Next Generation MULTICAST - Default MDT GRE (BGP AD - PIM C: Profile 3)

Contents

Introduction

This document describes the Default Multicast Distribution Tree (MDT) GRE (BGP AD - PIM C) for Multicast over VPN (mVPN). It uses an example and the implementation in Cisco IOS in order to illustrate the behavior.

What is Default MDT?

It is used to connect multicast to all the PE in one VRF. Default means it connects all the PE routers. By default, it carries all the traffic. The all the PIM control traffic and the Data plane traffic. Example: (*,G) Traffic and the (S,G) traffic. The default is the must. This Default MDT connects all the PE router to connect. This represents multipoint to multipoint. Anybody can send and everybody can receive from the tree.

What is Data MDT?

It’s optional and is created in demand. It carries specific (S,G) traffic. In the latest IOS release, you have threshold configured as 0 and infinite. Whenever a first packet hits the VRF, the Data MDT initialized, and if infinity then the Data MDT is never created, and the traffic moves forward in the

default MDT. The Data MDT is always the receiving tree, they never send any traffic. Data MDT is only for the (S,G) traffic.

The threshold at which the data MDT is created can be configured on a per-router or a per-VRF basis. When the multicast transmission exceeds the defined threshold, the sending PE router creates the data MDT and sends a User Datagram Protocol (UDP) message, which contains information about the data MDT to all routers on the default MDT. The statistics to determine whether a multicast stream has exceeded the data MDT threshold are examined once every second.

> **Note**: After a PE router sends the UDP message, it waits 3 more seconds before switching over; 13 seconds is the worst case switchover time and 3 seconds is the best case.

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. They are not created for (*, G) entries regardless of the value of the individual source data rate

- Permits PE to directly join a source tree for an MDT.
- No Rendezvous Points are needed in the network.
- RPs are a potential failure point and additional overheads.
- But they allow shared and BiDir trees (less state).
- Reduce forwarding delay.
- Avoid management overhead to administer group/RP mapping and redundant RPs for reliability.
- Trade off is more state required.
- (S, G) for each mVPN in a PE.

If there are 5 PEs each holding mVRF RED, there is 5 x (S, G) entries.

1. Configure the ip pim ssm range command on both P and PE routers (avoids unnecessary (*, G) entries being created).
2. SSM recommended for Data-MDTs.
3. Use BiDir if possible for Default-MDT (BiDir support is platform specific).

If SSM is not used for setting up Data MDTs:

- Each VRF needs to be configured with a unique set of multicast P-addresses; two VRFs in the same MD cannot be configured with the same set of addresses.

- Many more multicast P-addresses are needed.
- Complicated operations and management.
- SSM requires the PE to join a (S, G) not (*, G).

G is known as it is configured but PE does not directly know the value of S (S, G) of Default MDT propagated by MP-BGP.

The advantage of SSM is that it is not dependent on the use of an RP to derive the source PE router for a particular MDT group.

The IP address of the source PE and default MDT group is sent via Border Gateway Protocol (BGP)

There are two ways in which BGP can send this information:

- Extended community Cisco proprietary solutionNon-transitive attribute (not suitable for inter-AS)
- BGP address family MDT SAFI (66) **draft-nalawade-idr-mdt-safi**

  Note: GRE MVPNs were supported before using MDT SAFI; actually, even before MDT SAFI by using RD type 2. Technically, for Profile 3, MDT SAFI should not be configured, but both SAFIs are simultaneously supported for migration.

BGP

- Source PE and MDT Default Group encoded in NLRI of **MP_REACH_NLRI**.
- RD is the same as that of the MVRF for which the MDT Default Group is configured.
- RD Type is 0 or 1

```
▼ Path Attribute - MP_REACH_NLRI
  ▶ Flags: 0x80, Optional: Optional, Non-transitive, Complete
    Type Code: MP_REACH_NLRI (14)
    Length: 23
    Address family identifier (AFI): IPv4 (1)
    Subsequent address family identifier (SAFI): MCAST-VPN (5)
    Next hop network address (4 bytes)
    Number of Subnetwork points of attachment (SNPA): 0
  ▼ Network layer reachability information (14 bytes)
      Route Type: Intra-AS I-PMSI A-D route (1)
      Length: 12
▶ Path Attribute - ORIGIN: INCOMPLETE
▶ Path Attribute - AS_PATH: empty
▶ Path Attribute - MULTI_EXIT_DISC: 0
▶ Path Attribute - LOCAL_PREF: 100
▶ Path Attribute - COMMUNITIES: NO_EXPORT
▶ Path Attribute - EXTENDED_COMMUNITIES
▼ Path Attribute - PMSI_TUNNEL_ATTRIBUTE
  ▶ Flags: 0xc0, Optional, Transitive: Optional, Transitive, Complete
    Type Code: PMSI_TUNNEL_ATTRIBUTE (22)
    Length: 13
    Flags: 0
    Tunnel Type: PIM SSM Tree (3)
  ▶ MPLS Label Stack:  (withdrawn)
  ▼ Tunnel ID: < 1.1.1.1, 239.232.0.0 >
      PIM-SSM Tree tunnel Root Node: 1.1.1.1
      PIM-SSM Tree tunnel P-multicast group: 239.232.0.0
```

