

# Multicast in a VPN II

---

## In This Chapter

This section provides information about multicast in a VPRN service.

Topics in this section include:

- [Applicability on page 694](#)
- [Summary on page 695](#)
- [Overview on page 697](#)
- [Configuration on page 699](#)
- [Conclusion on page 748](#)

## Applicability

This section is applicable to all 7750 SR platforms and to 7450 ESS platforms when configured in mixed-mode. It was tested on release 9.0R5. The features are supported with IOM2, IOM3-XP and IMMs, and need chassis mode C or higher. There are no other pre-requisites for this configuration.

## Summary

Multicast VPN (MVPN) or Next Generation IP Multicast in an IP-VPN (NG-MVPNs) architectures describe a set of VRFs (Virtual Routing and Forwarding) or VPRNs (Virtual Private Routed Networks) that support the transport of multicast traffic across a provider network. MVPNs are defined in draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*, and in draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPN*.

Initial MVPN deployments were originally based on Draft-Rosen (draft-rosen-vpn-mcast-08.txt) which described the protocols and procedures required to support an IP Multicast VPN. There are a number of limitations with the Draft-Rosen implementation including, but not limited to:

- Draft-Rosen requires a set of multicast distribution trees (MDTs) per VPN, which requires a PIM state per MDT. There is no option to aggregate MDT across multiple VPNs
- Customer signaling, initially PE discovery and Data MDT signaling, were all PIM-based as there was no mechanism available to decouple these. Now, PE discovery is supported using a BGP MDT Address Family Identifier/Subsequent Address Family Identifier (AFI/SAFI), however, the data MDT still needs PIM.
- There is no mechanism for using MPLS to encapsulate multicast traffic in the VPN. GRE is the only encapsulation method available in Draft-Rosen.
- Draft-Rosen multicast trees are signaled using PIM only. MVPN allows the use of mLDP and RSVP P2MP LSPs.
- PE to PE protocol exchanges for Draft-Rosen is achieved using PIM only. MVPN allows for the use of BGP signaling as per unicast Layer 3 VPNs.

NG-MVPN addresses these limitations by extending the idea of the per-VRF tree by introducing the idea of Provider Multicast Service Interfaces (PMSIs). These are equivalent to the default MDTs of Draft-Rosen. NG-MVPN allows the decoupling of the mechanisms required to create a multicast VPN, such as PE auto-discovery (which PEs are members of which VPN), PMSI signaling (creation of tunnels between PEs) and customer multicast signaling (multicast signaling—IGMP/PIM—received from customer edge routers). Two types of PMSI exist:

- Inclusive (I-PMSI) — Contains all the PEs for a given MVPN, I-PMSI is the default multicast data path between all PEs of the same VPN.
- Selective (S-PMSI) — Contains only a subset of PEs of a given MVPN, used to optimize multicast stream distribution to only the PEs with active receivers for those streams.

The [Multicast in a VPN](#) section on page 635 contains the VPN configuration required for the provider multicast domain using PIM Any Source Multicast (ASM) with auto-discovery based on PIM or BGP auto-discovery (A/D), PIM used for the customer multicast signaling and PIM Source Specific Multicast (SSM) used for the S-PMSI creation. The customer domain configuration covers the following cases:

## Summary

- PIM ASM with the Rendezvous Point (RP) in the provider PE
- PIM ASM using anycast RP on the provider RPs
- PIM SSM

This section introduces some of the features that were not supported at the time of writing of the *Multicast in a VPRN I* (Release 7.0). It provides configuration details to implement:

- Multicast LDP (mLDP) and RSVP-TE Point to Multipoint (P2MP) for building customer trees (C-trees) which are using MPLS instead of PIM techniques.
- MVPN source redundancy.
- MDT AFI/SAFI (to fully interoperate with Cisco networks).

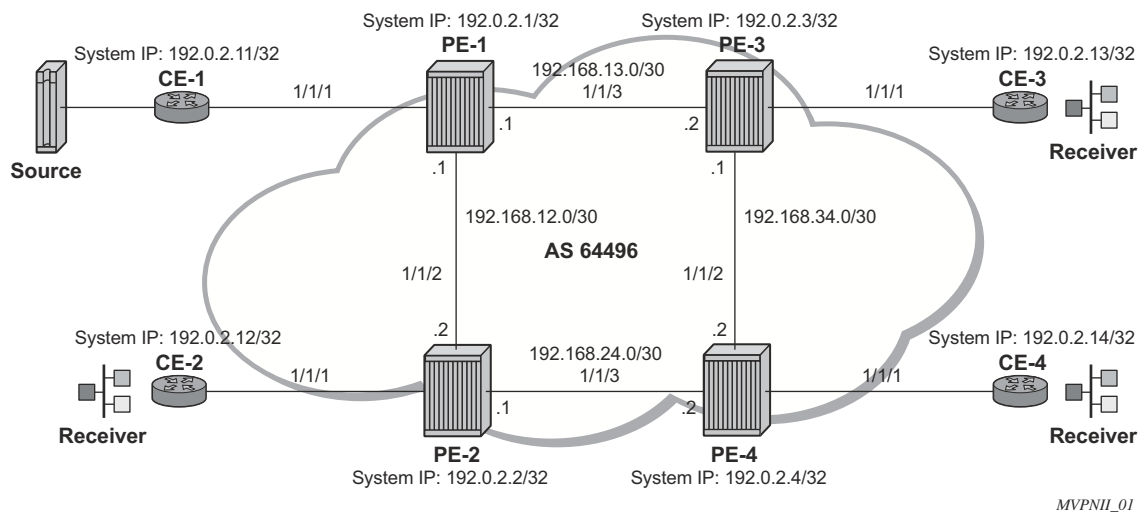
Note that PIM SSM is the only case addressed in this example, other PIM customer domain configurations are out of the scope, for more information refer to [Multicast in a VPN I on page 635](#)

## Overview

The network topology is displayed in [Figure 95](#). The setup consists of four 7750 SRs acting as Provider Edge (PE) routers within a single Autonomous System (AS).

- Full mesh ISIS in the AS (OSPF could be used instead)
- LDP on all interfaces in each AS (RSVP could be used instead)
- MP-iBGP sessions between the PE routers in the AS (Route Reflectors (RRs) could also be used).
- Layer 3 VPN on all PEs with identical route targets, in the form AS-number: vprn-service-id

Connected to each PE is a single 7750 SR acting as a Customer Edge (CE) router. CE-1 has a multicast source connected, and PEs 2 to 4 each have a single receiver connected which will receive the multicast streams from the source. In this setup, each receiver is IGMPv3 capable. If the receiver is IGMPv3 capable, it will issue IGMPv3 reports that may include a list of required source addresses.



**Figure 95: Network Topology**

When the receiver wishes to become a member of any group, the source address of the group must be known to the CE. As a result the source address must be IP reachable by each CE, so it is advertised by CE-1 to the PEs with attachment circuits in VPRN using BGP. Static routes are then configured on the receiver CEs to achieve IP reachability to the source address of the multicast group.

## Overview

Multicast traffic from the source is streamed towards router CE-1. Receivers connected to PE-2, PE-3 and PE-4 are interested in joining this multicast group.

CEs 1 to 4 are PIM enabled routers, which form a PIM adjacency with nearest PE. Between the PEs across the provider network there are no PIM adjacencies since BGP auto-discovery and BGP signalling are used. Selective PMSI using mLDP or RSVP P2MP are out of the scope of this section. Selective PMSI using PIM SSM is supported but cannot be used when I-PMSI is mLDP or RSVP with R9.0. I-PMSI and S-PMSI must use same tunnelling technology, either PIM/GRE or mLDP or RSVP P2MP.

## Configuration

The configuration is divided into the following sections:

- Provider Common Configuration
  - PE Global Configuration
- PE VPRN Configuration and PE VPRN Multicast Configuration for NG-MVPN
  - PMSI using mLDP
  - PMSI using RSVP-TE
  - UMH (Upstream Multicast Hop)
- PE VPRN Configuration and PE VPRN Multicast Configuration for Draft-Rosen using MDT AFI SAFI
  - Auto discovery using BGP MDT AFI SAFI as per Draft-Rosen version 9 with MDT using PIM SSM

## Provider Common Configuration

### PE Global Configuration

This section describes the common configuration required for each PE within the provider multicast domain, regardless of the MVPN PE auto-discovery or customer signaling methods. This includes Interior Gateway Protocol (IGP) and VPRN service configuration.

The configuration tasks can be summarized as follows:

- PE global configuration.
 

This includes configuration of the IGP (ISIS will be used); configuration of link layer LDP between PEs (LDP will be used here as the method to interconnect VPRNs); configuration of iBGP between PEs to facilitate VPRN route learning.
- VPRN configuration on the PEs.
 

This includes configuration of basic VPRN parameters (route-distinguisher, route target communities), configuration of attachment circuits towards CEs, configuration of VRF routing protocol and any routing policies.
- PIM within the VRF and MVPN parameters — I-PMSI
- CE configuration.

## Provider Common Configuration

**Step 1.** Configure the interfaces, the IGP (ISIS) in all PE nodes (where ISIS redistributes route reachability to all routers) and LDP in the interfaces (link layer LDP). To facilitate the ISIS configuration, all routers are Level2-Level1 capable within the same ISIS area-id, so there is only a single topology area in the network (all routers share the same topology). The configuration for PE-2 is displayed below.

```
PE-2>configure router
  interface "int-PE-2-PE-1"
    address 192.168.12.2/30
    port 1/1/2:1
  exit
  interface "int-PE-2-PE-4"
    address 192.168.24.1/30
    port 1/1/3:1
  exit
  interface "system"
    address 192.0.2.2/32
  exit
  autonomous-system 64496
  isis
    area-id 49.0001
    traffic-engineering
    interface "system"
      passive
    exit
    interface "int-PE-2-PE-1"
      interface-type point-to-point
    exit
    interface "int-PE-2-PE-4"
      interface-type point-to-point
    exit
    no shutdown
  exit
  ldp
    interface-parameters
      interface "int-PE-2-PE-1"
        exit
      interface "int-PE-2-PE-4"
        exit
    exit
    targeted-session
  exit
```

The configuration for the rest of nodes is similar. The IP addresses can be derived from [Figure 95](#).



**Step 2.** Verify that ISIS adjacencies and LDP peer sessions are formed.

```
*A:PE-1# show router isis adjacency
=====
ISIS Adjacency
=====
System ID                Usage State Hold Interface                MT Enab
-----
PE-2                     L1L2  Up    30  int-PE-1-PE-2                No
PE-3                     L1L2  Up    24  int-PE-1-PE-3                No
-----
Adjacencies : 2
=====
*A:PE-1# show router ldp session
=====
LDP Sessions
=====
Peer LDP Id              Adj Type   State           Msg Sent  Msg Recv  Up Time
-----
192.0.2.2:0             Both      Established     235      237      0d 00:08:57
192.0.2.3:0             Both      Established     237      245      0d 00:08:58
-----
No. of Sessions: 2
```

**Step 3.** Configure iBGP full mesh between the PEs for VPRN routing (Route Reflectors could also be an option).

```
*A:PE-1>config>router>bgp# info
-----
min-route-advertisement 1
rapid-withdrawal
rapid-update mvpn-ipv4 mdt-safi
group "internal"
  family vpn-ipv4 mvpn-ipv4 mdt-safi
  type internal
  neighbor 192.0.2.2
  exit
  neighbor 192.0.2.3
  exit
  neighbor 192.0.2.4
  exit
exit
no shutdown
-----
```

Note that the families configured under the group **internal** are vpn-ipv4, mvpn-ipv4 and mdt-safi, since the three families are referenced in this chapter.

