MVPN: Inter-AS Option B

In This Chapter

This section provides information about MVPN: Inter-AS Option B configurations.

Topics in this section include:

Applicability

This example is applicable to 7950 XRS, 7750 all variants, 7750 SR c4/12, 7450 mixed mode systems. Chassis mode C or D must be used. The configuration was tested on Release 11.0 R3.

Overview

This configuration note covers a basic technology overview, the network topology and configuration examples which are used for Multicast VPN Inter-AS option B.

Knowledge of the Alcatel-Lucent multicast and Layer 3 VPNs concepts are assumed throughout this document.

Overview

The Inter-AS MVPN feature allows the setup of Multicast Distribution Trees that span multiple Autonomous Systems.

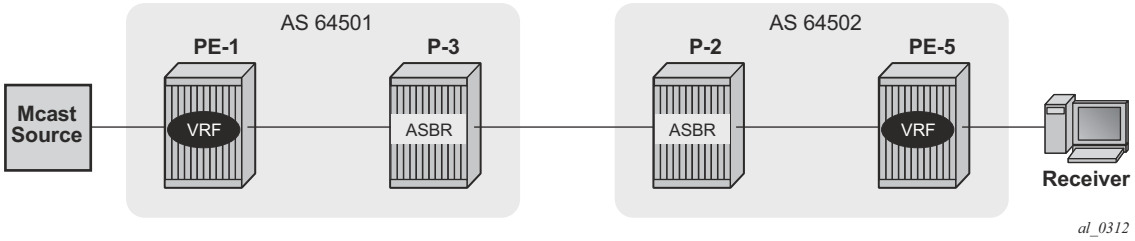

Figure 107: General Topology for Inter-AS MVPN

This example covers Draft-Rosen Inter-AS support (Option-B). Inter-AS Option B is supported for PIM SSM with Draft-Rosen MVPN using Multicast Distribution Tree (MDT) Subsequent Address Family (SAFI), using the BGP Connector attribute and PIM RPF vector.

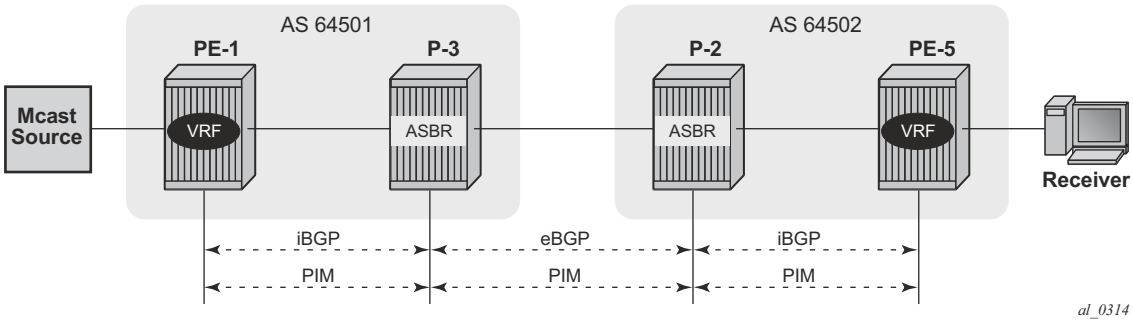

Figure 108: Protocols Used for Inter-AS MVPN

The following assumptions are made:

- PE-1 is named “sender PE” because the multicast source is directly connected to this router.
- PE-5 is named “receiver PE” because multicast receiver is directly connected to this router.

- P-2 and P-3 are named “ASBR” routers according to Inter-AS model.

The multicast receiver and source can be indirectly connected to PE routers via CE routers, but for the core multicast distribution these variations are conceptually the same. For simplicity, the PE and P router configurations will be provided.

There are several challenges which have to be solved in order to make complete inter-as solution operational:

Challenge 1:

In case of Inter-AS MVPN Option B, routing information towards the source PE is not available in a remote AS domain since IGP routes are not exchanged between ASs.

As a result a PIM-P Join would never be sent upstream (from the receiver PE to the sender PE in a different AS). However, the PIM-P join has to be propagated from PE-5 to PE-1. Therefore a solution is required to issue PIM-P Join and perform RPF.

Solution:

Use a PIM RPFV (Reverse Path Forwarding (RPF) vector) to segment the PIM-P propagation. In this example there are three segments:

- PE-5 -> ASBR P-2
- ASBR P-2 -> ASBR P-3
- ASBR P-3 -> PE-1

The RPF vector is added to a PIM join at the PE router when the following option is enabled:

```
*A:PE-5>config>router>pim# rpfv
  - no rpfv [mvpn]
  - rpfv mvpn
 <mvpn>               : Proxy RPF vector for inter-AS rosen mvpn
```

mvpn enables “mvpn RPF vector” processing for Inter-AS Option B MVPN based on RFC 5496 and RFC 6513. If a “core RPF” vector is received, it will be dropped before a message is processed.

All routers which are used for multicast traffic transportation must have this option enabled to allow RPF Vector processing. If the option is not enabled, the RPF Vector is dropped and the PIM Join is processed as if the PIM Vector is not present.

Details about RPF Vector can be found in the following RFCs: 5496, 5384, 6513.

