NGEN MVPN WITH PIM IMPLEMENTATION GUIDE

PARTNER VERSION – FOR DISTRIBUTION ONLY UNDER SIGNED NDA

Table of Contents

Table of Figures

Introduction

Unicast L3VPN is one of the commonly deployed services across most of the service providers and large enterprise customers. This service is based on RFC 2547bis, which does not have the framework for transporting of multicast traffic across the same infrastructure. L3VPN working group addresses the need with the procedure explained in following drafts.

- draft-ietf-l3vpn-2547bis-mcast-07.txt
- ietf-l3vpn-2547bis-mcast-bgp-05.txt

NGEN MVPN is based on above two drafts and adopted two important properties of unicast BGP/MPLS VPNs:

- BGP protocol is used for distributing all the necessary routing information to enable VPN multicast service. This allows Service Providers to leverage their knowledge and investment in managing BGP/MPLS VPN unicast service to offer VPN multicast services.
- Control plane independence from forwarding plane is provided (as required by RFC4834). This allows the separation of the control and data plane protocols and makes it easier to leverage newer transport technologies such as MPLS in delivering multicast VPN service.

The configurations presented in this document are targeted for service providers and large enterprise networks. More specifically, the configurations are for those wanting to leverage and enhance their existing unicast Layer 3 VPN 2547bis service offering to offer multicast service using the same MPLS data plane and BGP control plane architecture. Doing so reduces the operational expenses of using the same technology used to offer for both unicast and multicast services over the same infrastructure.

This implementation guide assists network designers and operation engineers who support service providers' large enterprise customers with Layer 3 VPN deployments using the Juniper Networks® M Series Multiservice Edge Routers, T Series Core Routers, and MX Series Ethernet Services Routers.

Scope

RFC4364, originally RFC 2547, has not specified a mechanism to provide multicast signaling and multicast data delivery through service provider networks for Layer 3 VPN services. Thus, a number of solutions have been discussed, implemented, and deployed.

This implementation guide focuses on the options where BGP is used for exchanging customer multicast routes among Provider Edge (PE) routers and PIM/GRE is used for creating multicast data tunnels through service provider core network.

Implementation

Figure 1 represents the complete view of the network topology used for this configuration guide.

Tables 1-3 identify the hardware and software used for the validation environment.

Physical Topology

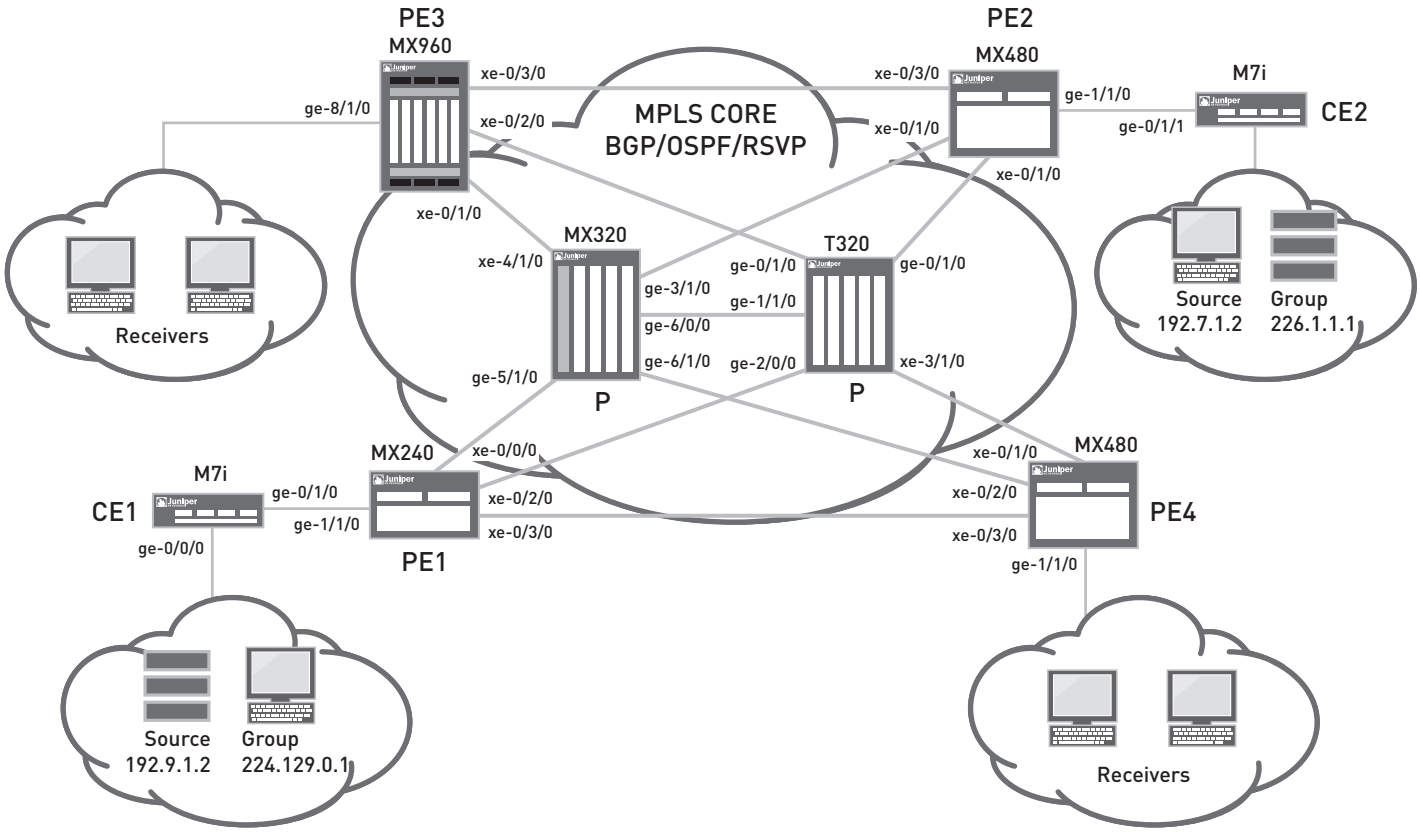

Figure 1: Network topology used for configuration

Hardware Used for Validation Environment

Juniper Networks Equipment

Table 1: Juniper Networks Routers Used for Validation Environment

EQUIPMENT	COMPONENTS
1x Juniper Networks T320 router 1x Juniper Networks M320 router	10 Gigabit Ethernet Xenpak with Type 3 FPCs
2x Juniper Networks MX480 routers 1x Juniper Networks MX960 router 1x Juniper Networks MX240 router	• 2x 40-port 1 Gigabit Ethernet Layer 2 / Layer 3 DPCs (DPCE-R-40GE-SFP or DPCE-R-Q-40GE-SFP) • 8 SFPs • 2x 4-port 10GbE L2/L3 DPCs (DPCE-R-4XGE-XFP or DPCE-R-Q-4XGE-XFP) • 4 XFPs

Testing Equipment Used for Validation Environment

Table 2: Testing Equipment Used for Validation Environment

EQUIPMENT	COMPONENTS
Agilent N2X tester	8 x 10/100/1000Mb ports

Software Used for Validation Environment

Table 3: Software Used for Validation Environment

EQUIPMENT	SOFTWARE
M Series, T Series, and MX Series routers	JUNOS Software Release 9.1R2

Configuration and Validation

Logical Topology

The topology is simplified with two PE routers and two customer edge (CE) routers for the explanation purposes. The other PE configurations are similar with only interface naming changes required as per the router's physical location.