Note that the mdt-safi parameter is not needed for NG-MVPN (mLDP/RSVP scenarios) and is only required for Draft-Rosen with MDT AFI SAFI.

Rapid withdrawal (configured on all PEs) disables the Minimum Route Advertisement Interval (MRAI) interval on sending BGP withdrawals. Rapid update (configured for MVPN-IPv4 and MDT AFI/SAFI address families) disables the MRAI interval on sending BGP update messages for the address family MVPN-IPv4 and MDT AFI/SAFI).

**Step 4.** Verify that BGP peer relationships are established.

```
*A:PE-1# show router bgp summary
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
BGP Admin State      : Up      BGP Oper State      : Up
Total Peer Groups    : 1      Total Peers          : 3
* Truncated info
=====
BGP Summary
=====
Neighbor
      AS PktRcvd InQ  Up/Down  State|Rcv/Act/Sent (Addr Family)
      PktSent OutQ
-----
192.0.2.2
      64496      28   0 00h12m04s 4/4/3 (VpnIPv4)
      29   0   1/1/1 (MvpnIpv4)
      0/0/0 (MdtSafi)
192.0.2.3
      64496      27   0 00h12m11s 2/2/3 (VpnIPv4)
```

```
192.0.2.4          30    0          1/1/1 (MvpnIpv4)
                  0/0/0 (MdtSafi)
64496             30    0 00h12m04s 2/2/3 (VpnIPv4)
                  29    0          2/2/1 (MvpnIpv4)
                  0/0/0 (MdtSafi)
-----
```

## PE VPRN Configuration and PE VPRN Multicast Configuration

A VPRN is created on each PE per service (the different services using mLDP, RSVP-TE and AFI/SAFI with PIM), these are the multicast VPRNs. PE-1 is the PE containing the attachment circuit towards CE-1. CE-1 is the CE nearest to the source. PE-2, PE-3 and PE-4 contain attachment circuits towards CE-2, CE-3 and CE-4 respectively. Each CE has a receiving host attached.

### PMSI using mLDP

Figure 96 shows the details of the topology for VPRN 1.

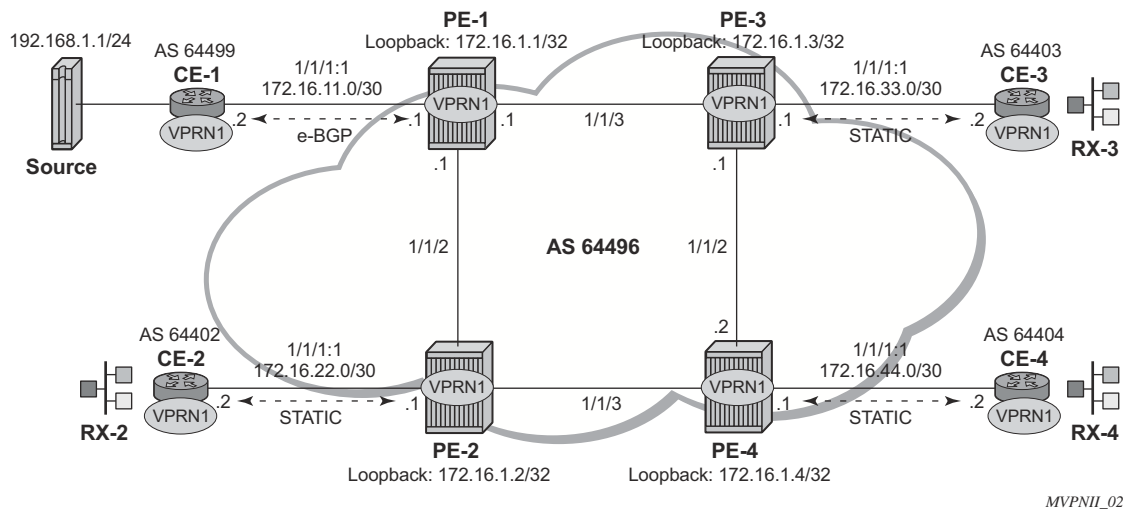


Figure 96: VPRN 1 Topology used for mLDP

## Unicast

**Step 1.** Create VPRN 1 on each PE, containing a route-distinguisher of 64496:10X (where X= number of PE) and vrf-target of 64496:100. The autonomous system number is 64496. For the next hop tunnel route resolution to connect the VPRNs between PEs, manually configured spoke SDPs are created (note that other methods like auto-bind LDP, RSVP-TE or auto-bind MPLS could also be used). LDP was already enabled in the overview section.

```
*A:PE-1>config>service# info
-----
customer 1 create
  description "Default customer"
exit
sdp 12 mpls create
  far-end 192.0.2.2
  ldp
  keep-alive
  shutdown
exit
no shutdown
exit
sdp 13 mpls create
  far-end 192.0.2.3
  ldp
  keep-alive
  shutdown
exit
no shutdown
exit
sdp 14 mpls create
  far-end 192.0.2.4
  ldp
  keep-alive
  shutdown
exit
no shutdown
exit
vprn 1 customer 1 create
  description "mLDP"
  autonomous-system 64496
  route-distinguisher 64496:101
  vrf-target target:64496:100

  spoke-sdp 12 create
    no shutdown
  exit
  spoke-sdp 13 create
    no shutdown
  exit
  spoke-sdp 14 create
    no shutdown
  exit
```

## Provider Common Configuration

**Step 2.** Create an attachment circuit interface towards the CE and a loopback (the loopback is not mandatory but it is configured to aid troubleshooting the routers).

```
PE-1>configure service vprn 1
    interface "loopback" create
        address 172.16.1.1/32
        loopback
    exit
    interface "int-PE-1-CE-1" create
        address 172.16.11.1/30
        sap 1/1/1:1 create
    exit
exit
```

**Step 3.** The source address of the multicast stream will need to be reachable by all routers (PEs and CEs) within the VPN. This will be advertised within BGP from CE-1 to PE-1. Create a BGP peering relationship with the CE.

```
PE-1>configure service vprn 1
    bgp
        group "external"
            type external
            peer-as 64499
            neighbor 172.16.11.2
        exit
    exit
    no shutdown
exit
```

**Step 4.** On CE-1, create a VPRN to support the connection of the source to CE-1 and to connect CE-1 to PE-1. Two attachment circuits are required as well as a BGP peering relationship with the PE. This uses a default BGP address family of ipv4.

```
A:CE-1>config>service>vprn# info
-----
autonomous-system 64499
route-distinguisher 64499:1
interface "int-CE-1-PE-1" create
    *Truncated info
interface "source" create
    *Truncated info
exit

bgp
    export "source"
    group "external"
        type external
        peer-as 64496
        neighbor 172.16.11.1
    exit
    exit
    no shutdown
exit
no shutdown
-----
```

**Step 5.** In order for the subnet on the CE connecting to the source to be advertised within BGP, a route policy is required. The subnet containing the multicast source is 192.168.1.0/24, so a prefix-list can be used to define a match, and then used within a route policy to inject into BGP. In this example, only the host (/32) of the source is advertised.

```
A:CE-1>config>router# info
* Truncated info
  policy-options
  begin
  prefix-list "source"
    prefix 192.168.1.1/32 exact
  exit
  policy-statement "source"
    entry 10
      from
        prefix-list "source"
      exit
      to
        protocol bgp
      exit
      action accept
      exit
    exit
  exit
  commit
exit
```

**Step 6.** Check that the route is seen in PE-1:

```
*A:PE-1>config>service# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type Proto Age Pref
Metric
-----
* Truncated info
192.168.1.1/32 Remote BGP 00h04m15s 170
172.16.11.2 0
-----
No. of Routes: 10
Flags: L = LFA nexthop available B = BGP backup route available
n = Number of times nexthop is repeated
=====
```

This prefix will also be automatically advertised within the BGP VPRN to all other PEs, and will be installed in VRF 1.

For example, on PE-4, the source 192.168.1.1/32 is received via BGP VPN with a next-hop of PE-1 (192.0.2.1):

```
*A:PE-4>config>service>vprn# show router 1 route-table
=====
```

## Provider Common Configuration

```
Route Table (Service: 1)
=====
Dest Prefix[Flags]                Type   Proto   Age      Pref
  Next Hop[Interface Name]                Metric
-----
* Truncated info
192.168.1.1/32                    Remote BGP VPN 00h00m05s 170
  192.0.2.1 (tunneled)                0
-----
No. of Routes: 10
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
```

Each CE containing a multicast receiver must be able to reach the source. As an example on CE-2, a static route will suffice and is configured with next hop of the PE-2 PE-CE interface.

```
A:CE-2# configure service vprn 1
...
static-route 192.168.1.1/32 next-hop 172.16.22.1
```

After **Steps 1 to 6**, all required unicast routing is provisioned. The following sections show the configuration of the multicast in the VPRN.



## Auto-Discovery and mLDP PMSI Establishment

The MP-BGP based auto-discovery is implemented with a new address family defined in RFC 4760 MP\_REACH\_NRLI/MP\_UNREACH\_NRLI attributes, with AFI 1 (IPv4) or 2 (IPv6) SAFI 5 (temporary value assigned by IANA). This is the mechanism by which each PE advertises the presence of an MVPN to other PEs. This can be achieved using PIM (like in Draft-Rosen) or using BGP. With the default parameter, BGP is automatically chosen because the PMSIs are mLDP and PIM is not an option in this case. Any PE that is a member of an MVPN will advertise to the other PEs using a BGP Multi-protocol Reachable Next-Hop Router Layer Information (NRLI) update that is sent to all PEs within the AS. This update will contain an Intra-AS I-PMSI auto-discovery Route type, also known as an Intra-AD. These use an address family mvpn-ipv4, so each PE must be configured to originate and accept such updates (note this was done earlier when configuring the families).

At this step (auto-discovery), the information about the PMSI is exchanged but the PMSI is not instantiated.

As each PE contains a CE which will be part of the multicast VRF, it is necessary to enable PIM on each interface containing the attachment circuit towards a CE, and to configure the I-PMSI multicast tunnel for the VRF. Note that S-PMSIs are not supported for mLDP with the 9.0R5 software release. In order for the BGP routes to be accepted into the VRF, a route-target community is required (vrf-target). This is configured in the **configure service vprn 1 mvpn** context and, in this case is set to the same value as the unicast vrf-target (the vrf-target community as the **configure service vprn 1 vrf-target** context).

On each PE, a VPRN instance is configured as follows:

```
A:PE-4# configure service vprn 1
...
pim
    interface "loopback"
    exit
    interface "int-PE-4-CE-4"
    exit
    rp
        static
        exit
        bsr-candidate
            shutdown
        exit
        rp-candidate
            shutdown
        exit
    exit
    no shutdown
exit
mvpn
    auto-discovery default
    c-mcast-signaling bgp
    exit
    provider-tunnel
        inclusive
```

## Provider Common Configuration

```
        mldp
        no shutdown
    exit
    exit
    exit
    vrf-target unicast
    exit
exit
```

When PIM SSM is used, the configuration always shows RP static with no RP entries (this is enabled by default when PIM is provisioned). In order for the BGP routes to be accepted into the VRF, a route-target community is required (vrf-target). Although it is not mandatory for the mvpn target to be equal to the unicast target, the recommendation is to use **vrf-target unicast** to avoid configuration mistakes and extra complexity.