PMSI attribute carry the Source address and the Group Address. In order to form the MT tunnel.

Multicast Addressing for SSM Group

232.0.0.0 – 232.255.255.255 has been reserved for global Source Specific Multicast applications.

239.0.0.0 – 239.255.255.255 is the administratively scoped IPv4 multicast address space range

The IPv4 Organization Local Scope - **239.192.0.0/14**

The Local Scope is the minimal enclosing scope, and hence is not further divisible.

The ranges **239.0.0.0/10, 239.64.0.0/10** and **239.128.0.0/10** are unassigned and available for expansion of this space.

These ranges should be left unassigned until the **239.192.0.0/14** space is no longer sufficient.

Recommendations

- Default-MDT should draw addresses from the 239/8 space starting with the range defined with the Organizational local scope of 239.192.0.0/14
- Data-MDT should draw addresses from the Organizational Local Scope.
- It is also possible to use the SSM global range 232.0.0.0 – 232.255.255.255
- As SSM always uses a unique (S, G) state, there is no possibility of overlap as the SSM multicast stream will be initiated by different sources (with different addresses) whether they

be within the provider network or the greater Internet.

- Use the same Data-MDT pool for every mVRF within a particular multicast domain (where the Default-MDT is common).

For example, all VRFs using Default-MDT 239.192.10.1 should use the same Data MDT 239.232.1.0/24 range

Overlay Signaling

Overlay signaling of Rosen GRE is shown in the image.

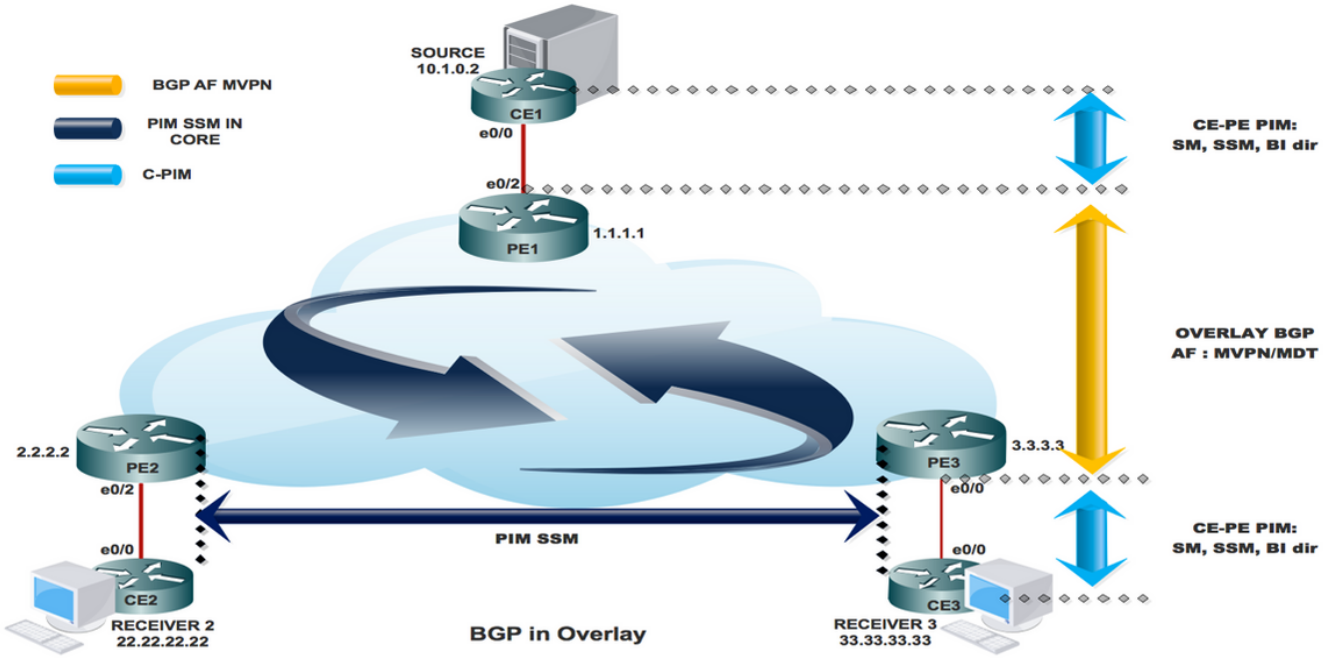

Topology

Topology of Rosen GRE is shown in the image.