Challenge 2:

With Inter-AS MVPN Option B, the BGP next-hop is modified by the local and remote ASBRs during re-advertisement of VPNv4 routes. When the BGP next-hop is changed, information regarding the originator of the prefix is lost when the advertisement reaches the receiver PE node. Therefore a solution is required to do a successful RPF check for the VPN source at receiver VPRN.

Note: This challenge does not apply to Model C since in Model C the BGP next-hop for VPN routes is not updated.

Solution:

A new transitive BGP attribute - Connector - is used to advertise an address of a sender PE node which is carried inside VPNv4 update. The BGP connector attribute allows the sender PE address information to be available to the receiver PE so that a receiver PE is able to associate VPNv4 advertisement to the corresponding source PE.

Inter-AS Option B will work when the following criteria are met:

- Draft-rosen MVPN is used with PIM SSM
- BGP MDT-SAFI address family is used
- PIM RPF Vector is configured
- BGP Connector attribute is used for vpn-ipv4 updates

SR OS inter-AS Option B is designed to be standard compliant based on the following RFCs:

- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*

Three global signaling stages can be identified when Inter-AS MVPN is configured:

Stage 1: BGP core signaling

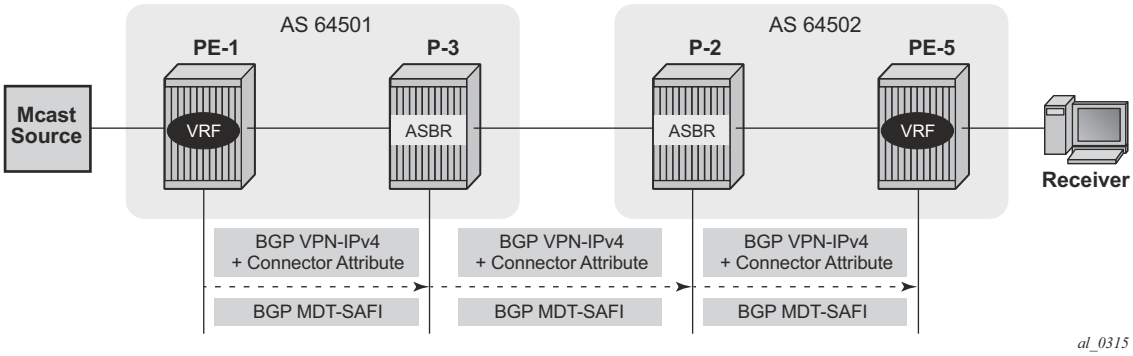

Figure 109: BGP Signaling Steps

The sender PE sends VPN-IPv4 and MDT-SAFI BGP updates for this particular MVPN:

- Every ASBR propagates VPN-IPv4 and MDT-SAFI BGP updates:
  → Next-Hop (NH) attribute is modified every time
  → Connector attribute stays untouched

When this stage is completed, all routers have necessary information:

- to start PIM signaling in the core network (PIM-P) to prepare the Default MDT
- to start PIM signaling of customer's multicast streams (PIM-C) inside VPN

Stage 2: Core PIM signaling

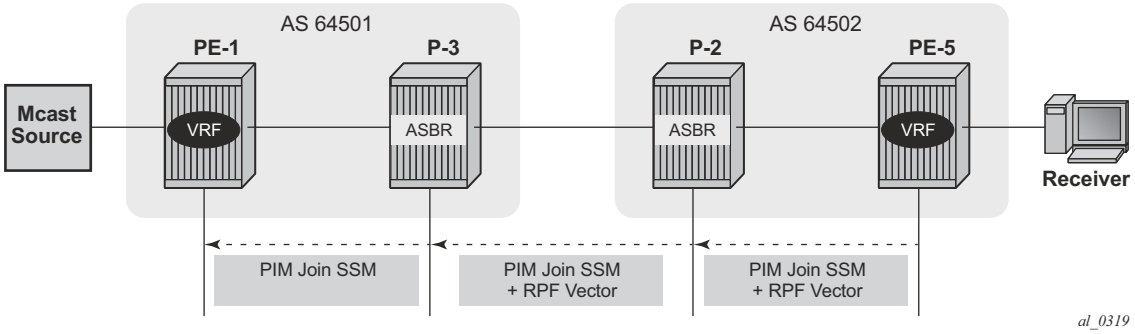

Figure 110: PIM-P Signaling Steps for Default MDT

PE-5 determines the reverse Path to the source based on the RPF Vector (ASBR P-2 IP address) and not based on the IP address of the multicast source (PE-1) which is unknown to it.

PE-5 inserts an RPF vector and sends a PIM-P Join to the immediate next-hop to reach ASBR P-2. Intermediate P-routers (if present) do not change the RPF vector.

P-2 finds itself in RPF Vector and has to make a decision based on MDT-SAFI BGP table:

- P-2 determines the reverse path to the multicast source based on the RPF Vector (ASBR P-3 IP address).
- If the multicast source and NH do not match, P-2 has to use RPFV.
- P-2 modifies the PIM-P Join received from PE-5 with ASBR P-3's IP address as the upstream (taken from Next-hop MDT-SAFI NLRI).
- P-3 can match the source IP with the NH in BGP MDT-SAFI. Therefore there is no need for RPF Vector to be used.
- P-3 removes the RPF vector and sends a normal PIM-P join towards PE-1.

When this stage is completed, the default MDT is established for this MVPN and PE routers have the necessary information to start PIM signaling inside VPRN (PIM-C).

