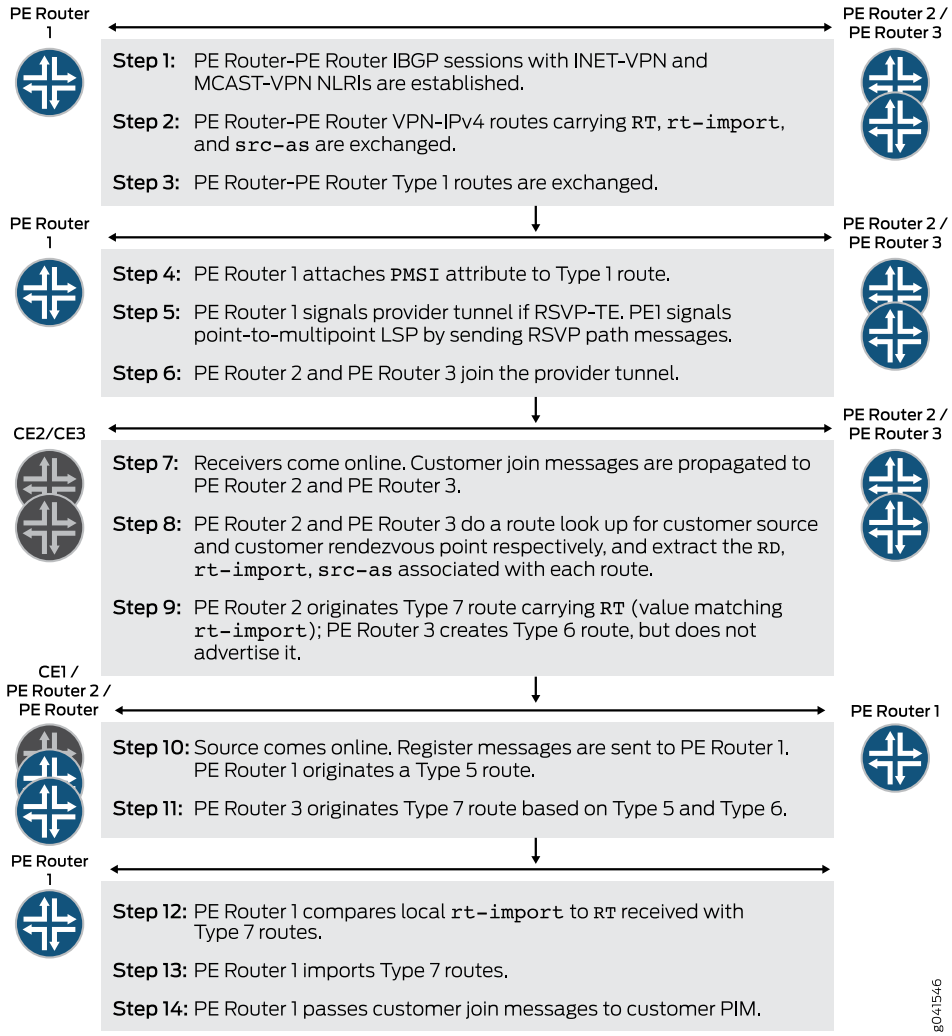


Next-generation MVPN services are configured on top of BGP-MPLS unicast VPN services.

You can configure a Juniper Networks PE router that is already providing unicast BGP-MPLS VPN connectivity to support multicast VPN connectivity in three steps:

1. Configure the provider edge (PE) routers to support the BGP multicast VPN address family by including the **signaling** statement at the **[edit protocols bgp group *group-name* family inet-mvpn]** hierarchy level. This address family enables PE routers to exchange MVPN routes.
2. Configure the PE routers to support the MVPN control plane tasks by including the **mvpn** statement at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level. This statement signals PE routers to initialize the MVPN module that is responsible for the majority of next-generation MVPN control plane tasks.
3. Configure the sender PE router to signal a provider tunnel by including the **provider-tunnel** statement at the **[edit routing-instances *routing-instance-name*]**

hierarchy level. You must also enable the tunnel signaling protocol (RSVP-TE or P-PIM) if it is not part of the unicast VPN service configuration. To enable the tunnel signaling protocol, include the **rsvp-te** or **pim-asm** statements at the **[edit routing-instances *routing-instance-name* provider-tunnel]** hierarchy level.

After these three statements are configured and each PE router has established internal BGP (IBGP) sessions using both INET-VPN and MCAST-VPN address families, four routing tables are automatically created. These tables are **bgp.l3vpn.0, bgp.mvpn.0**, **<routing-instance-name>.inet.0**, and **<routing-instance-name>.mvpn.0**. See Table 11 on page 20

### Table 11: Automatically Generated Routing Tables

| Automatically Generated Routing Table | Description |
|---|---|
| bgp.l3vpn.0 | Populated with VPN-IPv4 routes received from remote PE routers via the INET-VPN address family. The routes in the **bgp.l3vpn.0** table are in the form of **RD:IPv4-address** and carry one or more routing table communities. In a next-generation MVPN network, these routes also carry **rt-import** and **src-as** communities. |
| bgp.mvpn.0 | Populated by MVPN routes (Type 1 – Type 7). Received from remote PE routers via the MCAST-VPN address family. Routes in this table carry one or more routing table communities. |
| <routing-instance-name>.inet.0 | Populated by local and remote VPN unicast routes. The local VPN routes are typically learned from local CE routers via protocols such as BGP, OSPF, and RIP, or via a static configuration. The remote VPN routes are imported from the **bgp.l3vpn.0** table if their routing table matches one of the import routing tables configured for the VPN. When remote VPN routes are imported from the **bgp.l3vpn.0** table, their route distinguisher is removed, leaving them as regular unicast IPv4 addresses. |
| <routing-instance-name>.mvpn.0 | Populated by local and remote MVPN routes. The local MVPN routes are typically the locally originated routes, such as Type 1 intra-AS autodiscovery routes, or Type 7 C-multicast routes. The remote MVPN routes are imported from the **bgp.mvpn.0** table based on their route target. The import route target used for accepting MVPN routes into the **<routing-instance-name>.mvpn.0** table is different for C-multicast MVPN routes (Type 6 and Type 7) versus non-C-multicast MVPN routes (Type 1 – Type 5). |

**Related Documentation**

- Next-Generation MVPN Network Topology on page 2
- Generating Next-Generation MVPN VRF Import and Export Policies on page 21
- Generating Source AS and Route Target Import Communities on page 24
- Originating Type 1 Intra-AS Autodiscovery Routes on page 24
- Signaling Provider Tunnels and Data Plane Setup on page 28

## Generating Next-Generation MVPN VRF Import and Export Policies

In Junos OS, the policy module is responsible for VPN routing and forwarding (VRF) route import and export decisions. You can configure these policies explicitly, or Junos OS can generate them internally for you to reduce user-configured statements and simplify configuration. Junos OS generates all necessary policies for supporting next-generation multicast virtual private network (MVPN) import and export decisions. Some of these policies affect normal VPN unicast routes.

The system gives a name to each internal policy it creates. The name of an internal policy starts and ends with a "__" notation. Also the keyword **internal** is added at the end of each internal policy name. You can display these internal policies using the **show policy** command.

### Policies That Support Unicast BGP-MPLS VPN Services

A Juniper Networks provider edge (PE) router requires a vrf-import and a vrf-export policy to control unicast VPN route import and export decisions for a VRF. You can configure these policies explicitly at the **[edit routing-instances *routing-instance-name* vrf-import *import_policy_name*]** and **[edit routing-instances *routing-instance-name* vrf-export *export_policy_name*]** hierarchy level. Alternately, you can configure only the route target for the VRF at the **[edit routing-instances *routing-instance-name* vrf-target]** hierarchy level, and Junos OS then generates these policies automatically for you. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

The following list identifies the automatically generated policy names and where they are applied:

**Policy:** vrf-import

**Naming convention:** __vrf-import-<routing-instance-name>-internal__

**Applied to:** VPN-IPv4 routes in the bgp.l3vpn.0 table

**Policy:** vrf-export

**Naming convention:** __vrf-export-<routing-instance-name>-internal__

**Applied to:** Local VPN routes in the <routing-instance-name>.inet.0 table

Use the **show policy __vrf-import-vpna-internal__** command to verify that Router PE1 has created the following **vrf-import** and **vrf-export** policies based on a vrf-target of **target:10:1**. In this example, we see that the **vrf-import** policy is constructed to accept a route if the route target of the route matches **target:10:1**. Similarly, a route is exported with a route target of **target:10:1**.

```
user@PE1> show policy __vrf-import-vpna-internal__
Policy __vrf-import-vpna-internal__:
  Term unnamed:
    from community __vrf-community-vpna-common-internal__ [target:10:1]
    then accept
  Term unnamed:
    then reject
```

```
user@PE1> show policy __vrf-export-vpna-internal__
Policy __vrf-export-vpna-internal__:
  Term unnamed:
    then community + __vrf-community-vpna-common-internal__ [target:10:1] accept
```