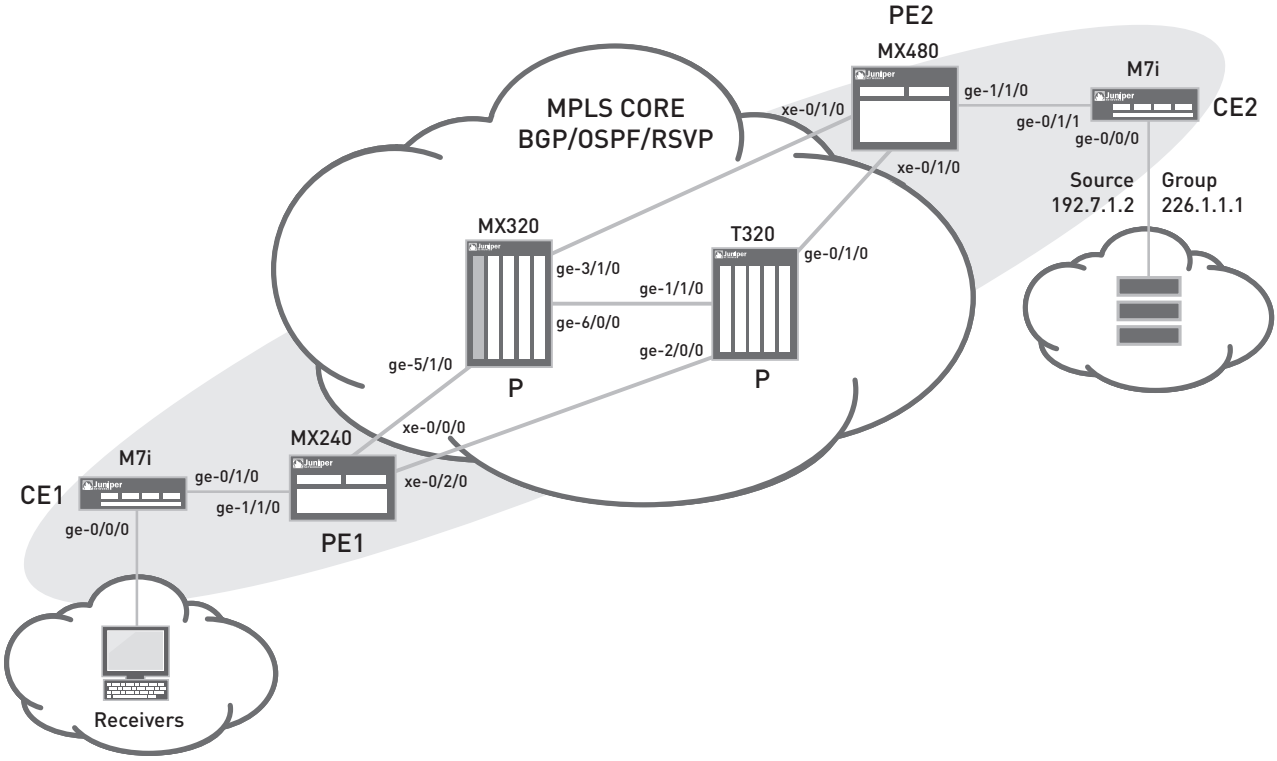

Figure 2: Simplified topology for base configuration

The core network is simulated using two Juniper Networks T320 and M320 Internet routers. Four Juniper Networks MX Series routers are connected to the T320 routers and M320 routers (used as the core router) as shown in Figure 2. Tester ports simulate source and receivers, also as shown in Figure 2.

Base Configuration before Enabling Multicast services for MPLS L3 VPN

In the given physical topology, the following logical design is realized on Juniper Networks routers and the test equipment prior the NGENMVPN configuration.

- All interfaces in Figure 1 are configured in OSPF area 0 and traffic engineering is enabled on OSPF area 0.
- All the PE/P routers are configured to support PIM.
- Full-mesh P2P RSVP-TE LSPs among four Juniper Networks PE routers are configured for unicast traffic tunneling
- One Layer 3 VPN service instance is configured.
- OSPF is enabled on all interfaces on all routers except FXP0, which is a management interface shown below. You can also specify individual interface names. The `traffic-engineering` command needs to be enabled to generate the OPAQUE LSAs (Link State Advertisements) to build the traffic engineering database, which is used to calculate the label switched paths (LSPs) paths. You can also configure IS-IS instead of OSPF and need to add `wide-metrics only` to build the traffic engineering database.

- A full mesh of IBGP is configured between PE routers to form the Layer 3VPN service. Initially, all IBGP sessions between Juniper Networks routers are configured for "inet unicast" and "inet-vpn unicast". Using `family inet-vpn` enables the multiprotocol BGP (MP-BGP) capabilities to exchange the VPNv4 routes between the PEs.
- `MPLS` is enabled on all core facing interfaces of PE routers and P routers.

NGEN MVPN is enabled in two steps:

1. Preparing the core for supporting tunnels used by multicast VPN services
2. Enabling individual Layer 3 VPN to provide multicast service on PE's. However, to further clarify each comment set in the final configuration, the enablement process is discussed in multiple steps.

The configuration statements for each step and final configuration are explained in subsequent sections.

Summary of Steps for Enabling MVPN

Preparing Core for NGEN MVPN Service

- **Step 1:** Configure the provider tunnel signaling protocol, PIM, on the service provider network (PE/P routers) and configure P-RP router.
- **Step 2:** Configure the command "inet-mvpn signaling" to enable MVPN capability for BGP.