Stage 3: Customer PIM signaling

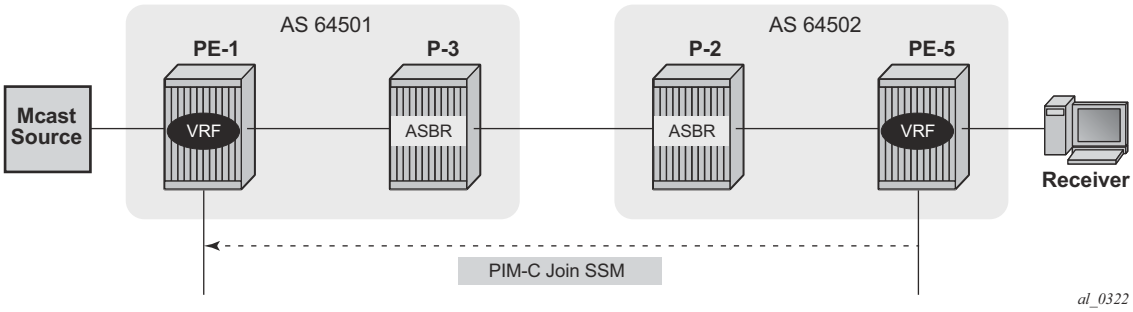

Figure 111: PIM-C Signaling

A PIM-C Join is sent to the source PE using the existing tunnel infrastructure to the RPF neighbor PE-1 provided by the BGP connector attribute of the vpn-ipv4 route of the multicast source.

When this stage is completed, the customer multicast flows throughout the network in a Default MDT.

Stage 4[1]: The Multicast stream threshold is reached.

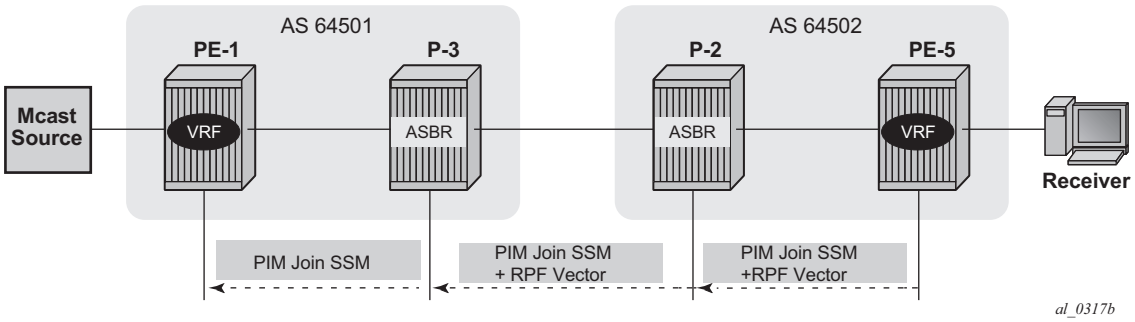

Figure 112: PIM-P Signaling Steps for Data MDT

The process is similar to the Default MDT setup:

- PE-5 determines reverse path to the source based on the RPF Vector (ASBR P-2's IP address) and not based on IP address of the multicast source (PE-1) which is unknown to it.
- PE-5 inserts an RPF vector and sends a PIM-P Join to the immediate next-hop to reach ASBR P-2.
- Intermediate P-routers (if present) do not change RPF vector.
- P-2 finds itself in the RPF Vector and has to make a decision based on the MDT-SAFI BGP table:
  - → P-2 determines reverse path to the multicast source based on the RPF Vector (ASBR P-3's IP address).
  - → If the multicast source and NH do not match, P-2 has to use the RPFV.
  - → P-2 modifies the PIM-P Join received from PE-5 with ASBR P-3's IP address as upstream (taken from Next-hop MDT-SAFI NLRI).
- P-3 can match the source IP with the NH in the BGP MDT-SAFI. Therefore there is no need for RPF Vector to be used.
- P-3 removes the RPF vector and sends a normal PIM-P join towards PE-1.

When this optional stage is completed, the customer multicast flows in a dedicated Data MDT.

1. This stage is optional and applicable when S-PMSI instance and S-PMSI threshold are configured.

Known interoperability issues:

The SR OS implementation was also designed to interoperate with Cisco routers' Inter-AS implementations that do not fully comply with the RFC 5384 and RFC5496.

When the following option is enabled:

```
configure router pim rpfv mvpn
```

Cisco routers need to be configured to include **RD** in an RPF vector using the following command for interoperability:

```
ip multicast vrf <name> rpf proxy rd vector
```

Configuration

The test topology is shown in Figure 113.