The values in this example are as follows:

- Internal import policy name: __vrf-import-vpna-internal__
- Internal export policy name: __vrf-export-vpna-internal__
- RT community used in both import and export policies: __vrf-community-vpna-common-internal__
- RT value: target:10:1

## Policies That Support Next-Generation MVPN Services

When you configure the **mvpn** statement at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level, Junos OS automatically creates three new internal policies: one for export, one for import, and one for handling Type 4 routes. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

The following list identifies the automatically generated policy names and where they are applied:

**Policy 1:** This policy is used to attach **rt-import** and **src-as** extended communities to VPN-IPv4 routes.

**Policy name:** __vrf-mvpn-export-inet-<routing-instance-name>-internal__

**Applied to:** All routes in the <routing-instance-name>inet.0 table

Use the **show policy __vrf-mvpn-export-inet-vpna-internal__** command to verify that the following export policy is created on Router PE1. Router PE1 adds **rt-import:10.1.1.1:64** and **src-as:65000:0** communities to unicast VPN routes through this policy.

```
user@PE1> show policy __vrf-mvpn-export-inet-vpna-internal__
Policy __vrf-mvpn-export-inet-vpna-internal__:
     Term unnamed:
    then community + __vrf-mvpn-community-rt_import-vpna-internal__
[rt-import:10.1.1.1:64 ] community + __vrf-mvpn-community-src_as-vpna-internal__
 [src-as:65000:0 ] accept
```

The values in this example are as follows:

- Policy name: __vrf-mvpn-export-inet-vpna-internal__
- rt-import community name: __vrf-mvpn-community-rt_import-vpna-internal__
- rt-import community value: rt-import:10.1.1.1:64
- src-as community name: __vrf-mvpn-community-src_as-vpna-internal__
- src-as community value: src-as:65000:0

**Policy 2:** This policy is used to import C-Mmulticast routes from the **bgp.mvpn.0** table to the **<routing-instance-name>.mvpn.0** table.

**Policy name:** __vrf-mvpn-import-cmcast-<routing-instance-name>-internal__

**Applied to:** C-multicast (MVPN) routes in the bgp.mvpn.0 table

Use the **show policy __vrf-mvpn-import-cmcast-vpna-internal__** command to verify that the following import policy is created on Router PE1. The policy accepts those C-multicast MVPN routes carrying a route target of **target:10.1.1.1:64** and installs them in the **vpna.mvpn.0** table.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-vpna-internal__
Policy __vrf-mvpn-import-cmcast-vpna-internal__:
  Term unnamed:
   from community __vrf-mvpn-community-rt_import-target-vpna-internal__
[target:10.1.1.1:64 ]
  then accept
 Term unnamed:
  then reject
```

The values in this example are as follows:

- Policy name: __vrf-mvpn-import-cmcast-vpna-internal__
- C-multicast import RT community: __vrf-mvpn-community-rt_import-target-vpna-internal__
- Community value: target:10.1.1.1:64

**Policy 3:** This policy is used for importing Type 4 routes and is created by default even if a selective provider tunnel is not configured. The policy affects only Type 4 routes received from receiver PE routers.

**Policy name:** __vrf-mvpn-import-cmcast-leafAD-global-internal__

**Applied to:** Type 4 routes in the bgp.mvpn.0 table

Use the **show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__** command to verify that the following import policy is created on Router PE1.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__
Policy __vrf-mvpn-import-cmcast-leafAD-global-internal__:
  Term unnamed:
    from community __vrf-mvpn-community-rt_import-target-global-internal__
[target:10.1.1.1:0 ]
    then accept
  Term unnamed:
    then reject
```

**Related Documentation**

- *Understanding MBGP Multicast VPN Extranets*
- *Example: Configuring MBGP Multicast VPN Extranets*
- *Hub-and-Spoke VPNs Using Next-Hop Tunnel Binding Overview*
- Enabling Next-Generation MVPN Services on page 18

## Generating Source AS and Route Target Import Communities

Both route target import (**rt-import**) and source autonomous system (**src-as**) communities contain two fields (following their respective keywords). In Junos OS, a provider edge (PE) router constructs the route target import community using its router ID in the first field and a per-VRF unique number in the second field. The router ID is normally set to the primary loopback IP address of the PE router. The unique number used in the second field is an internal number derived from the routing-instance table index. The combination of the two numbers creates a route target import community that is unique to the originating PE router and unique to the VPN routing and forwarding (VRF) instance from which it is created.

For example, Router PE1 creates the following route target import community: **rt-import:10.1.1.1:64**.

Since the route target import community is constructed using the primary loopback address and the routing-instance table index of the PE router, any event that causes either number to change triggers a change in the value of the route target import community. This in turn requires VPN-IPv4 routes to be re-advertised with the new route target import community. Under normal circumstances, the primary loopback address and the routing-instance table index numbers do not change. If they do change, Junos OS updates all related internal policies and re-advertises VPN-IPv4 routes with the new **rt-import** and **src-as** values per those policies.

To ensure that the route target import community generated by a PE router is unique across VRF tables, the Junos OS Policy module restricts the use of primary loopback addresses to next-generation multicast virtual private network (MVPN) internal policies only. You are not permitted to configure a route target for any VRF table (MVPN or otherwise) using the primary loopback address. The commit fails with an error if the system finds a user-configured route target that contains the IP address used in constructing the route target import community.

The global administrator field of the **src-as** community is set to the local AS number of the PE router originating the community, and the local administrator field is set to **0**. This community is used for inter-AS operations but needs to be carried along with all VPN-IPv4 routes.

For example, Router PE1 creates an **src-as** community with a value of **src-as:65000:0**.

**Related Documentation**

- Originating Type 1 Intra-AS Autodiscovery Routes on page 24
- Generating Next-Generation MVPN VRF Import and Export Policies on page 21
- Enabling Next-Generation MVPN Services on page 18

## Originating Type 1 Intra-AS Autodiscovery Routes

Every provider edge (PE) router that is participating in the next-generation multicast virtual private network (MVPN) is required to originate a Type 1 intra-AS autodiscovery route. In Junos OS, the MVPN module is responsible for installing the intra-AS

autodiscovery route in the local **<routing-instance-name>.mvpn.0** table. All PE routers advertise their local Type 1 routes to each other. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show route table vpna.mvpn.0** command to verify that Router PE1 has installed intra-AS AD routes in the **vpna.mvpn.0** table. The route is installed by the MVPN protocol (meaning it is the MVPN module that originated the route), and the mask for the entire route is /240.

```
user@PE1> show route table vpna.mvpn.0
vpna.mvpn.0: 6 destinations, 9 routes (6 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.1.1.1:1:10.1.1.1/240
                   *[MVPN/70] 04:09:44, metric2 1
                      Indirect
```

## Attaching Route Target Community to Type 1 Routes

Intra-AS AD routes are picked up by the BGP protocol from the **<routing-instance-name>.mvpn.0** table and advertised to the remote PE routers via the MCAST-VPN address family. By default, intra-AS autodiscovery routes carry the same route target community that is attached to the unicast VPN-IPv4 routes. If the unicast and multicast network topologies are not congruent, then you can configure a different set of import route target and export route target communities for non-C-multicast MVPN routes (C-multicast MVPN routes always carry a dynamic import route target).

Multicast route targets are configured by including the **import-target** and **export-target** statements at the **[edit routing-instances *routing-instance-name* protocols mvpn route-target]** hierarchy level.

Junos OS creates two additional internal policies in response to configuring multicast route targets. These policies are applied to non-C-multicast MVPN routes during import and export decisions. Multicast VPN routing and forwarding (VRF) internal import and export polices follow a naming convention similar to unicast VRF import and export policies. The contents of these polices are also similar to policies applied to unicast VPN routes.