Enabling NGEN MVPN for Individual VRFs

- **Step 3:** Enable customer multicast protocol by configuring C-multicast routing protocol on PE and CE routers (PIM or IGMP).
- **Step 4:** Enable the MVPN service for each Layer 3 VPN on PE routers by configuring "protocols MVPN" under the routing-instance.
- **Step 5:** Associate the provider tunnel with the service instance used by the MVPN (sender sites only).

Note: Repeat step 2 on all of PE routers on which MVPN needs to be enabled. This step is a common configuration for all the MVPN service instances and should be completed in a maintenance window because a new address family would reset BGP sessions.

The alternate option is to deploy a separate route-reflector(s) for MVPN. All the PE routers supporting MVPN can peer with new route reflector(s) dedicated for MVPN leaving existing BGP peers intact. In this case, step 2 does not need to be completed within a maintenance window, as there is no need to enable a new address family on the existing BGP session, which will avoid the traffic disturbtion on the existing customers.

The configuration statements for each step and final configuration are explained in subsequent sections.

Step-by-Step Configuration

This section describes the entire step-by-step configuration process.

Step 1: Enable PIM on P and PE Routers

This step prepares core and edge routers to create PIM-ASM tunnels.

- Enable Protocols PIM on all P and PE routers Interfaces in the core.
- Configure the Provider-rendezvous point (P-RP) with a static entry. In this configuration P router is acting as Provider RP.

You may use other techniques, such bootstrap router (BSR), to dynamically elect the RP, which is not discussed in this document. In this setup router P1 (Indica) is configured as a P-RP, and all the other routers with a static RP entry.

```
protocols {
   pim {
        rp {
            static {
                address 21.255.0.1;
            }
        }
        interface all;
        interface fxp0.0 {
```

```
            disable;
        }
    }
}
```

Step 2: Configure inet-mvpn Address Family for IBGP Sessions on PE Routers

Configuring the **inet-mvpn** address family allow PEs to enhance the BGP capability so as to support multicast extension. Doing so enables **mvpn** membership information to be automatically discovered (autodiscovery) and customer multicast routes to be exchanged without requiring full-mesh PIM adjacency among PE routers. As in the previous step, this step also prepares the infrastructure for individual VPNs to be ready to function once MVPN service enabled on them. Perform this configuration on all PE routers.

```
    bgp {
        group mesh {
            type internal;
            local-address 21.255.2.1;
              ...truncated...
            family inet-mvpn {
                signaling;
            }
        …truncated…
        }
```

Step 3: Configure Customer Multicast Routing Protocol

This configuration allows PE routers to exchange routing information between CE routers. Since Layer 3 VPN is deployed such that customer uses a Layer 3 router for the service provider connection, PIM mostly is used for the CE-PE multicast routing protocol. However, in some cases, a customer may want to use IGMP, which is supported by JUNOS® Software. The following configuration enables PIM on VPN for the customer facing interface.

There are two ways to configure the PIM Customer-RP either a PE router or a CE router. If a CE router has chosed to be a C-RP MSDP or anycast-RP should also be configured between the CE and PE routers so that the source-active messages gets communicated between PE and CE (C-RP). This scenario is not validated in this document.

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
        …truncated…
        protocols {
        …truncated…
            pim {
                rp {
                    local {
                        address 102.255.0.1;
                    }
                }
                interface all {
                    mode sparse;
                }
            }
        …truncated…
```

Although it is not used in this setup, following is an IGMP configuration that may be required for directly connected receivers.

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
       …truncated…
       protocols {
         Igmp {
           interface ge-8/1/0.1 {
             static {
               group 224.129.0.1 {
                 source 192.9.1.2;
               }
        }
}
```

Step 4: Enabling MVPN for individual service instance

In this step, enable MVPN on all PE routers for a given Layer 3 VPN service instance. The configuration sets the policies to correctly identify and process various BGP advertisements to build the multicast routing table.

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
       …truncated…
        protocols {
       …truncated…
            mvpn {

       …truncated…
```

Step 5: Associate provider tunnel with L3 VPN service instance.

You only need to perform this configuration on the PE routers that are connected to the senders if it is known that given sites include the multicast source. If all sites have receiver and sender, this configuration must be repeated for all PE routers.

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
…truncated…
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
        vrf-target target:100:50001;
…truncated…
```

The P-multicast group address of the tunnel identifier is a multicast group address assigned by the provider multicast address space. This address should be same for all intra-AS auto discovery routes originated by PEs of the same MVPN.

The **`provider tunnel`** configures the instance to use either RSVP-TE LSPs or PIM tunnels for the data forwarding. In this guide, PIM-SM/GRE is used as the tunnel protocol. Configure the `provider-tunnel` on only the sender site. A PE that is attached to only receiver sites in a MVPN does not need to originate a provider tunnel for that MVPN.

Tunnel PIC is required on all the PE routers to form the provider tunnel using PIM/GRE encapsulation. Tunnel PIC is also required on CE routers, if the" source" connected DR and the C-RP are on different routers to encapsulate and de-encapsulate the PIM Register messages.

On MX Series routers, there is no Tunnel PIC, and Tunnel PIC functionality can be enabled by configuring the chassis.

```
chassis {
    fpc 1 {
        pic 1 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }
}
```

The configuration above adds an 11th Gigabit Ethernet port on this PIC as mt-1/1/10, thus avoiding a burn-out a physical port.

Final Configuration

These configuration tables include base and new MVPN configurations once all the steps for MVPN have been completed. The following configuration includes base and new MVPN configurations once all the steps for MVPN are completed. The '**bold**' text below indicates the configuration that is added to enable MVPN top of base configuration. The configuration shown here does not include sender-site and receive site differentiation and applicable for all the cases. A provider-tunnel configuration *is used* **only** on the PE that serves the site which is known a priori to have (multicast) senders. The sites which are known a priori to have no (multicast) senders in our topology hasn't been associated with provider tunnel since The association is not required for the operation of NGEN MVPN. For the cases where it is not known a priori whether the site contains senders, one need to configure provider tunnel on the PE connected to that site.