The status of VPRN 1 on PE-1 is shown with the following output:

```
*A:PE-1# show router 1 mvpn
=====
MVPN 1 configuration data
=====
signaling          : Bgp          auto-discovery     : Default
UMH Selection      : Highest-IP    intersite-shared   : Enabled
vrf-import         : N/A
vrf-export         : N/A
vrf-target         : unicast
C-Mcast Import RT : target:192.0.2.1:2

ipmsi              : ldp
i-pmsi P2MP AdmSt  : Up

s-pmsi             : none
data-delay-interval: 3 seconds
enable-asm-mdt     : N/A
=====
```

The following shows a debug of an Intra-AD BGP update message received by PE-1 that was sent by PE-2. The message contains the PMSI tunnel type to be used (LDP P2MP LSP), LSP identification (root node, opaque value) and the type of BGP update (Type: Intra-AD Len: 12 RD: 64496:102 Orig: 192.0.2.2):

```
A 4 2011/10/06 01:25:42.81 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 91
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 8 Len: 4 Community:
    no-export
  Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:64496:100
```

```

Flag: 0xc0 Type: 22 Len: 22 PMSI:
  Tunnel-type LDP P2MP LSP (2)
  Flags [Leaf not required]
  MPLS Label 0
  Root-Node 192.0.2.2, LSP-ID 0x2001
Flag: 0x90 Type: 14 Len: 23 Multiprotocol Reachable NLRI:
  Address Family MVPN_IPV4
  NextHop len 4 NextHop 192.0.2.2
  Type: Intra-AD Len: 12 RD: 64496:102 Orig: 192.0.2.2
"

```

The set up has four PEs, so every PE should see each others peer Intra-AD route; the output below shows the routes received in PE-1:

```

*A:PE-1# show router bgp routes mvpn-ipv4 type intra-ad
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      LocalPref  MED
      RD              SourceAS          VPNLabel
      Nexthop        SourceIP
      As-Path        GroupIP
-----
u*>i  Intra-Ad          192.0.2.2        100        0
      64496:102      -                -
      192.0.2.2      -                -
      No As-Path    -
u*>i  Intra-Ad          192.0.2.3        100        0
      64496:103      -                -
      192.0.2.3      -                -
      No As-Path    -
u*>i  Intra-Ad          192.0.2.4        100        0
      64496:104      -                -
      192.0.2.4      -                -
      No As-Path    -
-----
Routes : 3
=====

```

The detailed output of the Intra-AD received from PE-2 shows the Tunnel-Type LDP P2MP LSP (LSP-ID is 8193), the originator id (192.0.2.2), and the route-distinguisher (64496:102):

```

*A:PE-1# show router bgp routes mvpn-ipv4 type intra-ad detail
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP MVPN-IPv4 Routes

```

## Provider Common Configuration

```

=====
Route Type      : Intra-Ad
Route Dist.    : 64496:102
Originator IP  : 192.0.2.2
Nexthop        : 192.0.2.2
From           : 192.0.2.2
Res. Nexthop   : 0.0.0.0
Local Pref.    : 100
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
Community      : no-export target:64496:100
Cluster        : No Cluster Members
Originator Id  : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
VPRN Imported  : 1
Interface Name : NotAvailable
Aggregator     : None
MED            : 0
Peer Router Id : 192.0.2.2
-----
PMSI Tunnel Attribute :
Tunnel-type   : LDP P2MP LSP
MPLS Label    : 0
Root-Node     : 192.0.2.2
Flags         : Leaf not required
LSP-ID       : 8193
-----

```

Because of the receiver-driven nature of mLDP, mLDP P2MP LSPs are setup unsolicited from the leaf PEs towards the head-end PE. The leaf PEs discover the head-end PE via I-PMSI/S-PMSI A-D routes. The Tunnel Identifier carried in the PMSI attribute is used as the P2MP Forwarding Equivalence Class (FEC) Element. The Tunnel Identifier consists of the head-end PE's address, along with a Generic LSP Identifier value. The Generic LSP Identifier value is automatically generated by the head-end PE. The previous show command displays the PMSI information with the detail of the Root-Node (192.0.2.2) and the LSP-ID (8193). The PMSI was created after receiving the A-D message from PE-2, where the following excerpt from the previous debug shows the same information (x2001 in HEX is equal to 8193 in decimal).

```

Flag: 0xc0 Type: 22 Len: 22 PMSI:
  Tunnel-type LDP P2MP LSP (2)
  Flags [Leaf not required]
  MPLS Label 0
  Root-Node 192.0.2.2, LSP-ID 0x2001

```

Once the mLDP P2MP LSPs are created, the I-PMSI is instantiated in the core:

```

*A:PE-1# show router 1 pim neighbor
=====
PIM Neighbor ipv4
=====
Interface          Nbr DR Prty    Up Time        Expiry Time    Hold Time
  Nbr Address
-----
int-PE-1-CE-1      1              0d 01:21:30   0d 00:01:16   105
  172.16.11.2
mpls-if-73729      1              0d 01:21:03   never           65535
  192.0.2.2
mpls-if-73730      1              0d 01:20:47   never           65535

```

```

192.0.2.3
mpls-if-73731      1          0d 01:20:32  never          65535
192.0.2.4
-----
Neighbors : 4
=====

*A:PE-1# show router 1 pim tunnel-interface
=====
PIM Interfaces ipv4
=====
Interface          Adm  Opr  DR Prty          Hello Intvl  Mcast Send
DR
-----
mpls-if-73728      Up   Up   N/A             N/A          N/A
192.0.2.1
mpls-if-73729      Up   Up   N/A             N/A          N/A
192.0.2.2
mpls-if-73730      Up   Up   N/A             N/A          N/A
192.0.2.3
mpls-if-73731      Up   Up   N/A             N/A          N/A
192.0.2.4
-----
Interfaces : 4
=====

```

Every PE has created an I-PMSI to the other PEs. Checking the mLDP P2MP LSPs that are originated, transit, or destination to PE-1:

```

*A:PE-1# show router ldp bindings fec-type p2mp active p2mp-id 8193 root 192.0.2.1
=====
LDP P2MP Bindings (Active)
=====
P2MP-Id      RootAddr
Interface    Op          IngLbl      EgrLbl  EgrIntf/  EgrNextHop
              Op          EgrIntf/   LspId
-----
8193         192.0.2.1
73728       Push
8193         192.0.2.1
73728       Push
8193         192.0.2.2
73729       Pop         262137      --      --         --
8193         192.0.2.2
73729       Swap       262137      262136  1/1/3:1   192.168.13.2
8193         192.0.2.3
73730       Swap       262142      262136  1/1/2:1   192.168.12.2
8193         192.0.2.3
73730       Pop         262142      --      --         --
8193         192.0.2.4
73731       Pop         262136      --      --         --
-----
No. of P2MP Active Bindings: 7
=====

```

The two first entries in the output show the P2MP LSP where PE-1 is the root headend (Push). The other two entries (Swap and Pop) correspond with transit and leaf for the P2MP LSPs originated

## Provider Common Configuration

by the other PEs. The command shows a P2MP-ID (8193) with an interface 73728 (matches with the **show router 1 pim tunnel interface** being the PIM interface created from PE-1) with two egress interfaces pointing to PE-2 and PE-3.

A similar command executed on PE-2 shows:

```
*A:PE-2# show router ldp bindings fec-type p2mp p2mp-id 8193 root 192.0.2.1 active
=====
LDP P2MP Bindings (Active)
=====
P2MP-Id      RootAddr
Interface    Op          IngLbl      EgrLbl  EgrIntf/   EgrNextHop
                               LspId
-----
8193         192.0.2.1
73729       Pop        262137      --      --         --
8193         192.0.2.1
73729       Swap      262137      262135  1/1/3:1    192.168.24.2
* Truncated info
-----
No. of P2MP Active Bindings: 7
=====
```

On PE-2, the first entry shows that PE-2 is a leaf of the P2MP LSP tree created by PE-1 (ingress label is 262137 which was the egress label to reach PE-2 and is popped). However, the second entry shows that PE-2 is transit for the P2MP LSP going to PE-4 (ingress label 262137, egress label 262135 next hop PE-4).

The same command on PE-4 shows:

```
*A:PE-4# show router ldp bindings active fec-type p2mp p2mp-id 8193 root 192.0.2.1
=====
LDP P2MP Bindings (Active)
=====
P2MP-Id      RootAddr
Interface    Op          IngLbl      EgrLbl  EgrIntf/   EgrNextHop
                               LspId
-----
8193         192.0.2.1
73731       Pop        262135      --      --         --
* Truncated info
-----
No. of P2MP Active Bindings: 5
=====
```

In the first entry the root is PE-1 and the action is Pop, being the ingress label 262135, showing that this is another leaf for the P2MP LSP started on PE-1.

To complete the information, checking on PE-3, the first entry there is a Pop where the root is PE-1, and the ingress label is 262137:

```
*A:PE-3# show router ldp bindings active fec-type p2mp p2mp-id 8193 root 192.0.2.1
```

```

=====
LDP P2MP Bindings (Active)
=====
P2MP-Id      RootAddr
Interface    Op          IngLbl     EgrLbl  EgrIntf/  EgrNextHop
                               LspId
-----
8193         192.0.2.1
73729       Pop        262137    --      --        --
* Truncated info
-----
No. of P2MP Active Bindings: 5
=====

```

As a summary, each root PE has a P2MP LSP with three leaves (the other PEs) and they are also transit points to the P2MP LSPs created in the other PEs. As an additional check, an OAM ping can show the different leaves that a P2MP LSP has:

```

*A:PE-1# oam p2mp-lsp-ping ldp 8193 sender-addr 192.0.2.1 detail
P2MP identifier 8193: 88 bytes MPLS payload
=====
Leaf Information
=====
From          RTT          Return Code
-----
192.0.2.2     =3.03ms     EgressRtr(3)
192.0.2.4     =5.17ms     EgressRtr(3)
192.0.2.3     =33.1ms     EgressRtr(3)
=====
Total Leafs responded = 3
          round-trip min/avg/max = 3.03 / 13.8 / 33.1 ms

Responses based on return code:
          EgressRtr(3)=3

```

An easy way to see the path that the LDP P2MP LSP follows for a specific leaf is the following command (such as LDP trace from PE-1 to PE-4):

```

*A:PE-1# oam ldp-treetrace prefix 192.0.2.4/32

ldp-treetrace for Prefix 192.0.2.4/32:

          192.168.24.2, ttl = 2 dst = 127.1.0.255 rc = EgressRtr status = Done
Hops:          192.168.12.2

ldp-treetrace discovery state: Done
ldp-treetrace discovery status: ' OK '
Total number of discovered paths: 1
Total number of failed traces: 0

```

The command shows that on PE-4 there is an active leaf of the P2MP LSP, and that there is an intermediate hop on PE-2.