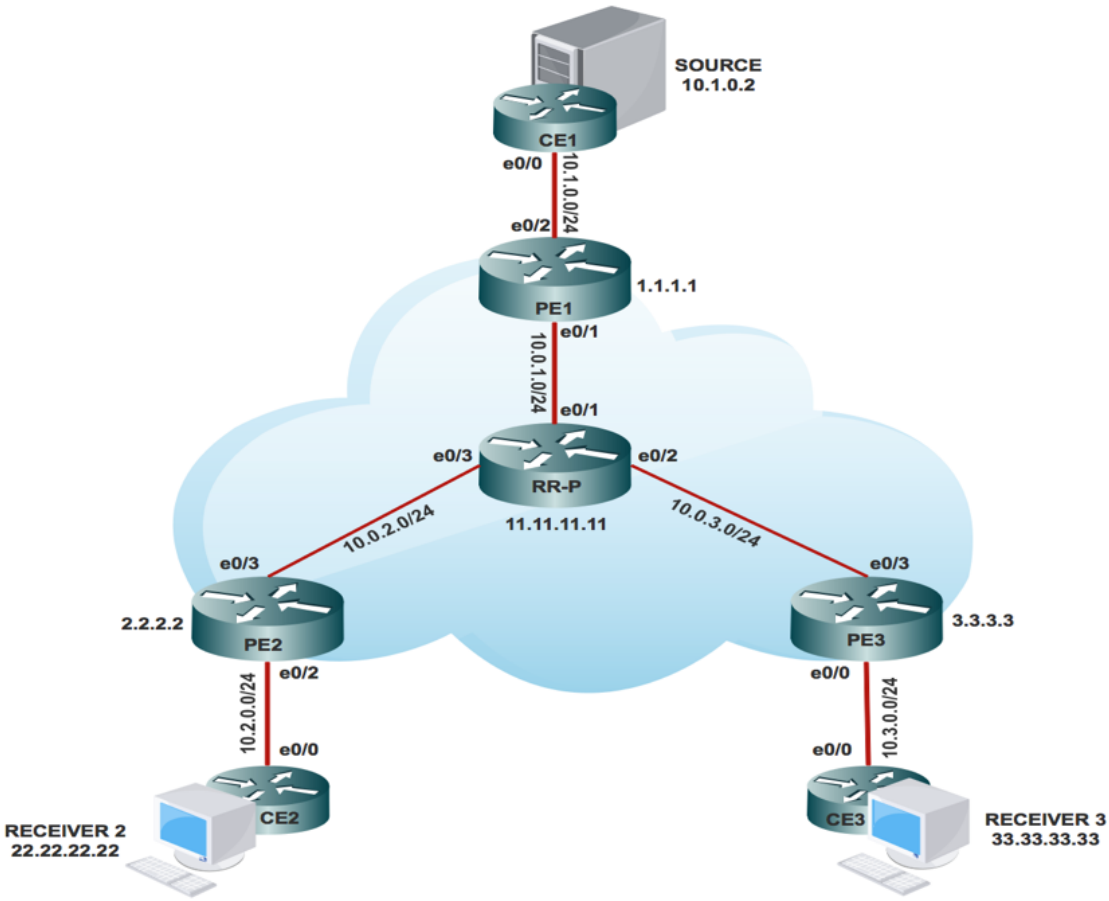

Multicast VPN Routing and Forwarding and Multicast Domains

MVPN introduces multicast routing information to the VPN routing and forwarding table. When a Provider Edge (PE) router receives multicast data or control packets from a Customer Edge (CE) router, forwarding is performed according to the information in the Multicast VPN Routing and Forwarding instance (MVRF). MVPN does not use label switching.

A set of MVRFs that can send multicast traffic to each other constitutes a multicast domain. For example, the multicast domain for a customer that wanted to send certain types of multicast traffic to all global employees would consist of all CE routers associated with that enterprise.