The following list identifies the default policy names and where they are applied:

**Multicast VRF import policy:**
__vrf-mvpn-import-target-<routing-instance-name>-internal__

**Multicast VRF export policy:** __vrf-mvpn-export-target-<routing-instance-name>-internal__

Use the **show policy __vrf-mvpn-import-target-vpna-internal__** command on Router PE1 to verify that Router PE1 has created the following internal MVPN policies if import-target and export-target are configured to be target:10:2:

```
user@PE1> show policy __vrf-mvpn-import-target-vpna-internal__
Policy __vrf-mvpn-import-target-vpna-internal__:
  Term unnamed:
   from community __vrf-mvpn-community-import-vpna-internal__ [target:10:2 ]
   then accept
```

```
  Term unnamed:
   then reject

user@PE1> show policy __vrf-mvpn-export-target-vpna-internal__
Policy __vrf-mvpn-export-target-vpna-internal__:
Term unnamed:
then community + __vrf-mvpn-community-export-vpna-internal__ [target:10:2 ] accept
```

The values in this example are as follows:

- Multicast import RT community: __vrf-mvpn-community-import-vpna-internal__
- Multicast export RT community: __vrf-mvpn-community-export-vpna-internal__ Value: target:10:2

## Attaching the PMSI Attribute to Type 1 Routes

The provider multicast service interface (PMSI) attribute is originated and attached to Type 1 intra-AS autodiscovery routes by the sender PE routers when the **provider-tunnel** statement is included at the **[edit routing-instances *routing-instance-name*]** hierarchy level. Since provider tunnels are signaled by the sender PE routers, this statement is not necessary on the PE routers that are known to have VPN multicast receivers only.

If the provider tunnel configured is Protocol Independent Multicast-Sparse Mode (PIM-SM) any-source multicast (ASM), then the PMSI attribute carries the IP address of the sender-PE and provider tunnel group address. The provider tunnel group address is assigned by the service provider (through configuration) from the provider's multicast address space and is not to be confused with the multicast addresses used by the VPN customer.

If the provider tunnel configured is the RSVP-Traffic Engineering (RSVP-TE) type, then the PMSI attribute carries the RSVP-TE point-to-multipoint session object. This point-to-multipoint session object is used as the identifier for the parent point-to-multipoint label-switched path (LSP) and contains the fields shown in Figure 10 on page 26.

### Figure 10: RSVP-TE Point-to-Multipoint Session Object Format

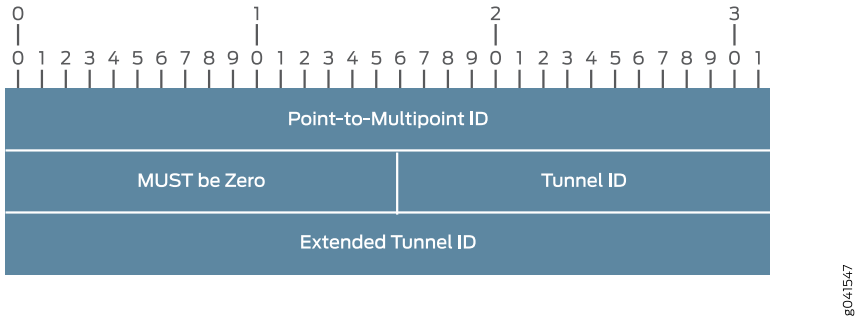


In Junos OS, the **P2MP ID** and **Extended Tunnel ID** fields are set to the router ID of the sender PE router. The **Tunnel ID** is set to the port number used for the point-to-multipoint RSVP session that is unique for the length of the RSVP session.

Use the **show rsvp session p2mp detail** command to verify that Router PE1 signals the following RSVP sessions to Router PE2 and Router PE3 (using port number 6574). In this example, Router PE1 is signaling a point-to-multipoint LSP named 10.1.1.1:65535:mvpn:vpna with two sub-LSPs. Both sub-LSPs **10.1.1.3:10.1.1.1:65535:mvpn:vpna** and **10.1.1.2:10.1.1.1:65535:mvpn:vpna** use the same RSVP port number (6574) as the parent point-to-multipoint LSP.

```
user@PE1> show rsvp session p2mp detail
Ingress RSVP: 2 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2

10.1.1.3
  From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: 10.1.1.3:10.1.1.1:65535:mvpn:vpna, LSPpath: Primary
  P2MP LSPname: 10.1.1.1:65535:mvpn:vpna
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299968
  Resv style: 1 SE, Label in: -, Label out: 299968
  Time left: -, Since: Wed May 27 07:36:22 2009
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 6574 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.12.100.6 (fe-0/2/3.0) 27 pkts
  RESV rcvfrom: 10.12.100.6 (fe-0/2/3.0) 27 pkts
  Explct route: 10.12.100.6 10.12.100.22
  Record route: <self> 10.12.100.6 10.12.100.22

10.1.1.2
  From: 10.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: 10.1.1.2:10.1.1.1:65535:mvpn:vpna, LSPpath: Primary
  P2MP LSPname: 10.1.1.1:65535:mvpn:vpna
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299968
  Resv style: 1 SE, Label in: -, Label out: 299968
  Time left: -, Since: Wed May 27 07:36:22 2009
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 6574 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.12.100.6 (fe-0/2/3.0) 27 pkts
  RESV rcvfrom: 10.12.100.6 (fe-0/2/3.0) 27 pkts
  Explct route: 10.12.100.6 10.12.100.9
  Record route: <self> 10.12.100.6 10.12.100.9
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

### Sender-Only and Receiver-Only Sites

In Junos OS, you can configure a PE router to be a sender-site only or a receiver-site only. These options are enabled by including the **sender-site** and **receiver-site** statements at the **[edit routing-instances *routing-instance-name* protocols mvpn]** hierarchy level.

- A sender-site only PE router does not join the provider tunnels advertised by remote PE routers
- A receiver-site only PE router does not send a PMSI attribute

The commit fails if you include the **receiver-site** and **provider-tunnel** statements in the same VPN.

**Related Documentation**

- Generating Source AS and Route Target Import Communities on page 24
- *Understanding MBGP Multicast VPN Extranets*
- Signaling Provider Tunnels and Data Plane Setup on page 28
- Generating Next-Generation MVPN VRF Import and Export Policies on page 21

## Signaling Provider Tunnels and Data Plane Setup

In a next-generation multicast virtual private network (MVPN), provider tunnel information is communicated to the receiver PE routers in an out-of-band manner. This information is advertised via BGP and is independent of the actual tunnel signaling process. Once the tunnel is signaled, the sender PE router binds the VPN routing and forwarding (VRF) table to the locally configured tunnel. The receiver PE routers bind the tunnel signaled to the VRF table where the Type 1 autodiscovery route with the matching provider multicast service interface (PMSI) attribute is installed. The same binding process is used for both Protocol Independent Multicast (PIM) and RSVP-Traffic Engineering (RSVP-TE) signaled provider tunnels.

### Provider Tunnels Signaled by PIM (Inclusive)

A sender provider edge (PE) router configured to use an inclusive PIM-sparse mode (PIM-SM) any-source multicast (ASM ) provider tunnel for a VPN creates a multicast tree (using the P-group address configured) in the service provider network. This tree is rooted at the sender PE router and has the receiver PE routers as the leaves. VPN multicast packets received from the local VPN source are encapsulated by the sender PE router with a multicast generic routing encapsulation (GRE) header containing the P-group address configured for the VPN. These packets are then forwarded on the service provider network as normal IP multicast packets per normal P-PIM procedures. At the leaf nodes, the GRE header is stripped and the packets are passed on to the local VRF C-PIM protocol for further processing.

In Junos OS, a logical interface called multicast tunnel (MT) is used for GRE encapsulation and de-encapsulation of VPN multicast packets. The multicast tunnel interface is created automatically if a Tunnel PIC is present.

- Encapsulation subinterfaces are created from an mt-x/y/z.[32768-49151] range.

- De-encapsulation subinterfaces are created from an mt-x/y/z.[49152-65535] range.

The multicast tunnel subinterfaces act as pseudo upstream or downstream interfaces between C-PIM and P-PIM.

In the following two examples, assume that the network uses PIM-SM (ASM) signaled GRE tunnels as the tunneling technology. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show interfaces mt-0/1/0 terse** command to verify that Router PE1 has created the following multicast tunnel subinterface. The logical interface number is 32768, indicating that this sub-unit is used for GRE encapsulation.

```
user@PE1> show interfaces mt-0/1/0 terse
Interface      Admin   Link   Proto   Local   Remote
        mt-0/1/0           up          up
        mt-0/1/0.32768     up          up      inet
                                               inet6
```

Use the **show interfaces mt-0/1/0 terse** command to verify that Router PE2 has created the following multicast tunnel subinterface. The logical interface number is 49152, indicating that this sub-unit is used for GRE de-encapsulation.