### Traffic Flow

The receiver RX-4, connected to CE-4, wishes to join the group 232.1.1.1 with source 192.168.1.1 and so sends an IGMPv3 report towards CE-4. CE-4 recognizes the report and sends a PIM join towards the source, hence it reaches PE-1 where the source is connected to through CE-1. The output below shows the debug seen on PE-4, where the PIM join is received from CE-4 and a BGP update Source Join is sent to all PEs (note that only the update sent to PE-1 is shown).

```
11 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPProcessSG
pimJPPProcessSG: (S,G)-> (192.168.1.1,232.1.1.1) type <S,G>, i/f int-PE-4-CE-4, u
pNbr 172.16.44.1 isJoin 1 isRpt 0 holdTime 210"

12 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Downstream: State NoInfo Event RxJoin, (S,G) (192.168.1.1,232.1.1.1) grou
pType <S,G>"

13 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Upstream: State NotJoined Event JoinDesiredTrue, (S,G) (192.168.1.1,232.1
.1.1) groupType <S,G>"

14 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpJoinDesiredTrue
No upstream interface. pSG (192.168.1.1,232.1.1.1) rpfType 3"

15 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpJoinDesiredTrue
pim 2 sg_type 2 refetch route type SPMSI pendingFetchMask 0x8"

16 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpJoinDesiredTrue
No upstream interface. pSG 0x5578f87c, (192.168.1.1,232.1.1.1) rpfType 3"

17 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Upstream: State Joined Event MribChange, (S,G) (192.168.1.1,232.1.1.1) gr
oupType <S,G>"

18 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpStateJMribChange
pSG 0x5578f87c, (192.168.1.1,232.1.1.1), type <S,G> oldMribNhopIp 0.0.0.0 oldRpf
NbrIp 0.0.0.0, oldRpfType NONE oldRpfIf 0 rptMribNhopIp 0.0.0.0, rptRpfNbrIp 0.0
.0.0 rtmReason 32"

19 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpStateJMribChange
pSG 0x5578f87c, (192.168.1.1,232.1.1.1), type <S,G> newMribNhopIp 192.0.2.1 newR
pNbrIp 192.0.2.1 newRpfType REMOTE newRpfIf 73731"

20 2011/10/13 15:41:30.83 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 69
  Flag: 0x40 Type: 1 Len: 1 Origin: 0"
```



```

Flag: 0x40 Type: 2 Len: 0 AS Path:
Flag: 0x80 Type: 4 Len: 4 MED: 0
Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:192.0.2.1:2
Flag: 0x90 Type: 14 Len: 33 Multiprotocol Reachable NLRI:
    Address Family MVPN_IPV4
    NextHop len 4 NextHop 192.0.2.4
    Type: Source-Join Len:22 RD: 64496:101 SrcAS: 64496 Src: 192.168.1.1 Grp
: 232.1.1.1
"

```

The following debug shows that PE-1 receives the BGP update Source Join with source 192.168.1.1 and group 232.1.1.1 and sends a PIM join towards CE-1:

```

57 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 69
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:192.0.2.1:2
    Flag: 0x90 Type: 14 Len: 33 Multiprotocol Reachable NLRI:
        Address Family MVPN_IPV4
        NextHop len 4 NextHop 192.0.2.4
        Type: Source-Join Len:22 RD: 64496:101 SrcAS: 64496 Src: 192.168.1.1 Grp
: 232.1.1.1
"

58 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPProcessSG
pimJPPProcessSG: (S,G)-> (192.168.1.1,232.1.1.1) type <S,G>, i/f mpls-if-73728, u
pNbr 192.0.2.1 isJoin 1 isRpt 0 holdTime 65535"

59 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Downstream: State NoInfo Event RxJoin, (S,G) (192.168.1.1,232.1.1.1) grou
pType <S,G>"

60 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Upstream: State NotJoined Event JoinDesiredTrue, (S,G) (192.168.1.1,232.1
.1.1) groupType <S,G>"

61 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSGUpJoinDesiredTrue
pim 2 sg_type 2 refetch route type SPMSI pendingFetchMask 0x8"

62 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimSendJoinPrunePdu
pimSendJoinPrunePdu: if 3, adj 172.16.11.2"

63 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]

```

## Provider Common Configuration

```
"PIM[Instance 2 vprn1]: pimSGEncodeGroupSet
pimEncodeGroupSet: encoding groupset for group 232.1.1.1, numJoinedSrcs 1, numPrunedSrcs 0"
```

```
64 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
```

```
"PIM[Instance 2 vprn1]: pimSGEncodeGroupSet
pimEncodeGroupSet: Encoding Join for source 192.168.1.1"
```

```
65 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
```

```
"PIM[Instance 2 vprn1]: pimSGEncodeGroupSet
pimEncodeGroupSet: num joined srcs 1, num pruned srcs 0"
```

```
66 2011/10/14 02:06:58.22 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
```

```
"PIM[Instance 2 vprn1]: pimSendJoinPrunePdu
pimSendJoinPrunePdu2: sending JP PDU with 1 groups."
```

The BGP update source join received by PE-1 is displayed with the command:

```
*A:PE-1# show router bgp routes mvpn-ipv4 type source-join
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      LocalPref  MED
      RD            SourceAS          SourceAS    VPNLabel
      Nexthop       SourceIP          SourceIP
      As-Path       GroupIP
-----
u*>i  Source-Join      -                 100        0
      64496:101      64496            64496      -
      192.0.2.4     192.168.1.1     192.168.1.1
      No As-Path    232.1.1.1
-----
Routes : 1
=====
```

To verify the traffic: on PE-1 there is a group 232.1.1.1 with source 192.168.1.1, the Reverse Path Forwarding (RPF) is CE-1, the multicast traffic is flowing from CE-1 to PE-1 using int-PE-1-CE-1 and the outgoing interface is using the PMSI mLDP mpls-if-73728.

```
*A:PE-1# show router 1 pim group detail
=====
PIM Source Group ipv4
=====
Group Address      : 232.1.1.1
Source Address     : 192.168.1.1
RP Address         : 0
Flags              :
MRIB Next Hop     : 172.16.11.2
MRIB Src Flags    : remote
Up Time           : 0d 00:12:04
Type               : (S,G)
Keepalive Timer   : Not Running
Resolved By       : rtable-u
```

```
Up JP State      : Joined          Up JP Expiry      : 0d 00:00:56
Up JP Rpt       : Not Joined StarG Up JP Rpt Override : 0d 00:00:00
```

```
Register State  : No Info
Reg From Anycast RP: No
```

```
Rpf Neighbor    : 172.16.11.2
Incoming Intf   : int-PE-1-CE-1
Outgoing Intf List : mpls-if-73728
```

```
Curr Fwding Rate : 0.0 kbps
Forwarded Packets : 0                Discarded Packets : 0
Forwarded Octets  : 0                RPF Mismatches    : 0
Spt threshold     : 0 kbps           ECMP opt threshold : 7
Admin bandwidth   : 1 kbps
```

```
-----
Groups : 1
=====
```

On PE-4, the same (S,G) arrives in the incoming interface mpls-if-73731 and the outgoing interface is int-PE-4-CE-4.

```
*A:PE-4# show router 1 pim group detail
```

```
=====
PIM Source Group ipv4
=====
```

```
Group Address      : 232.1.1.1
Source Address     : 192.168.1.1
RP Address         : 0
Flags              :
Type               : (S,G)
MRIB Next Hop     : 192.0.2.1
MRIB Src Flags    : remote
Up Time           : 0d 00:15:44
Keepalive Timer   : Not Running
Resolved By       : rtable-u
```

```
Up JP State      : Joined          Up JP Expiry      : 0d 00:00:16
Up JP Rpt       : Not Joined StarG Up JP Rpt Override : 0d 00:00:00
```

```
Register State  : No Info
Reg From Anycast RP: No
```

```
Rpf Neighbor    : 192.0.2.1
Incoming Intf   : mpls-if-73731
Outgoing Intf List : int-PE-4-CE-4
```

```
Curr Fwding Rate : 0.0 kbps
Forwarded Packets : 0                Discarded Packets : 0
Forwarded Octets  : 0                RPF Mismatches    : 0
Spt threshold     : 0 kbps           ECMP opt threshold : 7
Admin bandwidth   : 1 kbps
```

When the receiver is not interested in the channel group any more, the receiver RX-4 sends an IGMPv3 leave, PE-4 sends a PIM prune translated to a BGP MP\_UNREACH NLRI to all PEs. Note that, as mentioned before, rapid withdrawals are sent without waiting for the mrai (note that for simplicity, only one BGP update is shown in the output debug).

## Provider Common Configuration

```
33 2011/10/13 16:20:17.74 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPProcessSG
pimJPProcessSG: (S,G)-> (192.168.1.1,232.1.1.1) type <S,G>, i/f int-PE-4-CE-4, u
pNbr 172.16.44.1 isJoin 0 isRpt 0 holdTime 210"

34 2011/10/13 16:20:17.74 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Downstream: State Joined Event RxPrune, (S,G) (192.168.1.1,232.1.1.1) gro
upType <S,G>"

35 2011/10/13 16:20:17.74 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Downstream: State PrunePending Event PrunePendTimerExp, (S,G) (192.168.1.
1,232.1.1.1) groupType <S,G>"

36 2011/10/13 16:20:17.74 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 2 vprn1]
"PIM[Instance 2 vprn1]: pimJPPrintFsmEvent
PIM JP Upstream: State Joined Event JoinDesiredFalse, (S,G) (192.168.1.1,232.1.1
.1) groupType <S,G>"

37 2011/10/13 16:20:17.74 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 31
  Flag: 0x90 Type: 15 Len: 27 Multiprotocol Unreachable NLRI:
    Address Family MVPN_IPV4
      Type: Source-Join Len:22 RD: 64496:101 SrcAS: 64496 Src: 192.168.1.1 Grp
: 232.1.1.1
"
```

PMSI using RSVP-TE

Figure 97 shows the details of the topology for VPRN 2.

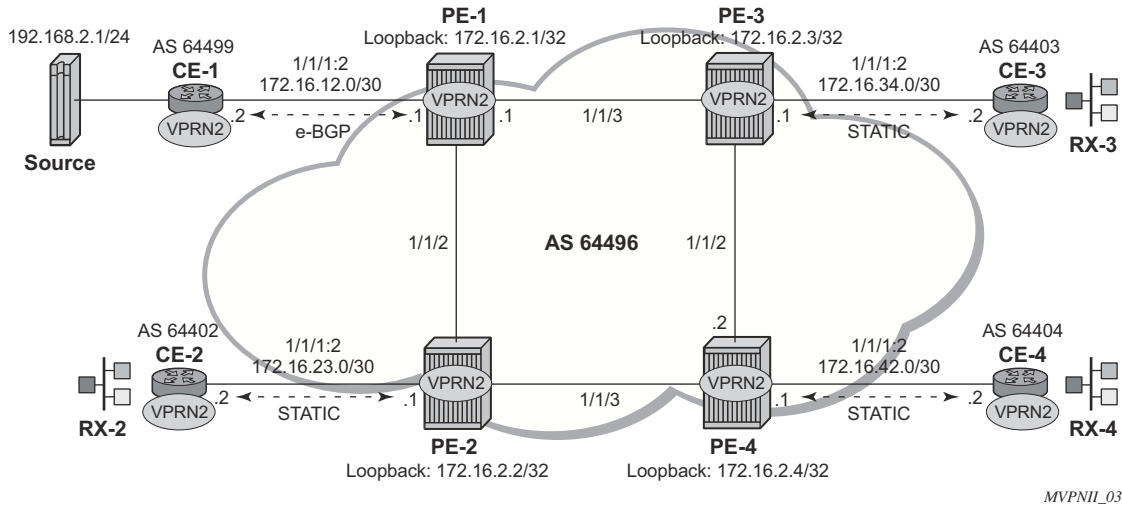


Figure 97: VPRN 2 Topology used for RSVP-TE P2MP

Unicast

For the sake of simplicity, check **Steps 1 to 6** in [PMSI using mLDP on page 704](#) for VPRN 2 creation information. The same steps are repeated for RSVP using for provisioning the details that appear in [Figure 97](#). The result is the configuration in all the PEs, taking as an example PE-1:

```
*A:PE-1>config>service>vprn# info
-----
description "P2MP RSVP"
autonomous-system 64496
route-distinguisher 64496:201
vrf-target target:64496:200
interface "loopback" create
    address 172.16.2.1/32
    loopback
exit
interface "int-PE-1-CE-1" create
    address 172.16.12.1/30
    sap 1/1/1:2 create
    exit
exit
bgp
    group "external"
```

## Provider Common Configuration

```
        type external
        peer-as 64499
        neighbor 172.16.12.2
        exit
    exit
    no shutdown
exit

spoke-sdp 12 create
    no shutdown
exit
spoke-sdp 13 create
    no shutdown
exit
spoke-sdp 14 create
    no shutdown
exit
no shutdown
```