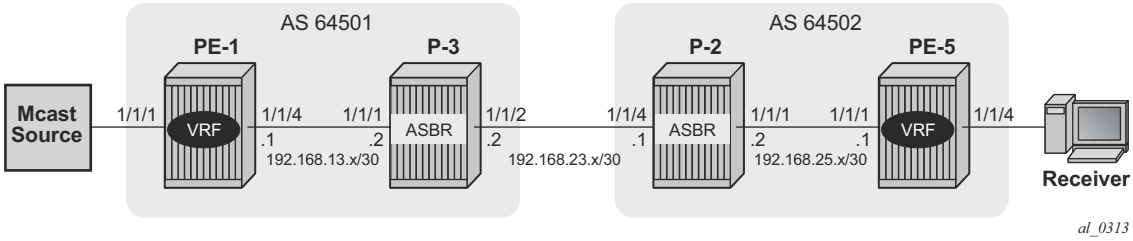

Figure 113: Test Topology Details

The following parameters are used in the test scenario:

- VPRN 1 is used
- Customer multicast group is 232.0.0.0/8
- Default MDT multicast group is 239.255.0.1
- Data MDT multicast group is 239.255.1.0/24
- Multicast source is 172.16.1.1
- PE-x routers have system IP addresses 192.0.2.x
- P-x routers have system IP addresses 192.0.2.x
- Interface between Router A and B has IP address 192.168.AB.x

Global BGP configuration for PE-1 router using the family mdt-safi with an iBGP neighbor to P-3. System interface IP address is used for iBGP session.

```
configure
  router
     bgp
        group "iBGP"
        family vpn-ipv4 mdt-safi
        type internal
        neighbor 192.0.2.3
           next-hop-self
        exit
  exit
```

Global BGP configuration for P-3 router using the family mdt-safi with an iBGP neighbor to PE-1 and an eBGP neighbor to P-2. System interface IP address is used for iBGP session and network interface IP address is used for eBGP session.

```
configure router bgp
    enable-inter-as-vpn
    group "eBGP"
        family vpn-ipv4 mdt-safi
        neighbor 192.168.23.1
          type external
          peer-as 64502
         exit
     exit
    group "iBGP"
      family vpn-ipv4 mdt-safi
       neighbor 192.0.2.1
          next-hop-self
          type internal
        exit
    exit
```

Global BGP configuration for P-2 router using the family mdt-safi with an iBGP neighbor to PE-5 and an eBGP neighbor to P-3. System interface IP address is used for iBGP session and network interface IP address is used for eBGP session.

```
configure router bgp
   enable-inter-as-vpn
   group "eBGP"
      family vpn-ipv4 mdt-safi
      neighbor 192.168.23.2
        type external
        peer-as 64501
      exit
    exit
    group "iBGP"
       family vpn-ipv4 mdt-safi
       neighbor 192.0.2.5
          next-hop-self
          type internal
        exit
    exit
```

Global BGP configuration for PE-5 router using the family mdt-safi with an iBGP neighbor to P-2. System interface IP address is used for iBGP session.

```
configure
  router
     bgp
        group "iBGP"
         family vpn-ipv4 mdt-safi
        type internal
        neighbor 192.0.2.2
           next-hop-self
         exit
       exit
```

Global PIM configuration for ALL routers.

```
configure router pim
            rpf-table both
            apply-to non-ies
            rp
                static
                exit
                bsr-candidate
                    shutdown
                exit
                rp-candidate
                    shutdown
                exit
            exit
            no shutdown
            rpfv mvpn
```

VPRN configuration for PE routers.

```
PE-x>config>service>vprn# info
----------------------------------------------
<snip>
            mvpn
                auto-discovery mdt-safi
                c-mcast-signaling pim
                    inclusive
                        pim ssm 239.255.0.1
                        exit
                    exit
                    selective
                        data-threshold 232.0.0.0/8 1
                        pim-ssm 239.255.1.0/24
                    exit
                exit
                vrf-target unicast
                exit
            exit
```

MVPN Verification and Debugging

BGP Core Signaling

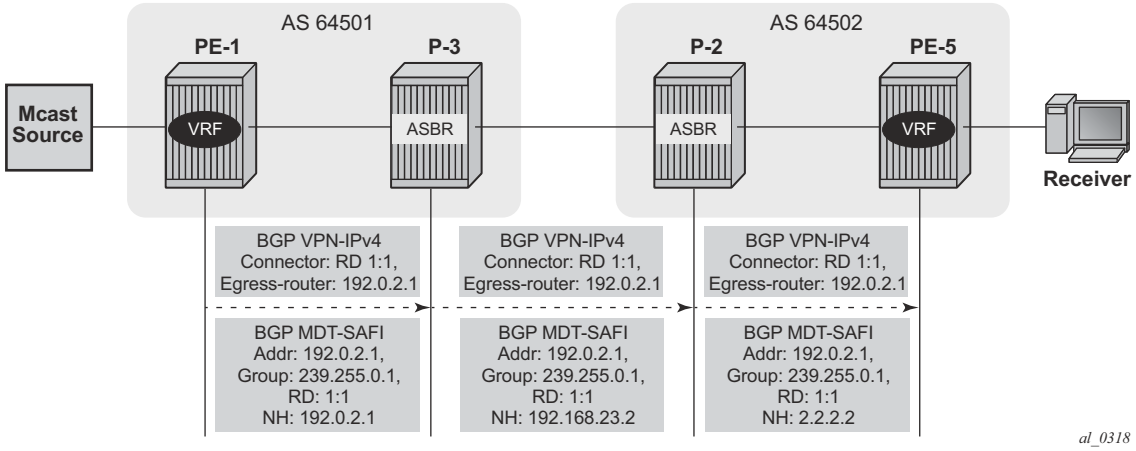

Figure 114: BGP Signaling Steps

On PE-1, the **debug router bgp update** output shows the BGP update messages which are sent to P-3. The VPN-IPv4 update contains a connector attribute and the MDT-SAFI update is used for signaling multicast group 239.255.0.1.

```
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 49 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.0.2.1
        172.16.1.0/30 RD 1:1 Label 262142
        192.0.2.1/32 RD 1:1 Label 262142
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:1:1
    Flag: 0xc0 Type: 20 Len: 14 Connector:
        RD 1:1, Egress-router 192.0.2.1

"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 62
    Flag: 0x90 Type: 14 Len: 26 Multiprotocol Reachable NLRI:
        Address Family MDT-SAFI
        NextHop len 4 NextHop 192.0.2.1
```

```
        [MDT-SAFI] Addr 192.0.2.1, Group 239.255.0.1, RD 1:1
    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:1:1
```