```
user@PE2> show interfaces mt-0/1/0 terse
Interface      Admin  Link   Proto   Local   Remote
       mt-0/1/0           up          up
       mt-0/1/0.49152     up          up       inet
                                               inet6
```

## P-PIM and C-PIM on the Sender PE Router

The sender PE router installs a local join entry in its P-PIM database for each VRF table configured to use PIM as the provider tunnel. The outgoing interface list (OIL) of this entry points to the core-facing interface. Since the P-PIM entry is installed as **Local**, the sender PE router sets the source address to its primary loopback IP address.

Use the **show pim join extensive** command to verify that Router PE1 has installed the following state in its P-PIM database.

```
user@PE1>  show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source
    Keepalive timeout: 339
    Downstream neighbors:
        Interface: fe-0/2/3.0
            10.12.100.6 State: Join Flags: S Timeout: 195

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

On the VRF side of the sender PE router, C-PIM installs a **Local Source** entry in its C-PIM database for the active local VPN source. The OIL of this entry points to **Pseudo-MVPN**, indicating that the downstream interface points to the receivers in the next-generation MVPN network. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show pim join extensive instance vpna 224.1.1.1** command to verify that Router PE1 has installed the following entry in its C-PIM database.

```
user@PE1> show pim join extensive instance vpna 224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
    Source: 192.168.1.2
    Flags: sparse,spt
    Upstream interface: fe-0/2/0.0
    Upstream neighbor: 10.12.97.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 0
    Downstream neighbors:
        Interface: Pseudo-MVPN
```

The forwarding entry corresponding to the C-PIM **Local Source** (or **Local RP**) on the sender PE router points to the multicast tunnel encapsulation subinterface as the downstream interface. This indicates that the local multicast data packets are encapsulated as they are passed on to the P-PIM protocol.

Use the **show multicast route extensive instance vpna group 224.1.1.1** command to verify that Router PE1 has the following multicast forwarding entry for group 224.1.1.1. The upstream interface is the PE-CE interface and the downstream interface is the multicast tunnel encapsulation subinterface:

```
user@PE1> show multicast route extensive instance vpna group 224.1.1.1
Family: INET

Group: 224.1.1.1
    Source: 192.168.1.2/32
    Upstream interface: fe-0/2/0.0
    Downstream interface list:
        mt-0/1/0.32768
    Session description: ST Multicast Groups
    Statistics: 7 kBps, 79 pps, 719738 packets
    Next-hop ID: 262144
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
```

### P-PIM and C-PIM on the Receiver PE Router

On the receiver PE router, multicast data packets received from the network are de-encapsulated as they are passed through the multicast tunnel de-encapsulation interface.

 

The P-PIM database on the receiver PE router contains two P-joins. One is for P-RP, and the other is for the sender PE router. For both entries, the OIL contains the multicast tunnel de-encapsulation interface from which the GRE header is stripped. The upstream interface for P-joins is the core-facing interface that faces towards the sender PE router.

Use the **show pim join extensive** command to verify that Router PE3 has the following state in its P-PIM database. The downstream neighbor interface points to the GRE de-encapsulation subinterface:

```
user@PE3> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.1.1.10
    Flags: sparse,rptree,wildcard
    Upstream interface: so-0/0/3.0
    Upstream neighbor: 10.12.100.21
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: mt-1/2/0.49152
            10.12.53.13 State: Join Flags: SRW Timeout: Infinity

Group: 239.1.1.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: so-0/0/3.0
    Upstream neighbor: 10.12.100.21
    Upstream state: Join to Source
    Keepalive timeout: 351
    Downstream neighbors:
        Interface: mt-1/2/0.49152
            10.12.53.13 State: Join Flags: S Timeout: Infinity

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

On the VRF side of the receiver PE router, C-PIM installs a join entry in its C-PIM database. The OIL of this entry points to the local VPN interface, indicating active local receivers. The upstream protocol, interface, and neighbor of this entry point to the next-generation-MVPN network. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show pim join extensive instance vpna 224.1.1.1** command to verify that Router PE3 has the following state in its C-PIM database:

```
user@PE3> show pim join extensive instance vpna 224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
    Source: *
    RP: 10.12.53.1
    Flags: sparse,rptree,wildcard
    Upstream protocol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to RP
```

```
    Downstream neighbors:
        Interface: so-0/2/0.0
            10.12.87.1 State: Join Flags: SRW Timeout: Infinity

Group: 224.1.1.1
    Source: 192.168.1.2
    Flags: sparse
    Upstream protocol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to Source
    Keepalive timeout:
    Downstream neighbors:
        Interface: so-0/2/0.0
            10.12.87.1 State: Join Flags: S Timeout: 195

Instance: PIM.vpna Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

The forwarding entry corresponding to the C-PIM entry on the receiver PE router uses the multicast tunnel de-encapsulation subinterface as the upstream interface.

Use the **show multicast route extensive instance vpna group 224.1.1.1** command to verify that Router PE3 has installed the following multicast forwarding entry for the local receiver:

```
user@PE3> show multicast route extensive instance vpna group 224.1.1.1
Family: INET

Group: 224.1.1.1
    Source: 192.168.1.2/32
    Upstream interface: mt-1/2/0.49152
    Downstream interface list:
        so-0/2/0.0
    Session description: ST Multicast Groups
    Statistics: 1 kBps, 10 pps, 149 packets
    Next-hop ID: 262144
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
```

## Provider Tunnels Signaled by RSVP-TE (Inclusive and Selective)

Junos OS supports signaling both inclusive and selective provider tunnels by RSVP-TE point-to-multipoint label-switched paths (LSPs). You can configure a combination of inclusive and selective provider tunnels per VPN.

- If you configure a VPN to use an inclusive provider tunnel, the sender PE router signals one point-to-multipoint LSP for the VPN.
- If you configure a VPN to use selective provider tunnels, the sender PE router signals a point-to-multipoint LSP for each selective tunnel configured.

Sender (ingress) PE routers and receiver (egress) PE routers play different roles in the point-to-multipoint LSP setup. Sender PE routers are mainly responsible for initiating the parent point-to-multipoint LSP and the sub-LSPs associated with it. Receiver PE

 

routers are responsible for setting up state such that they can forward packets received over a sub-LSP to the correct VRF table (binding a provider tunnel to the VRF).

### Inclusive Tunnels: Ingress PE Router Point-to-Multipoint LSP Setup

The point-to-multipoint LSP and associated sub-LSPs are signaled by the ingress PE router. The information about the point-to-multipoint LSP is advertised to egress PE routers in the PMSI attribute via BGP.

The ingress PE router signals point-to-multipoint sub-LSPs by originating point-to-multipoint RSVP path messages toward egress PE routers. The ingress PE router learns the identity of the egress PE routers from Type 1 routes installed in its **<routing-instance-name>.mvpn.0** table. Each RSVP path message carries an **S2L_Sub_LSP** object along with the point-to-multipoint session object. The **S2L_Sub_LSP** object carries a 4-byte sub-LSP destination (egress) IP address.

In Junos OS, sub-LSPs associated with a point-to-multipoint LSP can be signaled automatically by the system or via a static sub-LSP configuration. When they are automatically signaled, the system chooses a name for the point-to-multipoint LSP and each sub-LSP associated with it using the following naming convention.

Point-to-multipoint LSPs naming convention:

<ingress PE rid>:<a per VRF unique number>:mvpn:<routing-instance-name>

Sub-LSPs naming convention:

<egress PE rid>:<ingress PE rid>:<a per VRF unique number>:mvpn:<routing-instance-name>

Use the **show mpls lsp p2mp** command to verify that the following LSPs have been created by Router PE1:

Parent P2MP LSP: 10.1.1.1:65535:mvpn:vpna

Sub-LSPs: 10.1.1.2:10.1.1.1:65535:mvpn:vpna (Router PE1 to Router PE2) and

10.1.1.3:10.1.1.1:65535:mvpn:vpna (Router PE1 to Router PE3)

```
user@PE1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2
To              From            State Rt P  ActivePath   LSPname
10.1.1.2        10.1.1.1        Up     0 *
10.1.1.2:10.1.1.1:65535:mvpn:vpna
10.1.1.3        10.1.1.1        Up     0 *
10.1.1.3:10.1.1.1:65535:mvpn:vpna
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

The values in this example are as follows:

- I-PMSI P2MP LSP name: 10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE2): 10.1.1.2:10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mvpn:vpna

### Inclusive Tunnels: Egress PE Router Point-to-Multipoint LSP Setup

An egress PE router responds to an RSVP path message by originating an RSVP reservation (RESV) message per normal RSVP procedures. The RESV message contains the MPLS label allocated by the egress PE router for this sub-LSP and is forwarded hop by hop toward the ingress PE router, thus setting up state on the network. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show rsvp session** command to verify that Router PE2 has assigned label **299840** for the sub-LSP **10.1.1.2:10.1.1.1:65535:mvpn:vpna**:

```
user@PE2> show rsvp session
Total 0 displayed, Up 0, Down 0
Egress RSVP: 1 sessions
To              From      State     Rt Style     Labelin    Labelout    LSPname
10.1.1.2        10.1.1.1  Up  0 1 SE  299840          -
10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Use the **show mpls lsp p2mp** command to verify that Router PE3 has assigned label **16** for the sub-LSP **10.1.1.3:10.1.1.1:65535:mvpn:vpna**:

```
user@PE3> show mpls lsp p2mp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 1 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 1
To              From      State     Rt Style     Labelin    Labelout    LSPname
10.1.1.3        10.1.1.1 Up 0 1 SE       16           -
10.1.1.3:10.1.1.1:65535:mvpn:vpna
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

### Inclusive Tunnels: Egress PE Router Data Plane Setup

The egress PE router installs a forwarding entry in its **mpls** table for the label it allocated for the sub-LSP. The MPLS label is installed with a pop operation (a pop operation removes the top MPLS label), and the packet is passed on to the VRF table for a second route lookup. The second lookup on the egress PE router is necessary for the VPN multicast data packets to be processed inside the VRF table using normal C-PIM procedures.

 

Use the **show route table mpls label 16** command to verify that Router PE3 has installed the following label entry in its MPLS forwarding table:

```
user@PE3> show route table mpls label 16
+ = Active Route, - = Last Active, * = Both

16                 *[VPN/0] 03:03:17
                    to table vpna.inet.0, Pop
```

In Junos OS, VPN multicast routing entries are stored in the **<routing-instance-name>.inet.1** table, which is where the second route lookup occurs. In the example above, even though **vpna.inet.0** is listed as the routing table where the second lookup happens after the pop operation, internally the lookup is pointed to the **vpna.inet.1** table. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show route table vpna.inet.1** command to verify that Router PE3 contains the following entry in its VPN multicast routing table:

```
user@PE3> show route table vpna.inet.1
vpna.inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

224.1.1.1,192.168.1.2/32*[MVPN/70] 00:04:10
                         Multicast (IPv4)
```

Use the **show multicast route extensive instance vpna** command to verify that Router PE3 contains the following VPN multicast forwarding entry corresponding to the multicast routing entry for the Llocal join. The upstream interface points to **lsi.0** and the downstream interface (OIL) points to the **so-0/2/0.0** interface (toward local receivers). The **Upstream protocol** value is **MVPN** because the VPN multicast source is reachable via the next-generation MVPN network. The **lsi.0** interface is similar to the multicast tunnel interface used when PIM-based provider tunnels are used. The **lsi.0** interface is used for removing the top MPLS header.

```
user@PE3> show multicast route extensive instance vpna
Family: INET

Group: 224.1.1.1
    Source: 192.168.1.2/32
    Upstream interface: lsi.0
    Downstream interface list:
        so-0/2/0.0
    Session description: ST Multicast Groups
    Statistics: 1 kBps, 10 pps, 3472 packets
    Next-hop ID: 262144
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

The requirement for a double route lookup on the VPN packet header requires two additional configuration statements on the egress PE routers when provider tunnels are signaled by RSVP-TE.

First, since the top MPLS label used for the point-to-multipoint sub-LSP is actually tied to the VRF table on the egress PE routers, the penultimate-hop popping (PHP) operation is not used for next-generation MVPNs. Only ultimate-hop popping is used. PHP allows the penultimate router (router before the egress PE router) to remove the top MPLS label. PHP works well for VPN unicast data packets because they typically carry two MPLS labels: one for the VPN and one for the transport LSP.

After the LSP label is removed, unicast VPN packets still have a VPN label that can be used for determining the VPN to which the packets belong. VPN multicast data packets, on the other hand, carry only one MPLS label that is directly tied to the VPN. Therefore, the MPLS label carried by VPN multicast packets must be preserved until the packets reach the egress PE router. Normally, PHP must be disabled through manual configuration.

To simplify the configuration, PHP is disabled by default on Juniper Networks PE routers when you include the **mvpn** statement at the **[edit routing-instances *routing-interface-name* interface]** hierarchy level. PHP is also disabled by default when you include the **vrf-table-label** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level.

Second, in Junos OS, VPN labels associated with a VRF table can be allocated in two ways.

- Allocate a unique label for each VPN next hop (PE-CE interface). This is the default behavior.
- Allocate one label for the entire VRF table, which requires additional configuration. Only allocating a label for the entire VRF table allows a second lookup on the VPN packet's header. Therefore, PE routers supporting next-generation-MVPN services must be configured to allocate labels for the VRF table. There are two ways to do this as shown in Figure 11 on page 37.
  - One is by including a virtuall tunnel interface named **vt** at the **[edit routing-instances *routing-instance-name* interfaces]** hierarchy level, which requires a Tunnel PIC.
  - The second is by including the **vrf-table-label** statement at the **[routing-instances *routing-instance-name*]** hierarchy level, which does not require a Tunnel PIC.

Both of these options enable an egress PE router to perform two route lookups. However, there are some differences in the way in which the second lookup is done

If the **vt** interface is used, the allocated label is installed in the **mpls** table with a **pop** operation and a forwarding next hop pointing to the **vt** interface.

 

Figure 11: Enabling Double Route Lookup on VPN Packet Headers

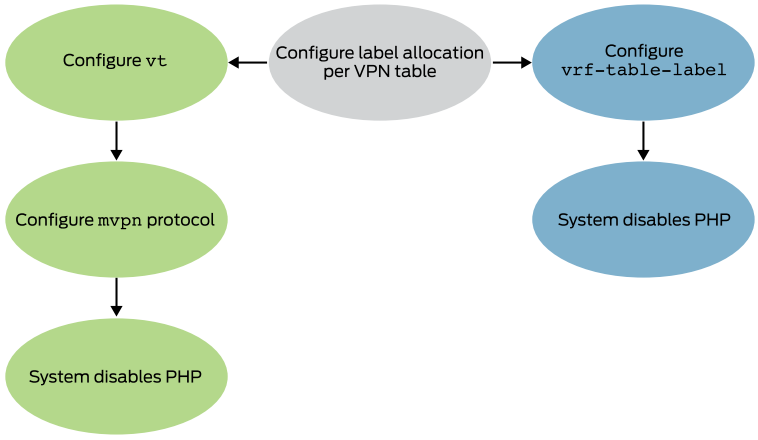


Use the **show route table mpls label 299840** command to verify that Router PE2 has installed the following entry and uses a **vt** interface in the **mpls** table. The label associated with the point-to-multipoint sub-LSP (**299840**) is installed with a pop and a forward operation with the **vt-0/1/0.0** interface being the next hop. VPN multicast packets received from the core exit the **vt-0/1/0.0** interface without their MPLS header, and the egress Router PE2 does a second lookup on the packet header in the **vpna.inet.1** table.

```
user@PE2> show route table mpls label 299840
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299840             *[VPN/0] 00:00:22
                    > via vt-0/1/0.0, Pop
```

If the **vrf-table-label** is configured, the allocated label is installed in the **mpls** table with a pop operation, and the forwarding entry points to the **<routing-instance-name>.inet.0** table (which internally triggers the second lookup to be done in the **<routing-instance-name>.inet.1** table).

Use the **show route table mpls label 16** command to verify that Router PE3 has installed the following entry in its **mpls** table and uses the **vrf-table-label** statement:

```
user@PE3> show route table mpls label 16
mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

16                 *[VPN/0] 03:03:17
                      to table vpna.inet.0, Pop
```

Configuring label allocation for each VRF table affects both unicast VPN and MVPN routes. However, you can enable per-VRF label allocation for MVPN routes only if per-VRF allocation is configured via **vt**. This feature is configured via multicast and unicast keywords at the **[edit routing-instances *routing-instance-name* interface vt-x/y/z.0]** hierarchy level.

Note that including the **vrf-table-label** statement enables per-VRF label allocation for both unicast and MVPN routes and cannot be turned off for either type of routes (it is either on or off for both).

If a PE router is a bud router, meaning it has local receivers and also forwards MPLS packets received over a point-to-multipoint LSP downstream to other P and PE routers, then there is a difference in how the **vrf-table-label** and **vt** statements work. When, the **vrf-table-label** statement is included, the bud PE router receives two copies of the packet from the penultimate router: one to be forwarded to local receivers and the other to be forwarded to downstream P and PE routers. When the **vt** statement is included, the PE router receives a single copy of the packet.