PE1 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH DOESN'T HAVE ANY SENDER)

Chassis Configuration

```
Not required for receiver site
```

Interface Configuration

```
Core Facing Interfaces

interfaces {
    xe-0/0/0 {
        unit 0 {
            family inet {
                address 21.0.3.2/30;
            }
            family mpls;
        }
    }
```

PE2 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH HAS ONE OR MORE SENDER)

Chassis Configuration

```
chassis {
    fpc 1 {
        pic 1 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }
}
```

Interface Configuration

```
Core Facing Interfaces

interfaces {
    xe-0/0/0 {
        unit 0 {
            family inet {
                address 21.2.5.1/30;
            }
            family mpls;
        }
    }
```

PE1 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH DOESN'T HAVE ANY SENDER)

Interface Configuration (cont.)

```
xe-0/1/0 {
        unit 0 {
            family inet {
                address 21.0.3.6/30;
            }
            family mpls;
        }
    }
    xe-0/2/0 {
        unit 0 {
            family inet {
                address 21.3.5.1/30;
            }
            family mpls;
        }
    }
    xe-0/3/0 {
        unit 0 {
            family inet {
                address 21.3.4.1/30;
            }
            family mpls;
        }
    }
```

PE-CE Interface

```
    ge-1/1/0 {
        vlan-tagging;
        unit 1 {
            vlan-id 1;
            family inet {
                address 40.40.40.1/30;
            }
        }
    }
```

Loopback Interface

```
    lo0 {
        unit 0 {
            family inet {
                address 21.255.3.1/32;
            }
        }
        unit 1 {
            family inet {
                address 103.255.0.1/32;
            }
        }
    }
}
```

PE2 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH HAS ONE OR MORE SENDER)

Interface Configuration (cont.)

```
xe-0/1/0 {
        unit 0 {
            family inet {
                address 21.2.5.5/30;
            }
            family mpls;
        }
    }
    xe-0/2/0 {
        unit 0 {
            family inet {
                address 21.0.2.2/30;
            }
            family mpls;
        }
    }
    xe-0/3/0 {
        unit 0 {
            family inet {
                address 21.1.2.2/30;
            }
            family mpls;
        }
     }
```

PE-CE Interface

```
    ge-1/1/0 {
        vlan-tagging;
        unit 1 {
            vlan-id 1;
            family inet {
                address 40.40.40.1/30;
            }
        }
    }
```

Loopback Interface

```
    lo0 {
        unit 0 {
            family inet {
                address 21.255.2.1/32;
            }
        }
Loopback for service instance:
        unit 1 {
            family inet {
                address 102.255.0.1/32;
            }
        }
    }
}
```

PE1 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH DOESN'T HAVE ANY SENDER)

Routing and Forwarding Options

```
routing-options {
    router-id 21.255.3.1;
    autonomous-system 100;
    forwarding-table {
        export pplb;
    }
}
```

Protocols

```
protocols {
    rsvp {
        interface all {
            aggregate;
            reliable;
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path r3r2_uni_1 {
            to 21.255.2.1;
            node-link-protection;
            primary r3r2_uni_1;
        }
        label-switched-path r3r1_uni_1 {
            to 21.255.1.1;
            node-link-protection;
            primary r3r1_uni_1;
        }
        label-switched-path r3r4_uni_1 {
            to 21.255.4.1;
            node-link-protection;
        }
    }
    bgp {
        group mesh {
            type internal;
            local-address 21.255.3.1;
            family inet {
                unicast;
            }
            family inet-vpn {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            include-mp-next-hop;
            neighbor 21.255.0.1;
            neighbor 21.255.5.1;
            neighbor 21.255.2.1;
            neighbor 21.255.1.1;
```

PE2 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH HAS ONE OR MORE SENDER)

Routing and Forwarding Options

```
routing-options {
    router-id 21.255.2.1;
    autonomous-system 100;
    forwarding-table {
        export pplb;
    }
}
```

Protocols

```
protocols {
    rsvp {
        interface all {
            aggregate;
            reliable;
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path r2r1_uni_1 {
            to 21.255.1.1;
            node-link-protection;
            primary r2r1_uni_1;
        }
        label-switched-path r2r3_uni_1 {
            to 21.255.3.1;
            node-link-protection;
            primary r2r3_uni_1;
        }
        label-switched-path r2r4_uni_1 {
            to 21.255.4.1;
            node-link-protection;
        }
    }
    bgp {
        group mesh {
            type internal;
            local-address 21.255.2.1;
            family inet {
                unicast;
            }
            family inet-vpn {
                unicast;
            }
            family inet-mvpn {
                signaling;
            }
            neighbor 21.255.0.1;
            neighbor 21.255.5.1;
            neighbor 21.255.1.1;
            neighbor 21.255.3.1;
            neighbor 21.255.4.1;
```

PE1 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH DOESN'T HAVE ANY SENDER)

Protocols (cont.)

```
            neighbor 21.255.4.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fxp0.0 {
                disable;
            }
            interface all {
            }
        }
    }
    pim {
        rp {
            static {
                address 21.255.0.1; # P-RP ##
            }
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
```

Policies

```
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
```

Service Instances

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
        interface ge-1/1/0.1;
        interface lo0.1;
        route-distinguisher 21.255.3.1:50001;


# NO PROVIDER TUNNEL CONFIG FOR RECEIVER SITE


        vrf-target target:100:50001;
        vrf-table-label;
        protocols {
            bgp {
                group l3vpn_50001 {
```

PE2 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH HAS ONE OR MORE SENDER)

Protocols (cont.)

```
  }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all {
            }
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        rp {
            static {
                address 21.255.0.1; # P-RP ##
            }
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
```

Policies

```
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
```

Service Instances

```
routing-instances {
    l3vpn_50001 {
        instance-type vrf;
        interface ge-1/1/0.1;
        interface lo0.1;
        route-distinguisher 21.255.2.1:50001;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
        vrf-target target:100:50001;
        vrf-table-label;
        protocols {
            bgp {
                group l3vpn_50001 {
```

PE1 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH DOESN'T HAVE ANY SENDER)

Service Instances (cont.)

```
                neighbor 103.41.0.2 {
                    peer-as 50001;
                }
            }
        }
        ospf {
            area 0.0.0.0 {
                interface all;
            }
        }
        pim {
            rp {
                static {
                    address 102.255.0.1;
                }
            }
            interface all {
                mode sparse;
            }
        }
        mvpn ;
    }
  }
}
```

PE2 CONFIGURATION (THE PE THAT SERVES THE SITES WHICH HAS ONE OR MORE SENDER)

Service Instances (cont.)