On P-3, the **debug router bgp update** output shows the BGP update messages which are sent to P-2. The VPN-IPv4 update contains an unmodified connector attribute and the MDT-SAFI update is used for signaling multicast group 239.255.0.1.

```
"Peer 1: 192.168.23.1: UPDATE
Peer 1: 192.168.23.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 126
    Flag: 0x90 Type: 14 Len: 81 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.168.23.2
        192.0.2.4/32 RD 1:1 Label 262142
        192.0.2.1/32 RD 1:1 Label 262142
        172.16.1.0/30 RD 1:1 Label 262142
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64501 >
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:1:1
    Flag: 0xc0 Type: 20 Len: 14 Connector:
        RD 1:1, Egress-router 192.0.2.1

"Peer 1: 192.168.23.1: UPDATE
Peer 1: 192.168.23.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 54
    Flag: 0x90 Type: 14 Len: 26 Multiprotocol Reachable NLRI:
        Address Family MDT-SAFI
        NextHop len 4 NextHop 192.168.23.2
        [MDT-SAFI] Addr 192.0.2.1, Group 239.255.0.1, RD 1:1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64501 >
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:1:1
```

On P-2, the **debug router bgp update** output shows the BGP update messages which are sent to PE-5. The VPN-IPv4 update contains an unmodified connector attribute and the MDT-SAFI update is used for signaling multicast group 239.255.0.1.

```
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 133
    Flag: 0x90 Type: 14 Len: 81 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.0.2.2
```

```
        192.0.2.4/32 RD 1:1 Label 262142
        172.16.1.0/30 RD 1:1 Label 262142
        192.0.2.1/32 RD 1:1 Label 262142
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64501 >
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:1:1
    Flag: 0xc0 Type: 20 Len: 14 Connector:
        RD 1:1, Egress-router 192.0.2.1

"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 61
    Flag: 0x90 Type: 14 Len: 26 Multiprotocol Reachable NLRI:
        Address Family MDT-SAFI
        NextHop len 4 NextHop 192.0.2.2
        [MDT-SAFI] Addr 192.0.2.1, Group 239.255.0.1, RD 1:1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64501 >
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:1:1
```

The BGP tables on PE1 and PE-5 are updated accordingly. The most interesting aspect here is MDT-SAFI routes received.

PE-5 has one MDT-SAFI update received from PE-1. The next-hop was modified accordingly to Option-B model.

```
*A:PE-5# show router bgp neighbor 192.0.2.2 received-routes mdt-safi
===============================================================================
 BGP Router ID:192.0.2.5          AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
===============================================================================
BGP MDT-SAFI Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop                    Group-Addr                        Label
      As-Path
-------------------------------------------------------------------------------
u*>i  1:1:192.0.2.1                                    100         None
      192.0.2.2                    239.255.0.1                      -
      64501
```

PE-1 has one MDT-SAFI update received from PE-5. The next-hop was modified accordingly to Option-B model.

```
*A:PE-1# show router bgp neighbor 192.0.2.4 received-routes mdt-safi
===============================================================================
 BGP Router ID:192.0.2.1          AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup

===============================================================================
BGP MDT-SAFI Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop                    Group-Addr                        Label
      As-Path
-------------------------------------------------------------------------------
u*>i  1:5:192.0.2.5                                    100         None
      192.0.2.4                    239.255.0.1                       -
      64502
```