---

In addition to the unicast, because RSVP is the signalling protocol to establish the P2MP LSPs, RSVP is configured on the interfaces. In addition, to use P2MP RSVP an LSP template is needed. The template defines the characteristics of the LSP to be created, for example, make-before-break, bandwidth, administrative groups, cspf, specific paths, etc. A basic template is used here. TE parameters specified in the template are commonly used in each RSVP PATH message for each of the branches of the P2MP RSVP LSP. The template is used in the mvpn context within the VPRN configuration (see [Auto-Discovery and RSVP PMSI Establishment on page 723](#)). The resignal timer for P2MP is configured to the minimum value of sixty minutes (60 — 10080 minutes):

```
*A:PE-1>config>router>mpls# info
```

---

```
p2mp-resignal-timer 60
interface "system"
exit
interface "int-PE-1-PE-2"
exit
interface "int-PE-1-PE-3"
exit
path "empty"
    no shutdown
exit
lsp-template "vrf2" p2mp
    default-path "empty"
    cspf
    fast-reroute facility
    exit
    no shutdown
exit
no shutdown
```

---

## Auto-Discovery and RSVP PMSI Establishment

The MP-BGP based auto-discovery is implemented with a new address family defined in RFC 4760 MP\_REACH\_NRLI/MP\_UNREACH\_NRLI attributes, with AFI 1 (IPv4) or 2 (IPv6) SAFI 5 (temporary value assigned by IANA). This is the mechanism by which each PE advertises the presence of an MVPN to other PEs. This can be achieved using PIM (like in Draft-Rosen) or using BGP. With the default parameter, BGP is automatically chosen because the PMSIs are RSVP and PIM is not an option in this case. Any PE that is a member of an MVPN will advertise to the other PEs using a BGP Multi-protocol Reachable Next-Hop Router Layer Information (NRLI) update that is sent to all PEs within the AS. This update will contain an Intra-AS I-PMSI auto-discovery Route type, also known as an Intra-AD. These use an address family mvpn-ipv4, so each PE must be configured to originate and accept such updates (note this was done earlier when configuring the families).

At this step (auto-discovery), the information about the PMSI is exchanged but the PMSI is not instantiated.

As each PE contains a CE which will be part of the multicast VRF, it is necessary to enable PIM on each interface containing the attachment circuit towards a CE, and to configure the I-PMSI multicast tunnel for the VRF. Note that S-PMSIs are not supported for RSVP with the 9.0R5 software release. In order for the BGP routes to be accepted into the VRF a route-target community is required (vrf-target). Although it is not mandatory for the mVPN target to be equal to the unicast target, the recommendation is to use vrf-target unicast to avoid configuration mistakes and extra complexity.

On each PE, a VPRN instance is configured as follows:

```
*A:PE-1>config>service>vprn# info
-----
      pim
      interface "loopback"
      exit
      interface "int-PE-1-CE-1"
      exit
      rp
          static
          exit
          bsr-candidate
              shutdown
          exit
          rp-candidate
              shutdown
          exit
      exit
      no shutdown
  exit
mvpn
  auto-discovery default
  c-mcast-signaling bgp
  provider-tunnel
      inclusive
      rsvp
```

## Provider Common Configuration

```
                lsp-template vrf2
                no shutdown
            exit
        exit
    exit
vrf-target unicast
exit
exit
```

The status of VPRN 2 on PE-1 is shown with the following output:

```
*A:PE-1>config>service>vprn# show router 2 mvpn
=====
MVPN 2 configuration data
=====
signaling          : Bgp          auto-discovery    : Default
UMH Selection     : Highest-Ip   intersite-shared  : Enabled
vrf-import       : N/A
vrf-export       : N/A
vrf-target       : unicast
C-Mcast Import RT : target:192.0.2.1:3

ipmsi             : rsvp vrf2
i-pmsi P2MP AdmSt : Up

s-pmsi           : none
data-delay-interval: 3 seconds
enable-asm-mdt   : N/A
=====
```

The following shows a debug of an Intra-AD BGP update message received by PE-1 that was sent by PE-4. The message contains the PMSI tunnel-type to be used (RSVP P2MP LSP), the P2MP LSP ID (encoded as extended Tunnel ID and P2MP-ID carried in the RSVP Session object), and the type of BGP update (Type: Intra-AD Len: 12 RD: 64496:204 Orig: 192.0.2.4):

```
41 2011/10/07 00:49:09.46 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 86
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 8 Len: 4 Community:
    no-export
  Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:64496:200
  Flag: 0xc0 Type: 22 Len: 17 PMSI:
    Tunnel-type RSVP-TE P2MP LSP (1)
    Flags [Leaf not required]
    MPLS Label 0
    P2MP-ID 0x2, Tunnel-ID: 61441, Extended-Tunnel-ID 192.0.2.4
  Flag: 0x90 Type: 14 Len: 23 Multiprotocol Reachable NLRI:
    Address Family MVPN_IPV4
    NextHop len 4 NextHop 192.0.2.4
```



```
Type: Intra-AD Len: 12 RD: 64496:204 Orig: 192.0.2.4
"
```

The set up has four PEs, so every PE should see the others peer Intra-AD route; the output below shows the routes received in PE-1:

```
*A:PE-1# show router bgp routes mvpn-ipv4 type intra-ad
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      LocalPref  MED
      RD           SourceAS          VPNLabel
      Nexthop      SourceIP
      As-Path      GroupIP
-----
u*>i  Intra-Ad         192.0.2.2         100        0
      64496:202     -                 -
      192.0.2.2     -                 -
      No As-Path    -
u*>i  Intra-Ad         192.0.2.3         100        0
      64496:203     -                 -
      192.0.2.3     -                 -
      No As-Path    -
u*>i  Intra-Ad         192.0.2.4         100        0
      64496:204     -                 -
      192.0.2.4     -                 -
      No As-Path    -
-----
Routes : 3
=====
```

The detailed output of the Intra-AD received from PE-4 shows the Tunnel-Type RSVP-TE P2MP LSP (P2MP-ID is 2), the originator id (192.0.2.4), and the route-distinguisher (64496:204):

```
*A:PE-1# show router bgp routes mvpn-ipv4 type intra-ad originator-ip 192.0.2.4 detail
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP MVPN-IPv4 Routes
=====
Route Type      : Intra-Ad
Route Dist.     : 64496:204
Originator IP   : 192.0.2.4
Nexthop        : 192.0.2.4
From           : 192.0.2.4
Res. Nexthop    : 0.0.0.0
```

## Provider Common Configuration

```

Local Pref.      : 100                      Interface Name : NotAvailable
Aggregator AS   : None                      Aggregator    : None
Atomic Aggr.    : Not Atomic                MED           : 0
Community       : no-export target:64496:200
Cluster         : No Cluster Members
Originator Id   : None                      Peer Router Id : 192.0.2.4
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : No As-Path
VPRN Imported   : 2

```

```

-----
PMSI Tunnel Attribute :
Tunnel-type          : RSVP-TE P2MP LSP      Flags           : Leaf not required
MPLS Label           : 0
P2MP-ID              : 2                    Tunnel-ID        : 61441
Extended-Tunne*     : 192.0.2.4
-----

```

For the I-PMSI, the headend PE firstly discovers all the leaf PEs via I-PMSI A-D routes, it then signals the P2MP LSP to all the leaf PEs using RSVP-TE (subsequently adding or removing S2L (source to leaf) paths as PEs are added or removed from the MVPN). Note that RSVP-TE S-PMSIs are not supported in the tested release 9.0R5.

As in the mLDP case, the demarcation of the domains are in the PE. The PE router participates in both the customer multicast domain and the provider's multicast domain. The customer's CEs are limited to a multicast adjacency with the multicast instance on the PE created to support that specific customer's IP-VPN. This way, customers are isolated from the provider's core multicast domain and other customer multicast domains while the provider's core P routers only participate in the provider's multicast domain and are isolated from all customers's multicast domains. C-trees to P-tunnels bindings are also discovered using BGP routes, instead of PIM join TLVs. MVPN c-multicast routing information is exchanged between PEs by using c-multicast routes that are carried using MCAST-VPN NLRIs.

Once the RSVP-TE P2MP LSPs are created, the I-PMSI is instantiated in the core:

```

*A:PE-1# show router 2 pim neighbor
=====
PIM Neighbor ipv4
=====
Interface          Nbr DR Prty    Up Time      Expiry Time   Hold Time
  Nbr Address
-----
int-PE-1-CE-1      1              0d 00:42:20  0d 00:01:34   105
  172.16.12.2
mpls-if-73741      1              0d 00:42:20  never          65535
  192.0.2.2
mpls-if-73742      1              0d 00:42:20  never          65535
  192.0.2.3
mpls-if-73743      1              0d 00:42:20  never          65535
  192.0.2.4
-----
Neighbors : 4
=====

```

```
*A:PE-1# show router 2 pim tunnel-interface
=====
PIM Interfaces ipv4
=====
Interface          Adm  Opr  DR Prty      Hello Intvl  Mcast Send
-----
DR
-----
mpls-if-73740      Up   Up   N/A         N/A         N/A
  192.0.2.1
mpls-if-73741      Up   Up   N/A         N/A         N/A
  192.0.2.2
mpls-if-73742      Up   Up   N/A         N/A         N/A
  192.0.2.3
mpls-if-73743      Up   Up   N/A         N/A         N/A
  192.0.2.4
-----
Interfaces : 4
```

The following command displays the PMSIs created on a PE, taking PE-3 as an example:

```
*A:PE-3# tools dump router 2 mvpn provider-tunnels
=====
MVPN 2 Inclusive Provider Tunnels Originating
=====
ipmsi (RSVP)          P2MP-ID Tunl-ID Ext-Tunl-ID
-----
vrf2-2-73732          2          61440    192.0.2.3
=====
MVPN 2 Selective Provider Tunnels Originating
=====
spmsi (RSVP)          P2MP-ID Tunl-ID Ext-Tunl-ID
-----
No Tunnels Found
=====
MVPN 2 Inclusive Provider Tunnels Terminating
=====
ipmsi (RSVP)          P2MP-ID Tunl-ID Ext-Tunl-ID
-----
mpls-if-73737          2          61452    192.0.2.1
mpls-if-73734          2          61440    192.0.2.2
mpls-if-73735          2          61441    192.0.2.4
=====
MVPN 2 Selective Provider Tunnels Terminating
=====
spmsi (RSVP)          P2MP-ID Tunl-ID Ext-Tunl-ID
-----
No Tunnels Found
=====
```

## Provider Common Configuration

Every PE has created an I-PMSI to the other PEs. As an example, PE-1 has established an LSP with name vrf2-2-73740 and PE-2, PE-3 and PE-4 as leaves. Note that the S2L path is empty as the template did not have any S2L path configured for simplicity.

```
*A:PE-1# show router mpls p2mp-lsp detail
=====
MPLS P2MP LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : vrf2-2-73740
LSP Type       : P2mpAutoLsp           LSP Tunnel ID : 61452
From           : 192.0.2.1
Adm State      : Up                    Oper State     : Up
LSP Up Time    : 0d 00:48:03           LSP Down Time : 0d 00:00:00
Transitions    : 1                    Path Changes   : 2
Retry Limit    : 0                    Retry Timer    : 30 sec
Signaling      : RSVP                 Resv. Style    : SE
Hop Limit      : 255                  Negotiated MTU : n/a
Adaptive       : Enabled              ClassType      : 0
FastReroute    : Enabled              Oper FR        : Enabled
FR Method      : Facility             FR Hop Limit   : 16
FR Bandwidth   : 0 Mbps               FR Node Protect : Disabled
FR Object      : Enabled
CSPF           : Enabled              ADSPEC         : Disabled
Metric         : Disabled             Use TE metric  : Disabled
Include Grps   :                      Exclude Grps   :
None
Least Fill    : Disabled

Auto BW        : Disabled
LdpOverRsvp    : Disabled             VprnAutoBind  : Disabled
IGP Shortcut   : Disabled             BGP Shortcut   : Disabled
BGPTransTun   : Disabled
Oper Metric    : Disabled
Prop Adm Grp   : Disabled             CSPFFirstLoose : Disabled

P2MPInstance: 2                       P2MP-Inst-type : Primary
S2l-Name       : empty                To              : 192.0.2.2
S2l-Name       : empty                To              : 192.0.2.3
S2l-Name       : empty                To              : 192.0.2.4
=====
```