### Inclusive Tunnels: Ingress and Branch PE Router Data Plane Setup

On the ingress PE router, local VPN data packets are encapsulated with the MPLS label received from the network for sub-LSPs.

Use the **show rsvp session** command to verify that on the ingress Router PE1, VPN multicast data packets are encapsulated with MPLS label **300016** (advertised by Router P1 per normal RSVP RESV procedures) and forwarded toward Router P1 down the sub-LSPs **10.1.1.3:10.1.1.1:65535:mvpn:vpna** and **10.1.1.2:10.1.1.1:65535:mvpn:vpna**.

```
user@PE1> show rsvp session
Ingress RSVP: 2 sessions
To          From        State     Rt Style   Labelin   Labelout   LSPname
10.1.1.3   10.1.1.1 Up 0 1 SE         -      300016
10.1.1.3:10.1.1.1:65535:mvpn:vpna
10.1.1.2 10.1.1.1   Up 0 1 SE         -      300016
10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

RFC 4875 describes a branch node as “an LSR that replicates the incoming data on to one or more outgoing interfaces.” On a branch Rrouter, the incoming data carrying an MPLS label is replicated onto one or more outgoing interfaces that can use different MPLS labels. Branch nodes keep track of incoming and outgoing labels associated with point-to-multipoint LSPs. Routers referenced in this topic are shown in “Next-Generation MVPN Network Topology” on page 2.

Use the **show rsvp session** command to verify that branch node P1 has the incoming label **300016** and outgoing labels **16** for sub-LSP **10.1.1.3:10.1.1.1:65535:mvpn:vpna** (to Router PE3) and **299840** for sub-LSP **10.1.1.2:10.1.1.1:65535:mvpn:vpna** (to Router PE2).

```
user@P1> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 2 sessions
To          From        State     Rt Style   Labelin   Labelout   LSPname
10.1.1.3 10.1.1.1 Up    0 1 SE     300016       16
10.1.1.3:10.1.1.1:65535:mvpn:vpna
10.1.1.2 10.1.1.1 Up    0 1 SE     300016   299840
```

```
10.1.1.2:10.1.1.1:65535:mvpn:vpna
Total 2 displayed, Up 2, Down 0
```

Use the **show route table mpls label 300016** command to verify that the corresponding forwarding entry on Router P1 shows that the packets coming in with one MPLS label (**300016**) are swapped with labels **16** and **299840** and forwarded out through their respective interfaces (**so-0/0/3.0** and **so-0/0/1.0** respectively toward Router PE2 and Router PE3).

```
user@P1> show route table mpls label 300016
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

300016                *[RSVP/7] 01:58:15, metric 1
                       > via so-0/0/3.0, Swap 16
                         via so-0/0/1.0, Swap 299840
```

## Selective Tunnels: Type 3 S-PMSI Autodiscovery and Type 4 Leaf Autodiscovery Routes

Selective provider tunnels are configured by including the **selective** statement at the **[edit routing-instances *routing-instance-name* provider-tunnel]** hierarchy level. You can configure a threshold to trigger the signaling of a selective provider tunnel. Including the **selective** statement triggers the following events.

First, the ingress PE router originates a Type 3 S-PMSI autodiscovery route. The S-PMSI autodiscovery route contains the route distinguisher of the VPN where the tunnel is configured and the (C-S, C-G) pair that uses the selective provider tunnel.

In this section assume that Router PE1 is signaling a selective tunnel for (**192.168.1.2, 224.1.1.1**) and Router PE3 has an active receiver.

Use the **show route table vpna.mvpn.0 | find 3:** command to verify that Router PE1 has installed the following Type 3 route after the selective provider tunnel is configured:

```
user@PE1> show route table vpna.mvpn.0 | find 3:
3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1/240
                          *[MVPN/70] 00:05:07, metric2 1
                             Indirect
```

Second, the ingress PE router attaches a PMSI attribute to a Type 3 route. This PMSI attribute is similar to the PMSI attribute advertised for inclusive provider tunnels with one difference: the PMSI attribute carried with Type 3 routes has its **Flags** bit set to **Leaf Information Required**. This means that the sender PE router is requesting receiver PE routers to send a Type 4 route if they have active receivers for the (C-S, C-G) carried in the Type 3 route. Also, remember that for each selective provider tunnel, a new point-to-multipoint and associated sub-LSPs are signaled. The PMSI attribute of a Type 3 route carries information about the new point-to-multipoint LSP.

Use the **show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn | find 3:** command to verify that Router PE1 advertises the following Type 3 route and the **PMSI** attribute. The point-to-multipoint session object included in the **PMSI** attribute has a

different port number (**29499**) than the one used for the inclusive tunnel (**6574**) indicating that this is a new point-to-multipoint tunnel.

```
user@PE1> show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn | find 3:
* 3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1/240 (1 entry, 1 announced)
BGP group int type Internal
    Route Distinguisher: 10.1.1.1:1
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:10:1
    PMSI: Flags 1:RSVP-TE:label[0:0:0]:Session_13[10.1.1.1:0:29499:10.1.1.1]
```

Egress PE routers with active receivers should respond to a Type 3 route by originating a Type 4 leaf autodiscovery route. A leaf autodiscovery route contains a route key and the originating router's IP address fields. The **Route Key** field of the leaf autodiscovery route contains the original Type 3 route that is received. The originating router's IP address field is set to the router ID of the PE router originating the leaf autodiscovery route.

The ingress PE router adds each egress PE router that originated the leaf autodiscovery route as a leaf (destination of the sub-LSP for the selective point-to-multipoint LSP). Similarly, the egress PE router that originated the leaf autodiscovery route sets up forwarding state to start receiving data through the selective provider tunnel.

Egress PE routers advertise Type 4 routes with a route target that is specific to the PE router signaling the selective provider tunnel. This route target is in the form of target:<rid of the sender PE>:0. The sender PE router (the PE router signaling the selective provider tunnel) applies a special internal import policy to Type 4 routes that looks for a route target with its own router ID. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

Use the **show route table vpna.mvpn | find 4:3:** command to verify that Router PE3 originates the following Type 4 route. The local Type 4 route is installed by the MVPN module.

```
user@PE3> show route table vpna.mvpn | find 4:3:
4:3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1:1.1.1.3/240
                           *[MVPN/70] 00:15:29, metric2 1
                              Indirect
```

Use the **show route advertising-protocol bgp 10.1.1.1 table vpna.mvpn detail | find 4:3:** command to verify that Router PE3 has advertised the local Type 4 route with the following route target community. This route target carries the IP address of the sender PE router (**10.1.1.1**) followed by a 0.

```
user@PE3> show route advertising-protocol bgp 10.1.1.1 table vpna.mvpn detail | find 4:3:
* 4:3:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1:10.1.1.1:10.1.1.3/240 (1 entry, 1
announced)
BGP group int type Internal
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:10.1.1.1:0
```

 

Use the **show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__** command to verify that Router PE1 (the PE router signaling the selective provider tunnel) has applied the following import policy to Type 4 routes. The routes are accepted if their route target matches **target:10.1.1.1:0**.

```
user@PE1> show policy __vrf-mvpn-import-cmcast-leafAD-global-internal__
Policy __vrf-mvpn-import-cmcast-leafAD-global-internal__:
Term unnamed:
from community __vrf-mvpn-community-rt_import-target-global-internal__
[target:10.1.1.1:0 ]
then accept
Term unnamed:
then reject
```

For each selective provider tunnel configured, a Type 3 route is advertised and a new point-to-multipoint LSP is signaled. Point-to-multipoint LSPs created by Junos OS for selective provider tunnels are named using the following naming conventions:

- Selective point-to-multipoint LSPs naming convention:

  <ingress PE rid>:<a per VRF unique number>:mv<a unique number>:<routing-instance-name>

- Selective point-to-multipoint sub-LSP naming convention:

  <egress PE rid>:<ingress PE rid>:<a per VRF unique>:mv<a unique number>:<routing-instance-name>

Use the **show mpls lsp p2mp** command to verify that Router PE1 signals point-to-multipoint LSP **10.1.1.1:65535:mv5:vpna** with one sub-LSP **10.1.1.3:10.1.1.1:65535:mv5:vpna**. The first point-to-multipoint LSP **10.1.1.1:65535:mvpn:vpna** is the LSP created for the inclusive tunnel.