Core PIM Signaling

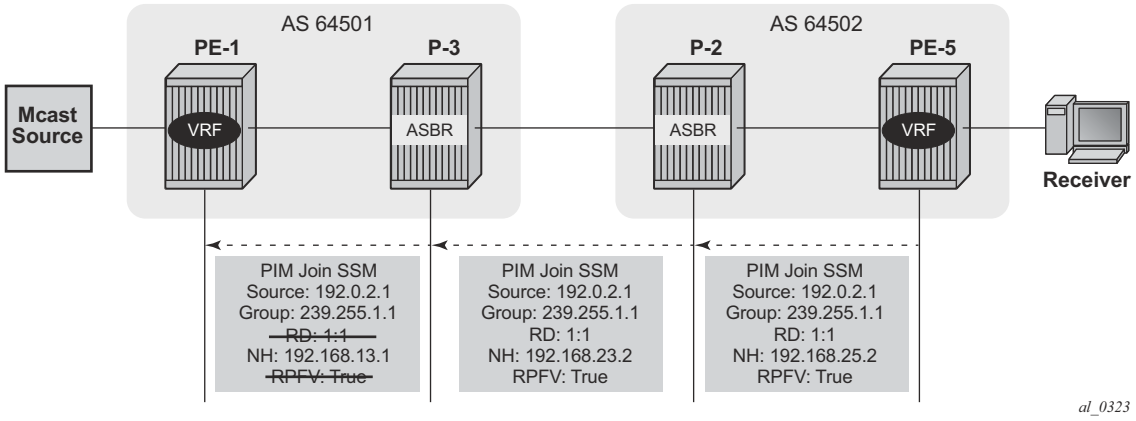

Figure 115: PIM-P Signaling Steps for Default MDT

On PE-5, the **debug router pim packet jp** output shows the PIM join/prune message which is sent to P-2. This message contains the original source of the multicast traffic (PE-1: 192.0.2.1) and the RPF Vector (P-2: 192.0.2.2).

```
"PIM[Instance 1 Base]: Join/Prune
[007 14:55:54.020] PIM-TX ifId 3 ifName int-PE5-P2  -> 224.0.0.13 Length: 48
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x8b5e
Upstream Nbr IP : 192.168.25.2 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.0.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G> JA={rpfvMvpn 192.0.2.2 1:1}
```

On P-2, **the debug router pim packet jp** output shows the PIM join/prune message which is propagated to P-3. The source of multicast traffic is untouched while the RPF Vector is modified for Inter-AS propagation.

```
"PIM[Instance 1 Base]: Join/Prune
[001 12:25:19.590] PIM-TX ifId 4 ifName int-P2-P3  -> 224.0.0.13 Length: 48
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x835e
Upstream Nbr IP : 192.168.23.2 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.0.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G> JA={rpfvMvpn 192.168.23.2 1:1}
```

On P-3, the **debug router pim packet jp** output shows the PIM join/prune message which is propagated to P-3. The source of multicast traffic is untouched while the RPF Vector is not present anymore.

```
"PIM[Instance 1 Base]: Join/Prune
```

```
[001 12:25:16.000] PIM-TX ifId 2 ifName int-P3-PE1  -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0xd694
Upstream Nbr IP : 192.168.13.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.0.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G>
```

As a result of this signaling, Default MDT is established between the two ASs. This can be checked with **show router pim group** command.

The PE-1 output shows the active multicast groups which are used as Default MDT.

```
*A:PE-1#show router pim group
===============================================================================
PIM Groups ipv4
===============================================================================
Group Address                Type                Spt Bit Inc Intf       No.Oifs
   Source Address               RP                          Inc Intf(S)
-------------------------------------------------------------------------------
239.255.0.1                   (S,G)               spt     system         2
   192.0.2.1
239.255.0.1                   (S,G)               spt     int-PE1-P3     1
   192.0.2.5
```

The PE-5 output shows active multicast groups which are used as Default MDT:

```
*A:PE-5# show router pim group
===============================================================================
PIM Groups ipv4
===============================================================================
Group Address                Type                Spt Bit Inc Intf       No.Oifs
   Source Address               RP                          Inc Intf(S)
-------------------------------------------------------------------------------
239.255.0.1                   (S,G)               spt     int-PE5-P2     1
   192.0.2.1
239.255.0.1                   (S,G)               spt     system         2
   192.0.2.5
```

The detailed information about the PIM-P group shows that the Default MDT is used to deliver traffic. Key parameters such as correct the incoming/outgoing interfaces and non-zero flow rate allow this conclusion to be made.

PE-5 has the incoming interface “int-PE5-P2”, outgoing interface “system” and flow rate of 5.4 kbps.

```
*A:PE-5#   show router pim group detail
Group Address      : 239.255.0.1
Source Address     : 192.0.2.1
RP Address         : 0
Advt Router        : 192.0.2.2
Upstream RPFV Nbr  : 192.168.25.2
RPFV Type          : Mvpn 1:1         RPFV Proxy         : 192.0.2.2

Flags              : spt               Type               : (S,G)
```

```
MRIB Next Hop      : 192.168.25.2
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:03:00
Up Time            : 0d 04:57:13        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:47
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       : 192.168.25.2
Incoming Intf      : int-PE5-P2
Outgoing Intf List : system

Curr Fwding Rate   : 5.4 kbps
Forwarded Packets  : 178895             Discarded Packets  : 0
Forwarded Octets   : 11814210           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

PE-1 has incoming the interface “system”, outgoing interfaces “system, int-PE1-P” and flow rate of 3.4 kbps.

```
A:PE-1# show router pim group detail
Group Address      : 239.255.0.1
Source Address     : 192.0.2.1
RP Address         : 0
Advt Router        : 192.0.2.1
Flags              : spt                Type               : (S,G)
MRIB Next Hop      :
MRIB Src Flags     : self
Keepalive Timer Exp: 0d 00:03:30
Up Time            : 0d 23:02:04        Resolved By        : rtable-m

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       :
Incoming Intf      : system
Outgoing Intf List : system, int-PE1-P3

Curr Fwding Rate   : 3.4 kbps
Forwarded Packets  : 826316             Discarded Packets  : 0
Forwarded Octets   : 34805244           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