Checking the RSVP-TE P2MP LSPs that are originated, transit, or destination to PE-1, the show command allows filtering by type, in this case showing the originated LSPs only:

```
*A:PE-1# show router mpls p2mp-info type originate
=====
MPLS P2MP LSPs (Originate)
=====
-----
S2L vrf2-2-73740::empty
-----
Source IP Address : 192.0.2.1           Tunnel ID      : 61452
P2MP ID           : 2                  Lsp ID        : 38402
```

```

S2L Name          : vrf2-2-73740::empty   To          : 192.0.2.2
Out Interface     : 1/1/2:1              Out Label   : 262138
Num. of S2ls     : 2
-----
S2L vrf2-2-73740::empty
-----
Source IP Address : 192.0.2.1              Tunnel ID   : 61452
P2MP ID          : 2                    Lsp ID     : 38402
S2L Name         : vrf2-2-73740::empty   To          : 192.0.2.3
Out Interface     : 1/1/3:1              Out Label   : 262128
Num. of S2ls     : 1
-----
S2L vrf2-2-73740::empty
-----
Source IP Address : 192.0.2.1              Tunnel ID   : 61452
P2MP ID          : 2                    Lsp ID     : 38402
S2L Name         : vrf2-2-73740::empty   To          : 192.0.2.4
Out Interface     : 1/1/2:1              Out Label   : 262138
Num. of S2ls     : 2
-----
P2MP Cross-connect instances : 3
=====

```

Following the path of the S2L from PE-1 to PE-4 (third entry S2L vrf2-2-73740), the outgoing interface is 1/1/2 that connects PE-1 to PE-2, so the LSP goes to PE-4 via PE-2.

```

*A:PE-2# show router mpls p2mp-info type transit
=====
MPLS P2MP LSPs (Transit)
=====
-----
S2L vrf2-2-73740::empty
-----
Source IP Address : 192.0.2.1              Tunnel ID   : 61452
P2MP ID          : 2                    Lsp ID     : 38402
S2L Name         : vrf2-2-73740::empty   To          : 192.0.2.4
Out Interface     : 1/1/3:1              Out Label   : 262136
Num. of S2ls     : 1
-----
S2L vrf2-2-73732::empty
-----
Source IP Address : 192.0.2.4              Tunnel ID   : 61441
P2MP ID          : 2                    Lsp ID     : 17924
S2L Name         : vrf2-2-73732::empty   To          : 192.0.2.1
Out Interface     : 1/1/2:1              Out Label   : 262133
Num. of S2ls     : 1
-----
P2MP Cross-connect instances : 2
=====

```

As transit, PE-2 shows that there is an LSP coming from PE-1 (vrf2-2-73740) and the outgoing interface is 1/1/3 that connects PE-2 with PE-4.

## Provider Common Configuration

Using the same command with a different filter on PE-4, 3 P2MP LSPs are terminated, one from each remote PE (PE-1, PE-2 and PE-3). On PE-4, an S2L vrf2-2-73740 from 192.0.2.1 and P2MP ID = 2 is traced.

```
*A:PE-4# show router mpls p2mp-info type terminate
=====
MPLS P2MP LSPs (Terminate)
=====
-----
S2L vrf2-2-73740::empty
-----
Source IP Address   : 192.0.2.1           Tunnel ID   : 61452
P2MP ID            : 2                 Lsp ID     : 38402
S2L Name           : vrf2-2-73740::empty To         : 192.0.2.4
In Interface       : 1/1/3:1          In Label   : 262136
Num. of S2ls      : 1
-----
S2L vrf2-2-73732::empty
-----
Source IP Address   : 192.0.2.2           Tunnel ID   : 61440
P2MP ID            : 2                 Lsp ID     : 8196
S2L Name           : vrf2-2-73732::empty To         : 192.0.2.4
In Interface       : 1/1/3:1          In Label   : 262132
Num. of S2ls      : 1
-----
S2L vrf2-2-73732::empty
-----
Source IP Address   : 192.0.2.3           Tunnel ID   : 61440
P2MP ID            : 2                 Lsp ID     : 59396
S2L Name           : vrf2-2-73732::empty To         : 192.0.2.4
In Interface       : 1/1/2:1          In Label   : 262142
Num. of S2ls      : 2
-----
P2MP Cross-connect instances : 3
=====
```

The following output shows P2MP LSP on PE-1 with more detail:

```
*A:PE-1# show router mpls p2mp-lsp "vrf2-2-73740" p2mp-instance "2" s2l "empty" detail
=====
MPLS LSP vrf2-2-73740 S2L empty (Detail)
=====
Legend :
  @ - Detour Available           # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  S - Strict                     L - Loose
  s - Soft Preemption
=====
LSP vrf2-2-73740 S2L empty
-----
LSP Name       : vrf2-2-73740           S2L LSP ID  : 38402
P2MP ID        : 2                     S2L Grp Id  : 1
Adm State      : Up                    Oper State   : Up
S2L State      : Active                  :
S2L Name       : empty                  To          : 192.0.2.2
S2L Admin      : Up                    S2L Oper    : Up
```

```

OutInterface: 1/1/2:1
S2L Up Time : 0d 01:22:50
RetryAttempt: 0
S2L Trans : 2
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  192.168.12.1(192.0.2.1) @
  -> 192.168.12.2(192.0.2.2)
ComputedHops:
  192.168.12.1(S) -> 192.168.12.2(S)
LastResignal: n/a

```

-----  
LSP vrf2-2-73740 S2L empty  
-----

```

LSP Name      : vrf2-2-73740
P2MP ID      : 2
Adm State    : Up
S2l State:   : Active
S2L Name     : empty
S2l Admin    : Up
OutInterface: 1/1/3:1
S2L Up Time  : 0d 01:22:50
RetryAttempt: 0
S2L Trans   : 2
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  192.168.13.1(192.0.2.1) @
  -> 192.168.13.2(192.0.2.3)
ComputedHops:
  192.168.13.1(S) -> 192.168.13.2(S)
LastResignal: n/a

```

-----  
LSP vrf2-2-73740 S2L empty  
-----

```

LSP Name      : vrf2-2-73740
P2MP ID      : 2
Adm State    : Up
S2l State:   : Active
S2L Name     : empty
S2l Admin    : Up
OutInterface: 1/1/2:1
S2L Up Time  : 0d 01:22:54
RetryAttempt: 0
S2L Trans   : 3
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  192.168.12.1(192.0.2.1) @
  -> 192.168.12.2(192.0.2.2) @
  -> 192.168.24.2(192.0.2.4)
ComputedHops:
  192.168.12.1(S) -> 192.168.12.2(S) -> 192.168.24.2(S)
LastResignal: n/a

```

## Provider Common Configuration

The last entry, vrf2-2-73740, provides the details of the S2L traced earlier, displaying the different hops (PE-1, PE-2, and PE-3), the fast reroute protection (link protection is supported only) and the labels used (262138 from PE-1 to PE-2, 262136 from PE-2 to PE-4). Note that on PE-1, although only one has been shown, both links PE-1 to PE-3 and PE-1 to PE-2 are fast reroute protected.

If any of the protected links between PE-1 and PE-2 or PE-3 are broken, fast reroute will be initiated. The protected bypass hops are displayed with the following command:

```
*A:PE-1# show router mpls bypass-tunnel protected-lsp p2mp detail
=====
MPLS Bypass Tunnels (Detail)
=====
-----
bypass-link192.168.12.2
-----
To                : 192.168.24.1          State           : Up
Out I/F           : 1/1/3:1                Out Label      : 262132
Up Time           : 0d 00:08:19         Active Time    : n/a
Reserved BW       : 0 Kbps              Protected LSP Count : 3
Type              : P2mp
Setup Priority     : 7                   Hold Priority   : 0
Class Type        : 0
Actual Hops       :
    192.168.13.1(S)  -> 192.168.13.2(S)  -> 192.168.34.2(S)
    -> 192.168.24.1(S)

Protected LSPs -
LSP Name          : vrf2-2-73740::empty
From              : 192.0.2.1                To              : 192.0.2.2
Avoid Node/Hop    : 192.168.12.2            Downstream Label : 262138
Bandwidth         : 0 Kbps

LSP Name          : vrf2-2-73740::empty
From              : 192.0.2.1                To              : 192.0.2.4
Avoid Node/Hop    : 192.168.12.2            Downstream Label : 262138
Bandwidth         : 0 Kbps

LSP Name          : vrf2-2-73732::empty
From              : 192.0.2.3                To              : 192.0.2.2
Avoid Node/Hop    : 192.168.12.2            Downstream Label : 262137
Bandwidth         : 0 Kbps
-----
bypass-link192.168.13.2
-----
To                : 192.168.34.1          State           : Up
Out I/F           : 1/1/2:1                Out Label      : 262131
Up Time           : 0d 00:08:10         Active Time    : n/a
Reserved BW       : 0 Kbps              Protected LSP Count : 2
Type              : P2mp
Setup Priority     : 7                   Hold Priority   : 0
Class Type        : 0
Actual Hops       :
    192.168.12.1(S)  -> 192.168.12.2(S)  -> 192.168.24.2(S)
    -> 192.168.34.1(S)

Protected LSPs -
LSP Name          : vrf2-2-73740::empty
```



From : 192.0.2.1 To : 192.0.2.3  
Avoid Node/Hop : 192.168.13.2 Downstream Label : 262128  
Bandwidth : 0 Kbps

LSP Name : vrf2-2-73732::empty  
From : 192.0.2.2 To : 192.0.2.3  
Avoid Node/Hop : 192.168.13.2 Downstream Label : 262127  
Bandwidth : 0 Kbps

=====

## Provider Common Configuration

### Traffic Flow

The receiver RX-4, connected to CE-4, wishes to join the group 232.2.2.2 with source 192.168.2.1 and so sends an IGMPv3 report towards CE-4. CE-4 recognizes the report and sends a PIM join towards the source, hence it reaches PE-1 where the source is connected to through CE-1. The output below shows the debug seen on PE-4, where the PIM join is received from CE-4 and a BGP update Source Join is sent to all PEs (note that only the update sent to PE-1 is shown).

```
40 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPProcessSG
pimJPPProcessSG: (S,G)-> (192.168.2.1,232.2.2.2) type <S,G>, i/f int-PE-4-CE-4, u
pNbr 172.16.42.1 isJoin 1 isRpt 0 holdTime 210"

41 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Downstream: State NoInfo Event RxJoin, (S,G) (192.168.2.1,232.2.2.2) grou
pType <S,G>"

42 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Upstream: State NotJoined Event JoinDesiredTrue, (S,G) (192.168.2.1,232.2
.2.2) groupType <S,G>"

43 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpJoinDesiredTrue
No upstream interface. pSG (192.168.2.1,232.2.2.2) rpfType 3"

44 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpJoinDesiredTrue
pim 3 sg_type 2 refetch route type SPMSI pendingFetchMask 0x8"

45 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpJoinDesiredTrue
No upstream interface. pSG 0x5578f87c, (192.168.2.1,232.2.2.2) rpfType 3"

46 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Upstream: State Joined Event MribChange, (S,G) (192.168.2.1,232.2.2.2) gr
oupType <S,G>"

47 2011/10/13 19:02:20.42 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpStateJMribChange
pSG 0x5578f87c, (192.168.2.1,232.2.2.2), type <S,G> oldMribNhopIp 0.0.0.0 oldRpf
NbrIp 0.0.0.0, oldRpfType NONE oldRpfIf 0 rptMribNhopIp 0.0.0.0, rptRpfNbrIp 0.0
.0.0 rtmReason 32"

48 2011/10/13 19:02:20.41 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpStateJMribChange
pSG 0x5578f87c, (192.168.2.1,232.2.2.2), type <S,G> newMribNhopIp 192.0.2.1 newR
pfNbrIp 192.0.2.1 newRpfType REMOTE newRpfIf 73737"

49 2011/10/13 19:02:20.41 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 69
  Flag: 0x40 Type: 1 Len: 1 Origin: 0"
```

```

Flag: 0x40 Type: 2 Len: 0 AS Path:
Flag: 0x80 Type: 4 Len: 4 MED: 0
Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:192.0.2.1:3
Flag: 0x90 Type: 14 Len: 33 Multiprotocol Reachable NLRI:
    Address Family MVPN_IPV4
    NextHop len 4 NextHop 192.0.2.4
    Type: Source-Join Len:22 RD: 64496:201 SrcAS: 64496 Src: 192.168.2.1 Grp
: 232.2.2.2
"