```
                neighbor 40.40.40.2 {
                    peer-as 50001;
                }
            }
        }
        ospf {
            area 0.0.0.0 {
                interface all;
            }
        }
        pim {
            rp {
                local {
                    address 102.255.0.1;
                }
            }
            interface all {
                mode sparse;
            }
        }
        mvpn ;
    }
  }
}
```

CE1 Configuration

```
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 192.9.1.1/24;
            }
        }
    }
    ge-0/1/0 {
        vlan-tagging;
        unit 1 {
            vlan-id 1;
            family inet {
                address 103.41.0.2/30;
            }
        }
    }
}
routing-options {
    autonomous-system 50001;
}
protocols {
    igmp {
        interface ge-0/0/0.0;
        interface ge-0/1/0.1;
    }
    bgp {
        group skoda {
            neighbor 103.41.0.1 {
                peer-as 100;
                as-override;
            }
        }
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        rp {
            static {
                address 102.255.0.1;
            }
        }
        interface ge-0/0/0.0;
        interface ge-0/1/0.0;
        interface ge-0/1/0.1;
    }
}
```

CE2 Configuration

```
interfaces {
    ge-5/0/0 {
        unit 0 {
            family inet {
                address 192.7.1.1/24;
            }
        }
    }
    ge-5/0/1 {
        vlan-tagging;
        unit 1 {
            vlan-id 1;
            family inet {
                address 40.40.40.2/30;
            }
        }
    }
}
routing-options {
    autonomous-system 50001;
}
protocols {
    bgp {
        group scorpio {
            neighbor 40.40.40.1 {
                peer-as 100;
            }
        }
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        rp {
            static {
                address 102.255.0.1;
            }
        }
        interface ge-0/0/0.0;
        interface ge-5/0/1.1;
    }
}
```

ConfigurationValidation

End-to-End Traffic Flow

The VPN multicast source 192.7.1.2 at CE2 (tester port) delivers multicast data to the IGMP receiver at CE1. Static IGMP joins are created on PE3 and PE4 for this group so they are also receiving the multicast stream from the same source as shown below.

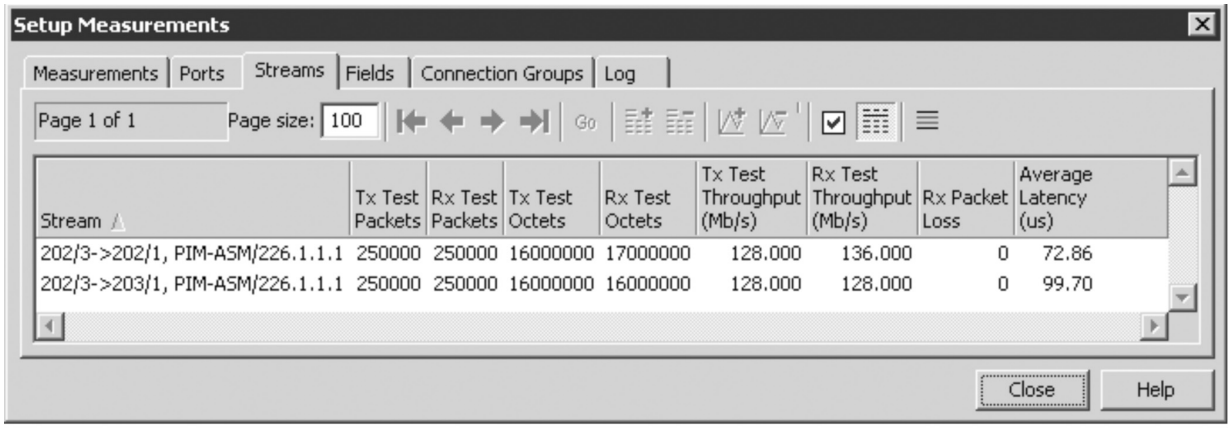

Below are commands that can be used to verify the NGEN MVPN setup. Detailed outputs of these show commands are attached in Appendix A.

Validation Commands for NGEN MVPN Control Plane

If everything works correctly, the VPN multicast source 192.7.1.2 at CE2 (tester port) delivers multicast data to the receiver at CE1.

This section describes different parts of the control plane. Assume that the source 192.9.1.2 started transmitting data to group 226.1.1.1 and the receiver connected to CE1 sent a join (*, 226.1.1.1) to PE1.

```
lab@PE1# run show bgp summary
lab@PE1> show route table l3vpn_50001.mvpn.0 detail
lab@PE1> show mvpn instance l3vpn_50001 extensive
lab@PE1 > show pim mdt instance l3vpn_50001
lab@PE1 > show interfaces mt*
lab@PE1 > show route table l3vpn_50001.mvpn.0 detail | find 5:
lab@PE2> show route table l3vpn_50001.mvpn.0 detail | find 5:
lab@PE2> show route table bgp.mvpn.0 detail
lab@PE2> show pim join extensive instance l3vpn_50001
lab@PE2> show igmp statistics
lab@PE2> show mvpn neighbor
lab@PE2> show mvpn c-multicast
```

Validation Commands for NGEN MVPN Data Plane

```
lab@PE2> show route table l3vpn_50001.inet.1 detail
lab@PE2>show route forwarding-table multicast destination 224.129.0.1 extensive vpn l3vpn_50001
lab@PE2> show multicast next-hops
lab@PE2> show multicast route extensive instance l3vpn_50001
lab@PE2> show mpls lsp p2mp
```

Detailed Control Plane Validation

This section describes different parts of the control plane. Assume that the source 192.7.1.2 is transmitting data to group 226.1.1.1 and the receiver connected to CE1 has sent a join (*, 226.1.1.1) to PE1.

The State of PE-PE IBGP Session

The IBGP session between the PE routers must be operational before the PE routers can exchange any `mvpn` routes. If the BGP session is working correctly, use the `show bgp summary` on each PE to see the `bgp.l3vpn.0` and `bgp.mvpn.0` tables.

```
lab@PE2# run show bgp summary
Groups: 2 Peers: 6 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 0          0          0          0          0          0
bgp.l3vpn.0           13         13          0          0          0          0
bgp.mvpn.0             6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
21.255.0.1        100      12487      12507       0       3 3d 22:55:06 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 0/0/0
  bgp.mvpn.0: 0/0/0
21.255.1.1        100      12502      12506       0       3 3d 22:55:06 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 4/4/0
  bgp.mvpn.0: 1/2/0
  l3vpn_50001.inet.0: 4/4/0
  l3vpn_50001.mvpn.0: 1/2/0
21.255.3.1        100      12484      12506       0       3 3d 22:55:06 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 5/5/0
  bgp.mvpn.0: 2/2/0
  l3vpn_50001.inet.0: 5/5/0
  l3vpn_50001.mvpn.0: 1/2/0
21.255.4.1        100        742        728       0       9     5:30:37 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 4/4/0
  bgp.mvpn.0: 1/2/0
  l3vpn_50001.inet.0: 4/4/0
  l3vpn_50001.mvpn.0: 1/2/0
21.255.5.1        100      12491      12505       0       3 3d 22:55:06 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 0/0/0
  bgp.mvpn.0: 0/0/0
40.40.40.2      50001       2921       2904       0       4    21:54:26 Establ
  l3vpn_50001.inet.0: 0/0/0