Configuration Tasks

1. Enable Multicast Routing on all nodes.

2. Enable Protocol Independent Multicast (PIM) Sparse Mode in all the interface.

3. With existing VRF configure Default MDT.

4. Configure the VRF on the interface Ethernet0/x.

5. Enable Multicast Routing on VRF.

6. Configure PIM SSM Default in all nodes inside the core.

7. Configure BGP Address Family MVPN.

8. Configure BSR RP in CE Node.

9. Pre-Configured:

```
VRF SSM-BGP
mBGP: Address family VPNv4
VRF Routing Protocol
```

Configuration Steps:

Enable Multicast Routing

On All Nodes

```
(conf) # ip multicast-routing
```

Enable "ip multicast-routing" in global mode on all nodes.

Enable PIM Sparse Mode

Enable on all connected Interface

```
(config)#interface Ethernet0/x
(config-if)#ip pim sparse-mode

(config)# interface lo0
(config-if)# ip pim sparse-mode
```

"x" represents the connected interface number on all nodes

Configure Default MDT Group in VRF

On PE1, PE2 and PE3

```
(config)#ip vrf SSM-BGP
(config-vrf)# mdt auto-discovery pim
(config-vrf)# mdt default 239.232.0.0
```

SERVICE PROVIDER : Group : 239.232.0.0 Source : 1.1.1.1

Configure the VRF on the interface Ethernet0/x

On PE1, PE2 and PE3

```
(config)#interface Ethernet0/x
(config-if)# ip vrf forwarding SSM-BGP
(config-if)# ip address 10.x.0.1 255.255.255.0
(config-if)# ip pim sparse-mode
```

"x" represent the interface number that PE connected to CE.

Enable Multicast Routing on VRF

On PE1, PE2 and PE3

```
(conf) # ip multicast-routing vrf SSM-BGP
```

Enable "ip multicast-routing m-GRE" in global mode.

Configure PIM SSM Default in all nodes inside the core.

On PE1, PE2, PE3 and RR-P Node

```
(config) # ip pim ssm default
```

Static RP configuration in the core in global mode.

Configure BSR RP in CE Node (Receiver)

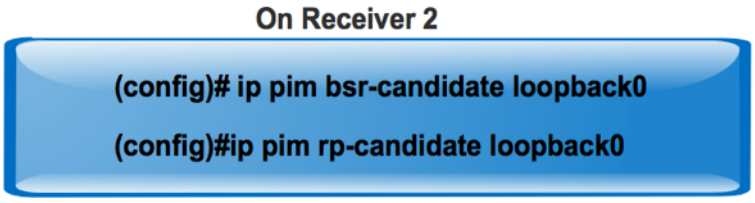

BSR RP configuration in the Receiver 2 in global mode.

Verify

Task 1: Verify Physical Connectivity.

Verify all the connected interface are **UP**.

Task 2: Verify BGP Address Family VPNv4 unicast.

- Verify that BGP is enabled in all the routers for AF VPNv4 unicast and BGP neighbors are **UP**.
- Verify that BGP VPNv4 unicast table has all the Customer prefixes.

Task 3: Verify BGP Address Family MVPN unicast.

- Verify that BGP is enabled in all the routers for AF IPV4 MVPN and BGP neighbors are **UP**.
- Verify that all PE discovery each other, with Type 1 Route.

Task 4: Verify Multicast Traffic end to end.

- Check PIM neighborship.
- Verify that multicast state is created in the VRF.
- Verify mRIB entry on PE1, PE2 and PE3.
- Verify that (S, G) mFIB entry, packet getting incremented in software forwarding.
- Verify ICMP packets getting reach from CE to CE.

Task 1: Verify Physical Connectivity

Verify all the connected interface are "UP"

```
#sh ip interface brief
```

Task 2: Verify Address Family VPNv4 unicast

Address Family VPNv4 unicast and BGP neighbors

```
# show running-config | s r bgp
# show bgp vpnv4 unicast summary all
```

VPNv4 unicast table has all the Customer prefixes

```
PE1#sh bgp vpnv4 unicast all

BGP table version is 31, local router ID is 1.1.1.1

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:100 (default for vrf m-GRE)

 *>i 22.22.22.22/32   2.2.2.2                  0    100      0 20 i
 *>i 33.33.33.33/32   3.3.3.3                  0    100      0 30 i
 *>  111.111.111.111/32
                       10.1.0.2                 0             0 10 i
```

Check on all the PE nodes (PE1, PE2 and PE3)

Task 3: Verify Address Family IPv4 MVPN

Address Family IPv4 MVPN and BGP neighbors

```
# show running-config | s r bgp
# #sh bgp ipv4 mvpn all
```