```

The following debug shows that PE-1 receives the BGP update Source Join with source 192.168.2.1 and group 232.2.2.2 and sends a PIM join towards CE-1:

```

67 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 69
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:192.0.2.1:3
    Flag: 0x90 Type: 14 Len: 33 Multiprotocol Reachable NLRI:
        Address Family MVPN_IPV4
        NextHop len 4 NextHop 192.0.2.4
        Type: Source-Join Len:22 RD: 64496:201 SrcAS: 64496 Src: 192.168.2.1 Grp
: 232.2.2.2
"

68 2011/10/14 05:11:35.05 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPProcessSG
pimJPPProcessSG: (S,G)-> (192.168.2.1,232.2.2.2) type <S,G>, i/f mpls-if-73740, u
pNbr 192.0.2.1 isJoin 1 isRpt 0 holdTime 65535"

69 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Downstream: State NoInfo Event RxJoin, (S,G) (192.168.2.1,232.2.2.2) grou
pType <S,G>"

70 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Upstream: State NotJoined Event JoinDesiredTrue, (S,G) (192.168.2.1,232.2
.2.2) groupType <S,G>"

71 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpJoinDesiredTrue
pim 3 sg_type 2 refetch route type SPMSI pendingFetchMask 0x8"

72 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSendJoinPrunePdu
pimSendJoinPrunePdu: if 3, adj 172.16.12.2"

73 2011/10/14 05:11:35.06 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]

```

## Provider Common Configuration

```
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: encoding groupset for group 232.2.2.2, numJoinedSrcs 1, numPrunedSrcs 0"

74 2011/10/14 05:11:35.05 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: Encoding Join for source 192.168.2.1"

75 2011/10/14 05:11:35.05 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: num joined srcs 1, num pruned srcs 0"

76 2011/10/14 05:11:35.05 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSendJoinPrunePdu
pimSendJoinPrunePdu2: sending JP PDU with 1 groups."

77 2011/10/14 05:12:34.76 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Upstream: State Joined Event JTimerExp, (S,G) (192.168.2.1,232.2.2.2) groupType <S,G>"

78 2011/10/14 05:12:34.76 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSendJoinPrunePdu
pimSendJoinPrunePdu: if 3, adj 172.16.12.2"

79 2011/10/14 05:12:34.76 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: encoding groupset for group 232.2.2.2, numJoinedSrcs 1, numPrunedSrcs 0"

80 2011/10/14 05:12:34.76 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: Encoding Join for source 192.168.2.1"

81 2011/10/14 05:12:34.77 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGEncodeGroupSet
pimEncodeGroupSet: num joined srcs 1, num pruned srcs 0"

82 2011/10/14 05:12:34.77 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSendJoinPrunePdu
pimSendJoinPrunePdu2: sending JP PDU with 1 groups."
```

The BGP update source join received by PE-1 is displayed with the following command:

```
*A:PE-1# show router bgp routes mvpn-ipv4 type source-join
=====
BGP Router ID:192.0.2.1      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

=====
BGP MVPN-IPv4 Routes
=====
Flag  RouteType      OriginatorIP      LocalPref  MED
      RD            SourceAS          VPNLabel
      Nexthop       SourceIP
      As-Path       GroupIP
-----
u*>i  Source-Join      -                 100        0
      64496:201      64496            -
      192.0.2.4     192.168.2.1
      No As-Path    232.2.2.2
-----
Routes : 1
=====
```

To verify the traffic: on PE-1 there is a group 232.2.2.2 with source 192.168.2.1, the RPF is CE-1, and the multicast traffic is flowing from CE-1 to PE-1 using int-PE-1-CE-1 and the outgoing interface is using the PMSI RSVP mpls-if-73740.

```
*A:PE-1# show router 2 pim group detail
=====
PIM Source Group ipv4
=====
Group Address      : 232.2.2.2
Source Address     : 192.168.2.1
RP Address         : 0
Flags              :
Type               : (S,G)
MRIB Next Hop     : 172.16.12.2
MRIB Src Flags    : remote
Keepalive Timer   : Not Running
Up Time           : 0d 00:10:48
Resolved By       : rtable-u

Up JP State       : Joined
Up JP Expiry      : 0d 00:00:12
Up JP Rpt         : Not Joined StarG
Up JP Rpt Override : 0d 00:00:00

Register State    : No Info
Reg From Anycast RP: No

Rpf Neighbor      : 172.16.12.2
Incoming Intf     : int-PE-1-CE-1
Outgoing Intf List : mpls-if-73740

Curr Fwding Rate  : 0.0 kbps
Forwarded Packets : 0
Discarded Packets : 0
Forwarded Octets  : 0
RPF Mismatches    : 0
Spt threshold     : 0 kbps
ECMP opt threshold : 7
```

## Provider Common Configuration

```
Admin bandwidth      : 1 kbps
```

```
-----  
Groups : 1  
=====
```

On PE-4, the same (S,G) arrives in the incoming interface mpls-if-73737 and the outgoing interface is int-PE-4-CE-4.

```
*A:PE-4# show router 2 pim group detail
```

```
=====
```

```
PIM Source Group ipv4  
=====
```

Group Address	: 232.2.2.2		
Source Address	: 192.168.2.1		
RP Address	: 0		
Flags	:	Type	: (S,G)
MRIB Next Hop	: 192.0.2.1		
MRIB Src Flags	: remote	Keepalive Timer	: Not Running
Up Time	: 0d 00:12:40	Resolved By	: rtable-u
Up JP State	: Joined	Up JP Expiry	: 0d 00:00:19
Up JP Rpt	: Not Joined StarG	Up JP Rpt Override	: 0d 00:00:00

```
Register State      : No Info  
Reg From Anycast RP: No
```

```
Rpf Neighbor        : 192.0.2.1  
Incoming Intf       : mpls-if-73737  
Outgoing Intf List : int-PE-4-CE-4
```

```
Curr Fwding Rate    : 0.0 kbps  
Forwarded Packets   : 0  
Forwarded Octets    : 0  
Spt threshold       : 0 kbps  
Admin bandwidth     : 1 kbps  
Discarded Packets   : 0  
RPF Mismatches      : 0  
ECMP opt threshold  : 7
```

```
-----  
Groups : 1  
=====
```

When the receiver is not interested in the channel group any more, the receiver RX-4 sends an IGMPv3 leave, PE-4 sends a PIM prune translated to a BGP MP\_UNREACH NLRI to all PEs. Note that, as mentioned before, rapid withdrawals are sent without waiting for the mrai (note that for simplicity, only one BGP update is shown in the output debug).

```
92 2011/10/13 19:16:36.67 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]  
"PIM[Instance 3 vprn2]: pimJPProcessSG  
pimJPProcessSG: (S,G)-> (192.168.2.1,232.2.2.2) type <S,G>, i/f int-PE-4-CE-4, u  
pNbr 172.16.42.1 isJoin 0 isRpt 0 holdTime 210"
```

```
93 2011/10/13 19:16:36.67 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]  
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent  
PIM JP Downstream: State Joined Event RxPrune, (S,G) (192.168.2.1,232.2.2.2) gro  
upType <S,G>"
```

```
94 2011/10/13 19:16:36.67 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
```

```

"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Downstream: State PrunePending Event PrunePendTimerExp, (S,G) (192.168.2.
1,232.2.2.2) groupType <S,G>"

95 2011/10/13 19:16:36.67 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimJPPrintFsmEvent
PIM JP Upstream: State Joined Event JoinDesiredFalse, (S,G) (192.168.2.1,232.2.2
.2) groupType <S,G>"

96 2011/10/13 19:16:36.67 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 31
  Flag: 0x90 Type: 15 Len: 27 Multiprotocol Unreachable NLRI:
    Address Family MVPN_IPV4
    Type: Source-Join Len:22 RD: 64496:201 SrcAS: 64496 Src: 192.168.2.1 Grp
: 232.2.2.2
"

```

### MVPN Source Redundancy

So far, the multicast traffic has been streamed towards router CE-1 from a single source. For redundancy purposes, the source can be redundant (two sources attached to different CEs that connect to a pair of PEs). To simulate the redundancy, CE-1 has been connected to both PE-1 and PE-2, using VPRN 2, and ECMP (Equal Cost Multi Path) is configured with the value of 2 in all PEs. With this configuration, any PE is able to reach the source through PE-1 and PE-2. The (S,G) is the same as the one used in P2MP RSVP TE (192.168.2.1, 232.2.2.2). [Figure 98](#) shows the VPRN 2 topology with the source redundancy.