[edit]

lab@PE1# run show bgp summary
Groups: 2 Peers: 6 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 0          0          0          0          0          0
bgp.l3vpn.0           11          7          0          0          0          0
bgp.mvpn.0             5          5          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/
Damped...
21.255.0.1        100      15609      15646       0       0 4d 22:38:26 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 0/0/0
  bgp.mvpn.0: 0/0/0
```

```
21.255.1.1        100      15629      15644       0       0 4d 22:38:35 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 0/4/0
  bgp.mvpn.0: 1/1/0
  l3vpn_50001.inet.0: 0/4/0
  l3vpn_50001.mvpn.0: 1/1/0
21.255.2.1        100      12704      12674       0       4 3d 22:55:48 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 3/3/0
  bgp.mvpn.0: 2/2/0
  l3vpn_50001.inet.0: 3/3/0
  l3vpn_50001.mvpn.0: 2/2/0
21.255.4.1        100        743        733       0       6     5:30:58 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 4/4/0
  bgp.mvpn.0: 2/2/0
  l3vpn_50001.inet.0: 4/4/0
  l3vpn_50001.mvpn.0: 1/2/0
21.255.5.1        100      15615      15645       0       0 4d 22:38:38 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 0/0/0
  bgp.mvpn.0: 0/0/0
103.41.0.2      50001      16038      15799       0       0 4d 22:38:34 Establ
l3vpn_50001.inet.0: 0/7/0
```

Originating a Type 1 Auto-discovery Route

A NGEN MVPN PE must originate a local auto-discovery (AD) route, which means the PE routers must install a local AD route in the L3VPN_50001.mvpn.0 table and advertise the local AD routes to each other. This route is used to find other PEs that has the same MVPN Site membership connected to them.

The AD routes should look like the output below. The loopback IP of this PE1 router is 21.255.2.1, so the AD route that PE2 announces to all other entire PEs in the network has 21.255.2.1 as its next hop. All PEs in the network advertise this route if they have a MVPN site configured on them.

lab@PE2> show route advertising-protocol bgp 21.255.3.1 detail table l3vpn_50001.mvpn.0

```
l3vpn_50001.mvpn.0: 6 destinations, 9 routes (6 active, 1 holddown, 0 hidden)
* 1:21.255.2.1:50001:21.255.2.1/240 (1 entry, 1 announced)
 BGP group mesh type Internal
     Route Distinguisher: 21.255.2.1:50001
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [100] I
     Communities: target:100:50001
     PMSI: Flags 0:PIM-SM:label[0:0:0]:Sender 21.255.2.1 Group 224.1.1.1
```

You can see the PE2 is advertising with the provider multicast service interface (PMSI) attribute, as it this PE is in sender site where the source is connected that has `provider-tunnel` configured.

Receiving a Type 1 AD Route

The AD route originated from PE1 is received on PE2 and the other PE routers, indicating that there is a MVPN site present on PE1.

The PE2 received the AD route from the PE1 as shown in the following output.

```
lab@PE2> show route receive-protocol bgp 21.255.3.1 detail table l3vpn_50001.mvpn.0
l3vpn_50001.mvpn.0: 6 destinations, 9 routes (6 active, 1 holddown, 0 hidden)
* 1:21.255.3.1:50001:21.255.3.1/240 (1 entry, 1 announced)
     Route Distinguisher: 21.255.3.1:50001
     Nexthop: 21.255.3.1
     Localpref: 100
     AS path: I
     Communities: target:100:50001
```

Note that there is no PMSI attribute attached to the AD route advertised from PE1 because there is no provider-tunnel configuration on PE1.

MVPN to provider tunnel Binding and the State of mt IFLs

The state of provider tunnel binding to an MVPN can be verified by issuing the `show mvpn instance` command. The output shows MVPN module's view of the tunnel binding.

```
lab@PE2> show mvpn instance extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: l3vpn_50001
  Provider tunnel: I-P-tnl:PIM-SM:21.255.2.1, 224.1.1.1
  Neighbor                               I-P-tnl
  21.255.1.1
  21.255.3.1
  21.255.4.1
  C-mcast IPv4 (S:G)            Ptnl                              St
  192.7.1.2/32:226.1.1.1/32     PIM-SM:21.255.2.1, 224.1.1.1           RM

lab@PE2>
```

The above output shows that PE2 has established a `provider-tunnel` using PIM-SM with the other PEs in the network.

The mt IFLs used for each tunnel are shown below. These `mt` interface should have the GRE as the encapsulation method.

```
lab@PE2> show pim mdt instance l3vpn_50001 extensive
Instance: PIM.l3vpn_50001
Tunnel direction: Outgoing
Default group address: 224.1.1.1
Default tunnel interface: mt-1/1/10.32768

Instance: PIM.l3vpn_50001
Tunnel direction: Incoming
Default group address: 224.1.1.1
Default tunnel interface: mt-1/1/10.49152
```

Originating a Type 5 Route

When the source within the site connected to PE1 starts to send multicast data, the PIM Designated Router connected to the source originates C-PIM Register message and sends it to the C-RP, which is also PE2 (as the PE2 **l3vpn_50001 instance** is configured as the C-RP).

As a result, PE2 must originate a type 5 (source active) SA route upon receiving the first data packet or the C-PIM register messages:

In this scenario, PE2 routing-instance l3vpn_50001 is configured as C-RP and source packets are sent to PE2 via L3VPN route.

```
lab@PE2> show route table l3vpn_50001.mvpn.0 detail | find 5:
5:21.255.2.1:50001:32:192.7.1.2:32:226.1.1.1/240 (1 entry, 1 announced)
        *PIM    Preference: 105
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 41
                State: <Active Int>
                Age: 7
                Task: PIM.l3vpn_50001
                Announcement bits (2): 1-mvpn global task 2-BGP RT Background
                AS path: I
```

As it can be seen above the route is installed by protocol PIM on the l3vpn_50001. mvpn.0 table.Then PE2 advertises the SA route to PE1 and other PE routers as seen below.

```
lab@PE2> show route advertising-protocol bgp 21.255.3.1 table l3vpn_50001 detail

l3vpn_50001.mvpn.0: 7 destinations, 8 routes (7 active, 1 holddown, 0 hidden)

* 5:21.255.2.1:50001:32:192.7.1.2:32:226.1.1.1/240 (1 entry, 1 announced)
 BGP group mesh type Internal
     Route Distinguisher: 21.255.2.1:50001
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [100] I
     Communities: target:100:50001
.
```

Note that the SA route is advertised with the unicast import route target (RT) of target 100:50001 just like the AD route.

PE1 installs the route in the bgp.mvpn.0 table first and then in the L3VPN_50001.mvpn.0 table, assuming the route is accepted by the VRF import policy.