```
user@PE1> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: 10.1.1.1:65535:mvpn:vpna, P2MP branch count: 2
To         From          State       Rt P        ActivePath            LSPname
10.1.1.3 10.1.1.1 Up 0 *                           10.1.1.3:10.1.1.1:65535:mvpn
:vpna
10.1.1.2 10.1.1.1 Up 0 *                           10.1.1.2:10.1.1.1:65535:mvpn
:vpna
P2MP name: 10.1.1.1:65535:mv5:vpna, P2MP branch count: 1
To         From          State       Rt P        ActivePath            LSPname
10.1.1.3 10.1.1.1 Up 0 *                            10.1.1.3:10.1.1.1:65535:mv5
:vpna
Total 3 displayed, Up 3, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

The values in this example are as follows.

- I-PMSI P2MP LSP name: 10.1.1.1:65535:mvpn:vpna
- I-PMSI P2MP sub-LSP name (to PE2): 10.1.1.2:10.1.1.1:65535:mvpn:vpna

- I-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mvpn:vpna
- S-PMSI P2MP LSP name: 10.1.1.1:65535:mv5:vpna
- S-PMSI P2MP sub-LSP name (to PE3): 10.1.1.3:10.1.1.1:65535:mv5:vpna

**Related Documentation**

- Next-Generation MVPN Data Plane on page 15
- Originating Type 1 Intra-AS Autodiscovery Routes on page 24
- Exchanging C-Multicast Routes on page 42

## Exchanging C-Multicast Routes

This section describes PE-PE distribution of Type 7 routes discussed in "Signaling Provider Tunnels and Data Plane Setup" on page 28.

In source-tree-only mode, a receiver provider edge (PE) router generates and installs a Type 6 route in its **<routing-instance-name>.mvpn.0** table in response to receiving a (C-*, C-G) message from a local receiver, but does not advertise this route to other PE routers via BGP. The receiver PE router waits for a Type 5 route corresponding to the C-join.

Type 5 routes carry information about active sources and can be advertised by any PE router. In Junos OS, a PE router originates a Type 5 route if one of the following conditions occurs:

- PE router starts receiving multicast data directly from a VPN multicast source.
- PE router is the candidate rendezvous point (router) (candidate RP) and starts receiving C-PIM register messages.
- PE router has a Multicast Source Discovery Protocol (MSDP) session with the candidate RP and starts receiving MSDP Source Active routes.

Once both Type 6 and Type 5 routes are installed in the **<routing-instance-name>.mvpn.0** table, the receiver PE router is ready to originate a Type 7 route

### Advertising C-Multicast Routes Using BGP

If the C-join received over a VPN interface is a source tree join (C-S, C-G), then the receiver PE router simply originates a Type 7 route (Step 7 in the following procedure). If the C-join is a shared tree join (C-*, C-G), then the receiver PE router needs to go through a few steps (Steps 1-7) before originating a Type 7 route.

Note that Router PE1 is the candidate RP that is conveniently located in the same router as the sender PE router. If the sender PE router and the PE router acting as (or MSDP peering with) the candidate RP are different, then the VPN multicast register messages first need to be delivered to the PE router acting as the candidate RP that is responsible for originating the Type 5 route. Routers referenced in this topic are shown in "Next-Generation MVPN Network Topology" on page 2.

1. A PE router that receives a (C-*, C-G) join message processes the message using normal C-PIM procedures and updates its C-PIM database accordingly.

   Enter the **show pim join extensive instance vpna 224.1.1.1** command on Router PE3 to verify that Router PE3 creates the C-PIM database after receiving the (*, 224.1.1.1) C-join message from Router CE3:

```
user@PE3> show pim join extensive instance vpna 224.1.1.1
Instance: PIM.vpna Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
    Source: *
    RP: 10.12.53.1
    Flags: sparse,rptree,wildcard
    Upstream protocol: BGP
  Upstream interface: Through BGP
  Upstream neighbor: Through MVPN
  Upstream state: Join to RP
  Downstream neighbors:
      Interface: so-0/2/0.0
            10.12.87.1 State: Join Flags: SRW Timeout: Infinity
```

2. The (C-*, C-G) entry in the C-PIM database triggers the generation of a Type 6 route that is then installed in the **<routing-instance-name>.mvpn.0** table by C-PIM. The Type 6 route uses the candidate RP IP address as the source.

   Enter the **show route table vpna.mvpn.0 detail | find 6:10.1.1.1** command on Router PE3 to verify that Router PE3 installs the following Type 6 route in the **vpna.mvpn.0** table:

```
user@PE3> show route table vpna.mvpn.0 detail | find 6:10.1.1.1
6:10.1.1.1:1:65000:32:10.12.53.1:32:224.1.1.1/240 (1 entry, 1 announced)
               *PIM     Preference: 105
          Next hop type: Multicast (IPv4), Next hop index: 262144
          Next-hop reference count: 11
          State: <Active Int>
          Age: 1d 1:32:58
          Task: PIM.vpna
          Announcement bits (2): 0-PIM.vpna 1-mvpn global task
            AS path: I
          Communities: no-advertise target:10.1.1.1:64
```

3. The route distinguisher and route target attached to the Type 6 route are learned from a route lookup in the **<routing-instance-name>.inet.0** table for the IP address of the candidate RP.

   Enter the **show route table vpna.inet.0 10.12.53.1 detail** command on Router PE3 to verify that Router PE3 has the following entry for **C-RP 10.12.53.1** in the **vpna.inet.0** table:

```
user@PE3> show route table vpna.inet.0 10.12.53.1 detail
vpna.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
10.12.53.1/32 (1 entry, 1 announced)
        *BGP     Preference: 170/-101
        Route Distinguisher: 10.1.1.1:1
        Next hop type: Indirect
        Next-hop reference count: 6
        Source: 10.1.1.1
        Next hop type: Router, Next hop index: 588
        Next hop: via so-0/0/3.0, selected
```

```
            Label operation: Push 16, Push 299808(top)
            Protocol next hop: 10.1.1.1
            Push 16
                     Indirect next hop: 8da91f8 262143
                     State: <Secondary Active Int Ext>
                     Local AS: 65000 Peer AS: 65000
                     Age: 4:49:25 Metric2: 1
                     Task: BGP_65000.10.1.1.1+179
                     Announcement bits (1): 0-KRT
                     AS path: I
                    Communities: target:10:1 src-as:65000:0 rt-import:10.1.1.1:64

                     Import Accepted
                     VPN Label: 16
                     Localpref: 100
                     Router ID: 10.1.1.1
                Primary Routing Table bgp.l3vpn.0
```

4. After the VPN source starts transmitting data, the first PE router that becomes aware of the active source (either by receiving register messages or the MSDP source-active routes) installs a Type 5 route in its **VRF mvpn** table.

   Enter the **show route table vpna.mvpn.0 detail | find 5:10.1.1.1** command on Router PE1 to verify that Router PE1 has installed the following entry in the **vpna.mvpn.0** table and starts receiving C-PIM register messages from Router CE1:

```
user@PE1> show route table vpna.mvpn.0 detail | find 5:10.1.1.1
5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
         *PIM    Preference: 105
                 Next hop type: Multicast (IPv4)
                 Next-hop reference count: 30
                 State: <Active Int>
                 Age: 1d 1:36:33
                 Task: PIM.vpna
                 Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP
RT Background
                 AS path: I
```

5. Type 5 routes that are installed in the **<routing-instance-name>.mvpn.0** table are picked up by BGP and advertised to remote PE routers.

   Enter the **show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn.0 | find 5:** command on Router PE1 to verify that Router PE1 advertises the following Type 5 route to remote PE routers:

```
user@PE1> show route advertising-protocol bgp 10.1.1.3 detail table vpna.mvpn.0 | find 5:
* 5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
         BGP group int type Internal
             Route Distinguisher: 10.1.1.1:1
             Nexthop: Self
             Flags: Nexthop Change
             Localpref: 100
             AS path: [65000] I
             Communities: target:10:1
```

6. The receiver PE router that has both a Type 5 and Type 6 route for (C-*, C-G) is now ready to originate a Type 7 route.

   Enter the **show route table vpna.mvpn.0 detail** command on Router PE3 to verify that Router PE3 has the following Type 5, 6, and 7 routes in the **vpna.mvpn.0** table.