Customer PIM Signaling

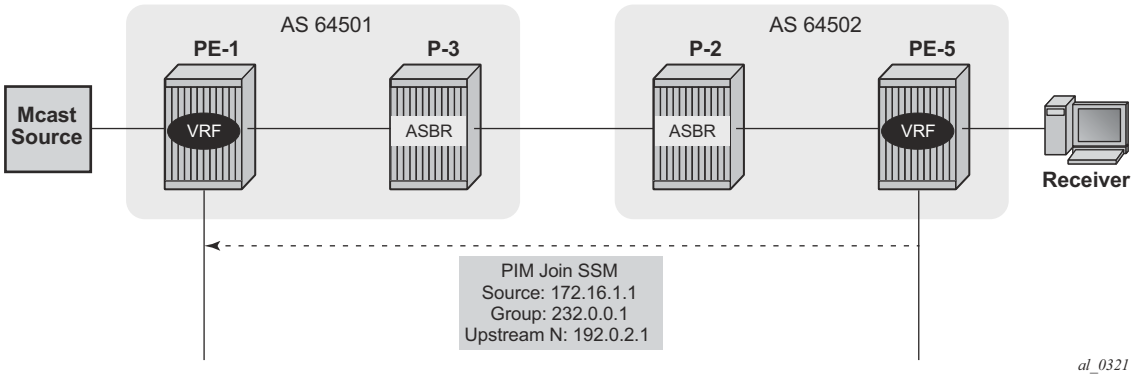

Figure 116: PIM-C Signaling

The PIM-C Join is sent to the sender PE using the existing tunnel infrastructure.

On PE-5, the **debug router 1 pim packet jp** output shows the PIM join/prune message which is sent to PE-1 using PMSI interface "1-mt-239.255.0.1" inside VPRN 1. All of this information and more can be found in the output of the **debug** command.

```
49 2013/05/02 11:23:32.11 UTC MINOR: DEBUG #2001 vprn1 PIM[Instance 9 vprn1]
"PIM[Instance 9 vprn1]: Join/Prune
[006 04:23:58.880] PIM-TX ifId 16390 ifName 1-mt-239.255.0.1  -> 224.0.0.13 Len
gth: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0xdbed
Upstream Nbr IP : 192.0.2.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 232.0.0.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                172.16.1.1/32 Flag S <S,G>
```

The detailed information about the PIM-C group for a particular VPRN shows that default MDT is used to deliver traffic. For this purpose the **show router 1 pim group detail** command is used. Key parameters such as the correct multicast group, correct incoming/outgoing interfaces and non-zero flow rate allow this conclusion to be made.

PE-1 has the incoming interface "int-source", outgoing interface "1-mt-239.255.0.1" and flow rate of 3.5 kbps.

```
*A:PE-1#show router 1 pim group detail
Group Address      : 232.0.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.4
Advt Router        : 192.0.2.4
Flags              : spt               Type              : (S,G)
MRIB Next Hop      : 172.16.1.1
```

```
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:03:22
Up Time            : 0d 06:39:09        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:50
Up JP Rpt          : Not Pruned         Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.1.1
Incoming Intf      : int-source
Outgoing Intf List : 1-mt-239.255.0.1

Curr Fwding Rate   : 3.5 kbps
Forwarded Packets  : 239467             Discarded Packets  : 0
Forwarded Octets   : 10057614           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

PE-5 has the incoming interface “1-mt-239.255.0.1”, outgoing interface “int-receiver” and flow rate of 3.5 kbps.

```
*A:PE-5 show router 1 pim group detail
Group Address      : 232.0.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.4
Advt Router        : 192.0.2.2
Flags              : spt                Type               : (S,G)
MRIB Next Hop      : 192.0.2.1
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:02:24
Up Time            : 0d 00:01:10        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:58
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      : 1-mt-239.255.0.1
Outgoing Intf List : int-receiver

Curr Fwding Rate   : 3.4 kbps
Forwarded Packets  : 649                Discarded Packets  : 0
Forwarded Octets   : 27258              RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