```
lab@PE1> show route table bgp.mvpn.0 detail
5:21.255.2.1:50001:32:192.7.1.2:32:226.1.1.1/240 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 4
                Source: 21.255.2.1
                Protocol next hop: 21.255.2.1
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:   100 Peer AS:   100
                Age: 5:05:07    Metric2: 2
                Task: BGP_100.21.255.2.1+179                AS path: I
```

```
                Communities: target:100:50001
                Localpref: 100
                Router ID: 21.255.2.1
                Secondary Tables: l3vpn_50001.mvpn.0
```

Note that the Secondary Tables field is set to L3VPN_50001.mvpn.0, indicating that the route is also installed in L3VPN_50001.mvpn.0 table.

The route entry for the SA route looks like the following in L3VPN_50001.mvpn.0 table.

```
lab@PE1> show route table l3vpn_50001.mvpn.0 detail | find 5:
5:21.255.2.1:50001:32:192.7.1.2:32:226.1.1.1/240 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 4
                Source: 21.255.2.1
                Protocol next hop: 21.255.2.1
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:   100 Peer AS:   100
                Age: 5:04:40    Metric2: 2
                Task: BGP_100.21.255.2.1+179
                Announcement bits (2): 0-PIM.l3vpn_50001 1-mvpn global task
                AS path: I
                Communities: target:100:50001
                Localpref: 100
                Router ID: 21.255.2.1
                Primary Routing Table bgp.mvpn.0
```

Originating a Type 6 Routes

PE1 should originate a type 6 route in response to receiving a (C-*, 226.1.1.1) from CE1. If PIM. l3vpn_50001 created a state for the join, it should also install a type 6 route in l3vpn_50001.mvpn.0 table. This will not be advertised to any other PE routers. This is similar to having (*,G) on ROSEN MVPN.

```
lab@PE1> show route table l3vpn_50001.mvpn.0 detail
6:21.255.2.1:50001:100:32:102.255.0.1:32:226.1.1.1/240 (1 entry, 1 announced)
        *PIM    Preference: 105
                Next hop type: Multicast (IPv4), Next hop index: 1048580
                Next-hop reference count: 5
                State: <Active Int>
                Age: 2:34:15
                Task: PIM.l3vpn_50001
                Announcement bits (1): 1-mvpn global task
                AS path: I
                Communities: no-advertise target:21.255.2.1:3
```

Originating a Type 7 Route

PE1 learns the IP address of the source sending data to group 226.1.1.1 via type 5 SA route from PE2 (C-RP). Now that it knows there is a local receiver and the source is reachable via MVPN, it originates a type 7 route and installs it in the L3VPN_50001.mvpn.0 table and also advertised to PE2 and other PE's if any.

```
lab@PE1> show route advertising-protocol bgp 21.255.2.1 detail table l3vpn_50001.mvpn.0
* 7:21.255.2.1:50001:100:32:192.7.1.2:32:226.1.1.1/240 (2 entries, 2 announced)
 BGP group mesh type Internal
     Route Distinguisher: 21.255.2.1:50001
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [100] I
     Communities: target:21.255.2.1:3

lab@PE1> show route table l3vpn_50001.mvpn.0 detail
l3vpn_50001.mvpn.0: 8 destinations, 10 routes (8 active, 2 holddown, 0 hidden)
7:21.255.2.1:50001:100:32:192.7.1.2:32:226.1.1.1/240 (2 entries, 2 announced)
        *MVPN   Preference: 70
                Next hop type: Indirect
                Next-hop reference count: 3
                Protocol next hop: 21.255.3.1
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 2:37:31    Metric2: 1
                Task: mvpn global task
                Announcement bits (3): 0-PIM.l3vpn_50001 1-mvpn global task 2-BGP RT Background
                AS path: I
                Communities: target:21.255.2.1:3
         PIM    Preference: 105
                Next hop type: Multicast (IPv4), Next hop index: 1048580
                Next-hop reference count: 5
                State: <Int>
                Inactive reason: Route Preference
                Age: 2:37:31
                Task: PIM.l3vpn_50001
                Announcement bits (1): 1-mvpn global task
                AS path: I
                Communities: target:21.255.2.1:3
```

There are two protocols installing the type 7 route in the ASM case. Since MVPN has a better route preference (70) than the PIM protocol (105), the MVPN route seems active and the PIM route seems inactive

The following output shows the state in the PIM.L3VPN_50001 database for the join after the SA route is received from PE2.

```
lab@PE1> show pim join extensive instance l3vpn_50001
Group: 226.1.1.1
    Source: *
    RP: 102.255.0.1
    Flags: sparse,rptree,wildcard
    Upstream protcol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: ge-1/1/0.1
            103.41.0.2 State: Join Flags: SRW Timeout: 175

Group: 226.1.1.1
    Source: 192.7.1.2
    Flags: sparse
    Upstream protcol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to Source
    Keepalive timeout:
    Downstream neighbors:
        Interface: ge-1/1/0.1
            103.41.0.2 State: Join Flags: S Timeout: 175
```

The first join is created when the original (*, 224.129.0.1) is received from CE and the type 6 route is installed. The second output is created after the SA route is received and the type 7 route is originated.

The type 7 routes are installed in the L3VPN_50001.mvpn.0 table because they were accepted by the internal VRF import policy that is applied to C-multicast (type 7) routes. The following output shows the policy on PE2.

```
lab@PE1> show policy
Configured policies:
__vrf-mvpn-import-cmcast-l3vpn_50001-internal__
__vrf-mvpn-import-cmcast-leafAD-global-internal__
__vrf-mvpn-export-target-l3vpn_50001-internal__
__vrf-mvpn-import-target-l3vpn_50001-internal__
__vrf-export-l3vpn_50001-internal__
__vrf-import-l3vpn_50001-internal__
pplb
__vrf-mvpn-export-inet-l3vpn_50001-internal__

lab@PE1>

lab@PE2> show policy __vrf-mvpn-import-cmcast-l3vpn_50001-internal__
Policy __vrf-mvpn-import-cmcast-l3vpn_50001-internal__:
    Term unnamed:
        from community __vrf-mvpn-community-rt_import-target-l3vpn_50001-internal__
[target:21.255.2.1:3 ]
        then accept
    Term unnamed:
        then reject

lab@PE1>
```


Since the imported RT of the received C-multicast route matches the RT specified in the internal VRF import policy, the route is accepted and installed in the L3VPN_50001.mvpn.0 table.

Verifying the P-PIM State on the Core

To verify the PIM state on all the core P and PE routers, use the `show pim neighbors` command as follows.