The Type 6 route is installed by C-PIM in Step 2. The Type 5 route is learned via BGP in Step 5. The Type 7 route is originated by the MVPN module in response to having both Type 5 and Type 6 routes for the same (C-*, C-G). The route target of the Type 7 route is the same as the route target of the Type 6 route because both routes (IP address of the candidate RP [10.12.53.1] and the address of the VPN multicast source [192.168.1.2]) are reachable via the same router [PE1]). Therefore, **10.12.53.1** and **192.168.1.2** carry the same route target import (**10.1.1.1:64**) community

```
user@PE3> show route table vpna.mvpn.0 detail
5:10.1.1.1:1:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
         *BGP     Preference: 170/-101
                  Next hop type: Indirect
                  Next-hop reference count: 4
                  Source: 10.1.1.1
                  Protocol next hop: 10.1.1.1
                  Indirect next hop: 2 no-forward
                  State: <Secondary Active Int Ext>
                  Local AS: 65000 Peer AS: 65000
                  Age: 1d 1:43:13 Metric2: 1
                  Task: BGP_65000.10.1.1.1+55384
                  Announcement bits (2): 0-PIM.vpna 1-mvpn global task
                  AS path: I
                  Communities: target:10:1
                  Import Accepted
                  Localpref: 100
                  Router ID: 10.1.1.1
                  Primary Routing Table bgp.mvpn.0

6:10.1.1.1:1:65000:32:10.12.53.1:32:224.1.1.1/240 (1 entry, 1 announced)
       *PIM      Preference: 105
                 Next hop type: Multicast (IPv4), Next hop index: 262144
                 Next-hop reference count: 11
                 State: <Active Int>
                 Age: 1d 1:44:09
                 Task: PIM.vpna
                 Announcement bits (2): 0-PIM.vpna 1-mvpn global task
                 AS path: I
                 Communities: no-advertise target:10.1.1.1:64

7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
       *MVPN         Preference: 70
                Next hop type: Multicast (IPv4), Next hop index: 262144
                Next-hop reference count: 11
                State: <Active Int Ext>
                Age: 1d 1:44:09 Metric2: 1
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP RT
 Background
                AS path: I
                Communities: target:10.1.1.1:64
```

7. The Type 7 route installed in the VRF MVPN table is picked up by BGP and advertised to remote PE routers.

   Enter the **show route advertising-protocol bgp 10.1.1.1 detail table vpna.mvpn.0 | find 7:10.1.1.1** command on Router PE3 to verify that Router PE3 advertises the following Type 7 route:

```
user@PE3> show route advertising-protocol bgp 10.1.1.1 detail table vpna.mvpn.0 | find 7:10.1.1.1
```

```
* 7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (1 entry, 1 announced)
         BGP group int type Internal
             Route Distinguisher: 10.1.1.3:1
             Nexthop: Self
             Flags: Nexthop Change
             Localpref: 100
             AS path: [65000] I
             Communities: target:10.1.1.1:64
```

8. If the C-join is a source tree join, then the Type 7 route is originated immediately (without waiting for a Type 5 route).

   Enter the **show route table vpna.mvpn.0 detail | find 7:10.1.1.1** command on Router PE2 to verify that Router PE2 originates the following Type 7 route in response to receiving a (**192.168.1.2, 232.1.1.1**) C-join:

```
user@PE2> show route table vpna.mvpn.0 detail | find 7:10.1.1.1
7:10.1.1.1:1:65000:32:192.168.1.2:32:232.1.1.1/240 (1 entry, 1 announced)
          *PIM    Preference: 105
                  Next hop type: Multicast (IPv4), Next hop index: 262146
                  Next-hop reference count: 4
                  State: <Active Int>
                  Age: 2d 18:59:56
                  Task: PIM.vpna
                  Announcement bits (3): 0-PIM.vpna 1-mvpn global task 2-BGP
RT Background
                  AS path: I
                  Communities: target:10.1.1.1:64
```

## Receiving C-Multicast Routes

A sender PE router imports a Type 7 route if the route is carrying a route target that matches the locally originated route target import community. All Type 7 routes must pass the **__vrf-mvpn-import-cmcast-<routing-instance-name>-internal__** policy in order to be installed in the **<routing-instance-name>.mvpn.0** table.

When a sender PE router receives a Type 7 route via BGP, this route is installed in the **<routing-instance-name>.mvpn.0** table. The BGP route is then translated back into a normal C-join inside the VRF table, and the C-join is installed in the local C-PIM database of the receiver PE router. A new C-join added to the C-PIM database triggers C-PIM to originate a Type 6 or Type 7 route. The C-PIM on the sender PE router creates its own version of the same Type 7 route received via BGP.

Use the **show route table vpna.mvpn.0 detail | find 7:10.1.1.1** command to verify that Router PE1 contains the following entries for a Type 7 route in the **vpna.mvpn.0** table corresponding to a (**192.168.1.2, 224.1.1.1**) join message. There are two entries; one entry is installed by PIM and the other entry is installed by BGP. This example also shows the Type 7 route corresponding to the (**192.168.1.2, 232.1.1.1**) join.

```
user@PE1> show route table vpna.mvpn.0 detail | find 7:10.1.1.1
7:10.1.1.1:1:65000:32:192.168.1.2:32:224.1.1.1/240 (2 entries, 2 announced)
*PIM     Preference: 105
         Next hop type: Multicast (IPv4)
         Next-hop reference count: 30
         State: <Active Int>
         Age: 1d 2:19:04
         Task: PIM.vpna
```

```
         Announcement bits (2): 0-PIM.vpna 1-mvpn global task
         AS path: I
         Communities: no-advertise target:10.1.1.1:64
BGP      Preference: 170/-101
         Next hop type: Indirect
         Next-hop reference count: 4
         Source: 10.1.1.3
         Protocol next hop: 10.1.1.3
         Indirect next hop: 2 no-forward
         State: <Secondary Int Ext>
         Inactive reason: Route Preference
         Local AS: 65000 Peer AS: 65000
         Age: 53:27      Metric2: 1
         Task: BGP_65000.10.1.1.3+179
         Announcement bits (2): 0-PIM.vpna 1-mvpn global task
         AS path: I
         Communities: target:10.1.1.1:64
         Import Accepted
         Localpref: 100
         Router ID: 10.1.1.3
         Primary Routing Table bgp.mvpn.0
7:10.1.1.1:1:65000:32:192.168.1.2:32:232.1.1.1/240 (2 entries, 2 announced)
*PIM     Preference: 105
         Next hop type: Multicast (IPv4)
         Next-hop reference count: 30
         State: <Active Int>
         Age: 2d 19:21:17
         Task: PIM.vpna
         Announcement bits (2): 0-PIM.vpna 1-mvpn global task
         AS path: I
         Communities: no-advertise target:10.1.1.1:64
BGP      Preference: 170/-101
         Next hop type: Indirect
         Next-hop reference count: 4
         Source: 10.1.1.2
         Protocol next hop: 10.1.1.2
         Indirect next hop: 2 no-forward
         State: <Secondary Int Ext>
         Inactive reason: Route Preference
         Local AS: 65000 Peer AS: 65000
         Age: 53:27      Metric2: 1
         Task: BGP_65000.10.1.1.2+49165
         Announcement bits (2): 0-PIM.vpna 1-mvpn global task
         AS path: I
         Communities: target:10.1.1.1:64
         Import Accepted
         Localpref: 100
         Router ID: 10.1.1.2
         Primary Routing Table bgp.mvpn.0
```

Remote C-joins (Type 7 routes learned via BGP translated back to normal C-joins) are installed in the VRF C-PIM database on the sender PE router and are processed based on regular C-PIM procedures. This process completes the end-to-end C-multicast routing exchange.

Use the **show pim join extensive instance vpna** command to verify that Router PE1 has installed the following entries in the C-PIM database:

```
user@PE1> show pim join extensive instance vpna
```

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

Group: 224.1.1.1
    Source: 192.168.1.2
    Flags: sparse,spt
    Upstream interface: fe-0/2/0.0
    Upstream neighbor: 10.12.97.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 201
    Downstream neighbors:
    Interface: Pseudo-MVPN

Group: 232.1.1.1
    Source: 192.168.1.2
    Flags: sparse,spt
    Upstream interface: fe-0/2/0.0
    Upstream neighbor: 10.12.97.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout:
    Downstream neighbors:
    Interface: Pseudo-MVPN

Instance: PIM.vpna Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

**Related Documentation**

- Signaling Provider Tunnels and Data Plane Setup on page 28
- Distributing C-Multicast Routes on page 10
- *Understanding MBGP Multicast VPN Extranets*

## Conclusion

Next-generation multicast virtual private networks (NG-MVPNs) gives service providers a new way of offering multicast VPN service. The strength of next-generation MVPN comes from its architecture that brings together the multicast protocols used at the edge and the BGP-MPLS technology deployed in the core. In particular, the use of BGP for transferring multicast routes across geographic locations and point-to-multipoint LSPs for distributing multicast data bring VPN multicast service offering to the same reliable and scalable level as the VPN unicast service.

**Related Documentation**

- Next-Generation MVPN Network Topology on page 2
- Next-Generation MVPN Concepts and Terminology on page 3
- *Understanding MBGP Multicast VPN Extranets*