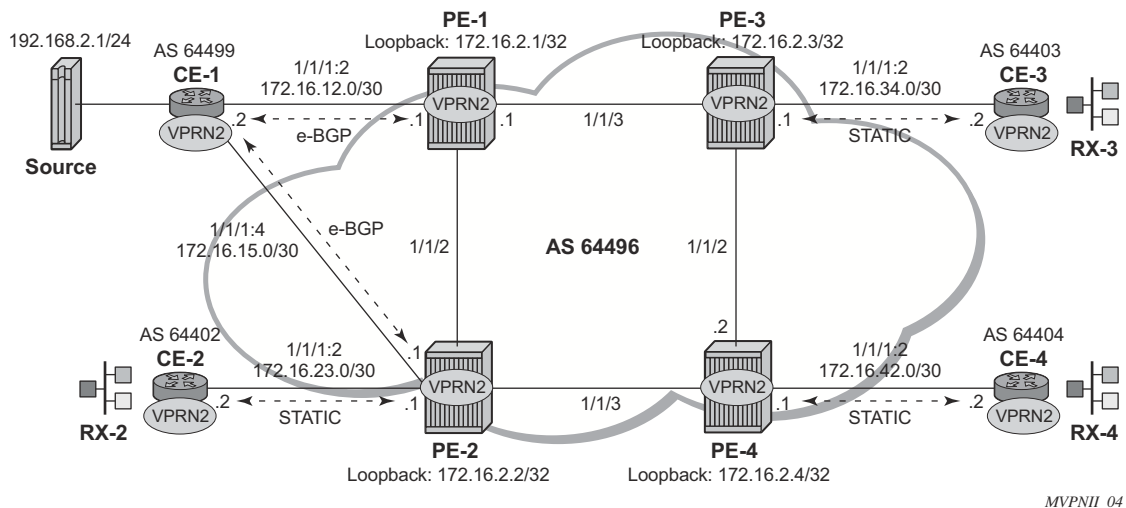


Figure 98: VPRN 2 Topology used for MVPN Source Redundancy

## Provider Common Configuration

The configuration change with respect to the previous section (P2MP RSVP-TE PMSIs) is an additional interface created in both CE-1 and PE-2 (int-CE-1-PE-2 on CE-1 and int-PE-2-CE-1 on PE-2), the addition of these interfaces to pim and also the creation an e-BGP session between the two routers. The following is the configuration on PE-2 (CE-1 configuration changes are not displayed for brevity).

```
*A:PE-2>config>service>vprn# info
-----
description "P2MP RSVP"
ecmp 2
autonomous-system 64496
route-distinguisher 64496:202
vrf-target target:64496:200
interface "loopback" create
    address 172.16.2.2/32
    loopback
exit
interface "int-PE-2-CE-2" create
    address 172.16.23.1/30
    sap 1/1/1:2 create
    exit
exit
interface "int-PE-2-CE-1" create
    address 172.16.16.1/30
    sap 1/1/4:2 create
    exit
exit
bgp
    group "external"
        type external
        peer-as 64499
        neighbor 172.16.16.2
        exit
    exit
    no shutdown
exit
pim
    interface "loopback"
    exit
    interface "int-PE-2-CE-2"
    exit
    interface "int-PE-2-CE-1"
    exit
    rp
        static
        exit
        bsr-candidate
            shutdown
        exit
        rp-candidate
            shutdown
        exit
    exit
    no shutdown
exit
mvpn
    auto-discovery default
    c-mcast-signaling bgp
```



```

provider-tunnel
  inclusive
  rsvp
    lsp-template vrf2
    no shutdown
  exit
exit
exit
vrf-target unicast
exit
exit
spoke-sdp 12 create
  no shutdown
exit
spoke-sdp 23 create
  no shutdown
exit
spoke-sdp 24 create
  no shutdown
exit
no shutdown

```

-----

Checking the routes on PE-4, the source is reachable through PE-1 and PE-2 as ECMP is set to 2. Note that if the configuration of the VPRN is provisioned with **autobind-mpls** instead of static spoke-SDPs, the command **ignore-nh-metric** is also needed.

```

*A:PE-4# show router 2 route-table
=====
Route Table (Service: 2)
=====
Dest Prefix[Flags]                               Type   Proto   Age      Pref
  Next Hop[Interface Name]                       Metric
-----
* Truncated info

192.168.2.1/32                                   Remote BGP VPN 00h19m55s 170
    192.0.2.1 (tunneled)                          0
192.168.2.1/32                                   Remote BGP VPN 00h19m55s 170
    192.0.2.2 (tunneled)                          0
-----
No. of Routes: 11
Flags: L = LFA nexthop available   B = BGP backup route available
      n = Number of times nexthop is repeated
=====

```

When PE-4 receives a c-join/prune, PE-4 needs to find the **upstream multicast hop** for the (S,G). This is the Upstream Multihop Selection (UMH) and is configurable. The values are highest-ip, hash-based and tunnel-status.

```

*A:PE-4>config>service>vprn>mvpn# umh-selection
- no umh-selection
- umh-selection {highest-ip|hash-based|tunnel-status}

```

## Provider Common Configuration

The default is highest-ip, which is the selection of the highest /32 IP addresses (in this set up, PE-2 is preferred versus PE-1). A BGP c-join is sent with the route target equal to the VRF import extended community distributed by PE-2 for the subnet of the source (see PE-4 debug show below).

```
* Truncated info
12 2011/10/14 08:42:53.87 UTC MINOR: DEBUG #2001 vprn2 PIM[Instance 3 vprn2]
"PIM[Instance 3 vprn2]: pimSGUpStateJMribChange
pSG 0x5578f87c, (192.168.2.1,232.2.2.2), type <S,G> newMribNhopIp 192.0.2.2 newR
pfNbrIp 192.0.2.2 newRpfType REMOTE newRpfIf 73741"

13 2011/10/14 08:42:53.87 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 69
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:192.0.2.2:3
  Flag: 0x90 Type: 14 Len: 33 Multiprotocol Reachable NLRI:
    Address Family MVPN_IPV4
    NextHop len 4 NextHop 192.0.2.4
    Type: Source-Join Len:22 RD: 64496:202 SrcAS: 64496 Src: 192.168.2.1 Grp
: 232.2.2.2
"
```

The second option is hash-based, where the UMH is selected (both PEs are potentially possible UMHs) after hashing the stream's source and group addresses. For this example, PE-2 is also preferred.

The third option, tunnel-status, is based on the status of the P2MP RSVP tunnel (not available in mLDP or PIM). The roots PE-1 and PE-2 are sending BFD messages to the leaf PE-4 (in fact this is UFD, unidirectional forwarding detection). PE-4's c-join for the (S,G) is sent to both PE-1 and PE-2, and in return the traffic is forwarded from both PE-1 and PE-2 for the c-group onto the I-PMSI; hence PE-4 receives two copies of the c-(S,G) stream. By configuration, the stream from the primary PE-1 is selected by PE-4 to be forwarded to receiver RX-4. If BFD messages are no longer received over the primary P2MP LSP, then stream from the standby PE-2 is selected and forwarded to the receiver.

The configuration on PE-1 and PE-2 is similar and is as follows (only PE-2 is shown):

```
PE-2>configure service vprn 2
      mvpn
        auto-discovery default
        c-mcast-signaling bgp
        umh-selection tunnel-status
        provider-tunnel
          inclusive
          rsvp
            lsp-template vrf2
            enable-bfd-root 100
```

```

                no shutdown
            exit
        exit
    exit
    vrf-target unicast
    exit
exit

```

PE-1 and PE-2 are root. On PE-4 BFD is configured as leaf and the primary PE (PE-1) and backup PE (PE-2) are also provisioned:

```

PE-4>configure service vprn 2
    mvpn
        auto-discovery default
        c-mcast-signaling bgp
        umh-selection tunnel-status
        umh-pe-backup
            umh-pe 192.0.2.1 standby 192.0.2.2
        exit
        provider-tunnel
            inclusive
            rsvp
                lsp-template vrf2
                enable-bfd-leaf
                no shutdown
            exit
        exit
    exit
    vrf-target unicast
    exit
exit

```

This BFD (UFD) configuration on the root establishes a session with the leaf where BFD packets are only transmitted:

```

*A:PE-1>config>service>vprn# show router 2 bfd session
=====
BFD Session
=====
Interface          State          Tx Intvl  Rx Intvl  Multipl
 Remote Address    Protocols      Tx Pkts   Rx Pkts   Type
-----
mpls-if-73747     Up (3)         100        0          3
 127.0.0.0        pim            9955       0          central
-----
No. of BFD sessions: 1

```

On PE-4, two BFD sessions are received, one from each root (note that BFD packets are only received):

```

*A:PE-4>config>service>vprn# show router 2 bfd session
=====
BFD Session
=====

```

## Provider Common Configuration

Interface Remote Address	State Protocols	Tx Intvl Tx Pkts	Rx Intvl Rx Pkts	Multipl Type
mpls-if-73747 192.0.2.1	Up (3) pim	0 0	100 681	3 central
mpls-if-73748 192.0.2.2	Up (3) pim	0 0	100 470	3 central

No. of BFD sessions: 2  
=====

PE-4 delivers the multicast traffic from the primary configured UMH, PE-1. If, as an example of a failure condition, PE-1 goes down (reboot), PE-4 will switch to the PE-2 P2MP LSP.

---

## MDT AFI SAFI

In the Draft-Rosen up to version 6, the default MDT is PIM Spare Mode only, and there is no auto-discovery mechanism available. From SR-OS Release 7.0 and beyond, it is possible to configure PIM SSM with auto-discovery, using AFI 1 and SAFI 5. The Draft-Rosen version 7 allows use of MDT PIM SM or SSM, and auto-discovery based on AFI 1 and SAFI 66 to distribute the default MDT. Draft-Rosen version 9 adds a new MDT NLRI. The SR-OS has added the capability and support of MDT SAFI as specified in draft-nalawade-idr-mdt-safi-03 and draft-rozen-vpn-mcast-15.

MDT SAFI is used to discover PEs in a specific MVPN, so that PIM SSM can be used for default MDT. The basic idea is the same as MVPN BGP auto-discovery, but it uses a different BGP SAFI. BGP messages in which AFI=1 and SAFI=66 are "MDT-SAFI" messages. The NLRI format is 8-byte-RD:IPv4-address followed by the MDT group address. The IPv4 address identifies the PE that originated this route and the RD identifies a VRF in that PE. The group address must be an IPv4 multicast group address and is used to build the P-tunnels.

All PEs attached to a given MVPN must specify the same group-address. MDT-SAFI routes do not carry RTs and the group address is used to associate a received MDT-SAFI route with a VRF.

MDT SAFI can only be used when the implicit provider tunnel is PIM GRE based with a specific IPv4 group address.

Note: For additional information on the use of PIM PMSIs, refer to [Multicast in a VPN I on page 635](#).

[Figure 99](#) shows the topology of VPRN 3.




Figure 99: VPRN 3 Topology used for MVPN Source Redundancy

In this scenario, there is no MPLS based PMSI, there is PIM in the core for the control plane and the data traffic is GRE encapsulated. PIM needs to be configured in the base router on interface system and on the other interfaces pointing to other PEs. PIM is used for c-signaling. In addition, auto-discovery is provisioned to use mdt-safi and a PIM SSM inclusive PMSI with address 239.1.1.1 as the default MDT. The configuration, in all PEs, is like the following on PE-4:

```
*A:PE-4>config router
  pim
    interface "system"
    exit
    interface "int-PE-4-PE-2"
    exit
    interface "int-PE-4-PE-3"
    exit
  rp
    static
    exit
    bsr-candidate
    shutdown
    exit
    rp-candidate
    shutdown
    exit
  exit
  no shutdown
exit

*A:PE-4>config service vprn 3
  description "PIM SSM"
  autonomous-system 64496
```

## Provider Common Configuration

```
route-distinguisher 64496:304
vrf-target target:64496:300
interface "loopback" create
    address 172.16.3.4/32
    loopback
exit
interface "int-PE-4-CE-4" create
    address 172.16.43.1/30
    sap 1/1/1:3 create
    exit
exit
pim
    interface "loopback"
    exit
    interface "int-PE-4-CE-4"
    exit
    rp
        static
        exit
        bsr-candidate
            shutdown
        exit
        rp-candidate
            shutdown
        exit
    exit
    no shutdown
exit
mvpn
    auto-discovery mdt-safi
    provider-tunnel
        inclusive
        pim ssm 239.1.1.1
        exit
    exit
    vrf-target unicast
    exit
exit
spoke-sdp 14 create
    no shutdown
exit
spoke-sdp 24 create
    no shutdown
exit
spoke-sdp 34 create
    no shutdown
exit
no shutdown
```

The following debug output shows a BGP update with MDT AFI SAFI:

```
2 2011/10/14 16:48:14.58 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 62
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 16 Len: 8 Extended Community:
    target:64496:300
  Flag: 0x90 Type: 14 Len: 26 Multiprotocol Reachable NLRI:
    Address Family MDT-SAFI
    NextHop len 4 NextHop 192.0.2.4
    [MDT-SAFI] Addr 192.0.2.4, Group 239.1.1.1, RD 64496:304
```

## Conclusion

This chapter provides information to configure multicast within a VPRN with next generation multicast VPN techniques. Specifically, the use of MPLS I-PMSIs (mLDP and P2MP RSVP-TE), MVPN source redundancy and the complete set of features needed to interoperate with Draft-Rosen in live deployments are covered.