```
lab@PE2> show pim neighbors
Instance: PIM.master

Interface           IP V Mode        Option      Uptime Neighbor addr
xe-0/0/0.0           4 2             HPLG   1d 06:01:14 21.2.5.2
xe-0/1/0.0           4 2             HPLG   1d 06:01:14 21.2.5.6
xe-0/2/0.0           4 2             HPLG   1d 06:01:14 21.0.2.1
xe-0/3/0.0           4 2             HPLG   1d 06:00:48 21.1.2.1

lab@PE2>
```

You can verify the PIM state on the customer PIM network using the following output.

```
lab@PE2> show pim neighbors instance l3vpn_50001
Instance: PIM.l3vpn_50001

Interface           IP V Mode        Option      Uptime Neighbor addr
ge-1/1/0.1           4 2             HPLG      23:14:45 40.40.40.2
```

The provider PIM state is established using the `provider-tunnel` group address specified under the `routing-instance`. You can verify the data flow on the provider PIM join state.

All the C-PIM Multicast traffic destined to 226.1.1.1 is tunneled using the P-PIM of the `provider-tunnel` group address 224.1.1.1 as follows.

Provider PIM Traffic

```
lab@PE2> show multicast route extensive
Family: INET

Group: 224.1.1.1
    Source: 21.255.2.1/32
    Upstream interface: local
    Downstream interface list:
        xe-0/0/0.0 xe-0/3/0.0
    Session description: ST Multicast Groups
    Statistics: 17494 kBps, 249911 pps, 49417181 packets
    Next-hop ID: 1048587
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0

Family: INET6
```

Customer PIM Traffic

```
lab@PE2> show multicast route extensive instance l3vpn_50001
Family: INET

Group: 226.1.1.1
    Source: 192.7.1.2/32
    Upstream interface: ge-1/1/0.1
    Downstream interface list:
        mt-1/1/10.32768
    Session description: Unknown
    Statistics: 11500 kBps, 250003 pps, 2395241688 packets
    Next-hop ID: 1048582
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

The following output shows the PIM protocol view of the core and customer PIM on the `routing-instance`.

```
lab@PE2> show route table l3vpn_50001.inet.1

l3vpn_50001.inet.1: 1 destinations, 2 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

226.1.1.1,192.7.1.2/32*[MVPN/70] 03:56:51
                        Multicast (IPv4)
                      [PIM/105] 03:57:12
                        Multicast (IPv4)

lab@PE2> show route table inet.1

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

224.1.1.1,21.255.2.1/32*[PIM/105] 06:46:04
                        Multicast (IPv4)

lab@PE2>
```

Detailed Data Plane Validation

You can verify the forwarding state of the network using similar steps as used for the control plane. In JUNOS Software, it is the MVPN module's responsibility to download the VPN multicast routing information to the forwarding plane.

You can display the forwarding state for 226.1.1.1 on PE2 several ways. For example, from PIM's perspective, the following output is what is installed in the kernel.

```
lab@PE2> show route table l3vpn_50001.inet.1 extensive

l3vpn_50001.inet.1: 1 destinations, 2 routes (1 active, 0 holddown, 0 hidden)
226.1.1.1.192.7.1.2/64 (2 entries, 1 announced)
TSI:
KRT in-kernel 226.1.1.1.192.7.1.2/64 -> {[1048582]}
        *MVPN   Preference: 70
                Next hop type: Multicast (IPv4), Next hop index: 1048582
                Next-hop reference count: 3
                State: <Active Int>
                Age: 1d 2:37:29
                Task: mvpn global task
                Announcement bits (1): 0-KRT
                AS path: I
         PIM    Preference: 105
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 41
                State: <Int>
                Inactive reason: Route Preference
                Age: 1d 2:37:50
                Task: PIM.l3vpn_50001
                AS path: I
```

The route was installed by the MVPN (module), and the next-hop index of the downstream interface is **1048582**.

From the kernel's point of view, the multicast route 226.1.1.1 looks like the following output.

```
lab@PE2> show route forwarding-table multicast destination 226.1.1.1 vpn l3vpn_50001 extensive
Routing table: l3vpn_50001.inet [Index 3]
Internet:

Destination:  226.1.1.1.192.7.1.2/64
  Learn VLAN: 258                     Route type: user
  Route reference: 0                  Route interface-index: 76
  IFL generation: 0                   Epoch: 255
  Sequence Number: 0                  Learn Mask: 0x0
  IPC generation: 54520               L2 Flags: none
  Flags: cached, check incoming interface, accounting, sent to PFE, rt nh decoupled
  Next-hop type: indirect             Index: 1048582  Reference: 2
  Next-hop type: routed multicast     Index: 555      Reference: 1
```

The next-hop index of the downstream interface matches in the output of both commands. The downstream interfaces and their next-hop index are also tracked by PIM. You can display all the downstream interfaces known to PIM using the `show multicast next-hops` command.

```
lab@PE2> show multicast next-hops detail
Family: INET
ID          Refcount KRefcount Downstream interface
1048582            3         1 mt-1/1/10.32768
1048586            2         1 xe-0/0/0.0
                               xe-0/3/0.0
```

You can display PIM's view of the forwarding entries using the `show multicast route extensive instance l3vpn_50001` command.

You can also the see the multicast traffic flow using the following command. In the setup, the multicast traffic is flowing at the rate of **250000 pps**.

```
lab@PE2> show multicast route instance l3vpn_50001 extensive
Family: INET

Group: 226.1.1.1
    Source: 192.7.1.2/32
    Upstream interface: ge-1/1/0.1
    Downstream interface list:
        mt-1/1/10.32768
    Session description: Unknown
    Statistics: 11500 kBps, 250000 pps, 23528147649 packets
    Next-hop ID: 1048582
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface
```

Summary

This guide discusses the implementation of NG-MVPN using PIM-ASM as the Provider tunnel mechanism. It provides detailed step by step configuration and implementation methods to enable Multicast service on top of the L3VPN unicast service offerings. It also provides the detailed "show commands outputs:" and "troubleshooting steps "to validate successful implementation.

About Juniper

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER
(888.586.4737)
or 408.745.2000
Fax: 408.745.2100

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin,
Ireland
Phone: 35.31.8903.600
Fax: 35.31.8903.601

8010027-001-EN May 2009