When Multicast Stream Threshold is Reached

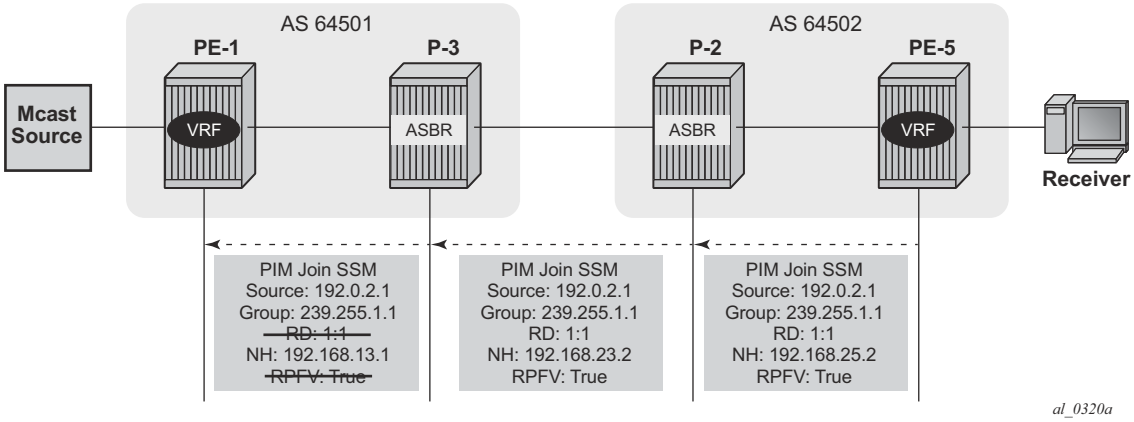

Figure 117: PIM-P Signaling Steps for Data MDT

On PE-5, the **debug router pim packet jp** output shows the PIM join/prune message which is sent to P-2. This message contains the original source of multicast traffic (PE-1: 192.0.2.1) and the RPF Vector (P-2: 192.0.2.2). Note a new multicast group (239.255.1.1) which is signalled for purposes of the Data MDT.

```
"PIM[Instance 1 Base]: Join/Prune
[000 09:48:16.140] PIM-TX ifId 3 ifName int-PE5-P2  -> 224.0.0.13 Length: 48
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x3aae
Upstream Nbr IP : 192.168.25.2 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G> JA={rpfvMvpn 192.0.2.2 1:1}
```

On P-2, the **debug router pim packet jp** output shows the PIM join/prune message which is propagated to P-3. The source of multicast traffic is untouched while the RPF Vector is modified for Inter-AS propagation.

```
"PIM[Instance 1 Base]: Join/Prune
[001 22:30:36.630] PIM-TX ifId 4 ifName int-P2-P3  -> 224.0.0.13 Length: 48
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x32ae
Upstream Nbr IP : 192.168.23.2 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G> JA={rpfvMvpn 192.168.23.2 1:1}
```

On P-3, the **debug router pim packet jp** output shows the PIM join/prune message which is propagated to P-3. The source of multicast traffic is untouched while the RPF Vector is not present anymore.

```
"PIM[Instance 1 Base]: Join/Prune
[001 22:30:32.770] PIM-TX ifId 2 ifName int-P3-PE1  -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x85e4
Upstream Nbr IP : 192.168.13.1 Resvd: 0x0, Num Groups 1, HoldTime 210
        Group: 239.255.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
        Join Srcs:
                192.0.2.1/32 Flag S <S,G>
```

As a result of this signaling, the Data MDT is established between the two ASs. This can be checked with **show router pim group** command.

The PE-1 output shows an additional multicast group (239.255.1.3), which was created in the global routing table (GRT).

```
*A:PE-1# show router pim group
===============================================================================
PIM Groups ipv4
===============================================================================
Group Address                Type                Spt Bit Inc Intf       No.Oifs
   Source Address              RP                          Inc Intf(S)
-------------------------------------------------------------------------------
239.255.0.1                   (S,G)               spt     system         2
   192.0.2.1
239.255.0.1                   (S,G)               spt     int-PE1-P3     1
   192.0.2.5
239.255.1.3                   (S,G)                       system         1
   192.0.2.1
```

The PE-5 output shows an additional multicast group (239.255.1.3), which was created in the GRT.

```
*A:PE-5# show router pim group
===============================================================================
PIM Groups ipv4
===============================================================================
Group Address                Type                Spt Bit Inc Intf       No.Oifs
   Source Address              RP                          Inc Intf(S)
-------------------------------------------------------------------------------
239.255.0.1                   (S,G)               spt     int-PE5-P2     1
   192.0.2.1
239.255.0.1                   (S,G)               spt     system         2
   192.0.2.5
239.255.1.3                   (S,G)                       int-PE5-P2     1
   192.0.2.1
```

The detailed information about the PIM group in a VPRN shows that the Data MDT is used to receive traffic instead of Default MDT.

The PE-5 output for multicast groups in a VPRN 1 has slightly changed: new line “Incoming SPMSI Intf” was added. This indicates that the S-PMSI instance and dedicated Data MDT are used for this particular multicast group. The non-zero rate for the multicast flow is also an indication that multicast traffic is forwarded.

```
*A:PE-5#show router 1 pim group detail
===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.0.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.4
Advt Router        : 192.0.2.2
Flags              : spt                Type               : (S,G)
MRIB Next Hop      : 192.0.2.1
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:01:10
Up Time            : 0d 00:30:21        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:40
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      : 1-mt-239.255.0.1
Incoming SPMSI Intf: 1-mt-239.255.0.1*
Outgoing Intf List : int-receiver

Curr Fwding Rate   : 3.4 kbps
Forwarded Packets  : 18187              Discarded Packets  : 0
Forwarded Octets   : 763854             RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

The **show router 1 pim s-pmsi detail** command can also be used to verify existence of the S-PMSI instance for the VPRN 1. The output is short, but very informative: the multicast group inside VPRN, multicast source IP, multicast group which is used for S-PMSI tunneling and current flow rate can be found.

```
*A:PE-5#show router 1 pim s-pmsi detail
===============================================================================
PIM Selective provider tunnels
===============================================================================
Md Source Address  : 192.0.2.1          Md Group Address   : 239.255.1.1
Number of VPN SGs  : 1                  Uptime             : 0d 00:29:57
MT IfIndex         : 24580              Egress Fwding Rate : 3.4 kbps
VPN Group Address  : 232.0.0.1          VPN Source Address : 172.16.1.1
State              : RX Joined
Expiry Timer       : 0d 00:02:23
```

Conclusion

Inter-AS MVPN offers flexibility for the operators who can use it to provide additional value added services to their customers. Before implementing this feature in the network the following are required:

- The RPF vector must be enabled on every router for inter-AS MVPN.
- Can be used only with Draft-Rosen mVPN with PIM SSM and MDT SAFI.