IPv4 MVPN table has all the PE routes with Type 1 routes

```
PE1#sh bgp ipv4 mvpn all
BGP table version is 15, local router ID is 1.1.1.1

Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf SSM-BGP)
 *>   [1][1:1][1.1.1.1]/12
                      0.0.0.0                            32768 ?
 *>i  [1][1:1][2.2.2.2]/12
                      2.2.2.2                  0    100      0 ?
 *>i  [1][1:1][3.3.3.3]/12
                      3.3.3.3                  0    100      0 ?
Route Distinguisher: 2:2
 *>i  [1][2:2][2.2.2.2]/12
                      2.2.2.2                  0    100      0 ?
Route Distinguisher: 3:3
     Network          Next Hop            Metric LocPrf Weight Path
 *>i  [1][3:3][3.3.3.3]/12
                      3.3.3.3                  0    100      0 ?
```

Check on all the PE nodes (PE1, PE2 and PE3)

Verify that (S,G) mFIB entry, packet getting incremented

```
PE1#sh ip mfib vrf SSM-BGP 225.1.1.1 verbose

I/O Item Flags:
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept,

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   FS Pkt Count/PS Pkt Count
VRF SSM-BGP
 (10.1.0.2,225.1.1.1) Flags: K DDE
   SW Forwarding: 10/0/100/0, Other: 2/1/1
   Ethernet0/2 Flags: RA A MA
   Tunnel0, MDT/239.232.0.0 Flags: RF F NS
     CEF: Adjacency with MAC: 4500000000000000FF2FC9E401010101EFE8000000000800
     Pkts: 10/0
```

Verify that multicast state is created in the VRF

```
PE1#sh ip mroute vrf SSM-BGP verbose
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, p - PIM Joins on route,


(10.1.0.2, 225.1.1.1), 00:00:03/00:02:56, flags: Tp
  Incoming interface: Ethernet0/2, RPF nbr 10.1.0.2
  Outgoing interface list:
    Tunnel0, GRE MDT: 239.232.0.0 (default), Forward/Sparse, 00:00:03/00:03:26, p
```

Check on all the PE nodes (PE1, PE2 and PE3)

Verify that (S,G) mFIB entry, packet getting incremented

```
PE1#sh ip mfib vrf SSM-BGP 225.1.1.1 verbose

I/O Item Flags:
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept,

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   FS Pkt Count/PS Pkt Count
VRF SSM-BGP
 (10.1.0.2,225.1.1.1) Flags: K DDE
   SW Forwarding: 10/0/100/0, Other: 2/1/1
   Ethernet0/2 Flags: RA A MA
   Tunnel0, MDT/239.232.0.0 Flags: RF F NS
     CEF: Adjacency with MAC: 4500000000000000FF2FC9E401010101EFE8000000000800
     Pkts: 10/0
```

mRIB in the Service Provider Core.

```
PE1#sh ip mroute verbose
IP Multicast Routing Table
Flags: s - SSM Group, C - Connected,
       T - SPT-bit set,
       I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       p - PIM Joins on route,

(1.1.1.1, 239.232.0.0), 01:00:33/00:03:03, flags: sTp
  Incoming interface: Loopback0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse, 01:00:33/00:03:03, p

(3.3.3.3, 239.232.0.0), 01:00:33/stopped, flags: sTIZ
  Incoming interface: Ethernet0/1, RPF nbr 10.0.1.2
  Outgoing interface list:
    MVRF SSM-BGP, Forward/Sparse, 01:00:33/00:02:26

(2.2.2.2, 239.232.0.0), 01:00:33/stopped, flags: sTIZ
  Incoming interface: Ethernet0/1, RPF nbr 10.0.1.2
  Outgoing interface list:
    MVRF SSM-BGP, Forward/Sparse, 01:00:33/00:02:26
```

Check on all the PE nodes (PE1, PE2 and PE3)

Verify ICMP packets getting reach from CE to CE

```
SOURCE1#ping 225.1.1.1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 225.1.1.1, timeout is 2 seconds:

Reply to request 0 from 10.3.0.2, 29 ms
Reply to request 0 from 10.3.0.2, 29 ms
```

How Tunnel Interfaces are Created?

MDT Tunnel Creation

Once we configure mdt default 239.232.0.0

Tunnel 0 came up and assigned its Loopback 0 address as source.

%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up

```
PIM(1): Check DR after interface: Tunnel0 came up!
PIM(1): Changing DR for Tunnel0, from 0.0.0.0 to 1.1.1.1 (this system)
%PIM-5-DRCHG: VRF SSM-BGP: DR change from neighbor 0.0.0.0 to 1.1.1.1 on interface Tunnel0
```

This image shows MDT Tunnel Creation.

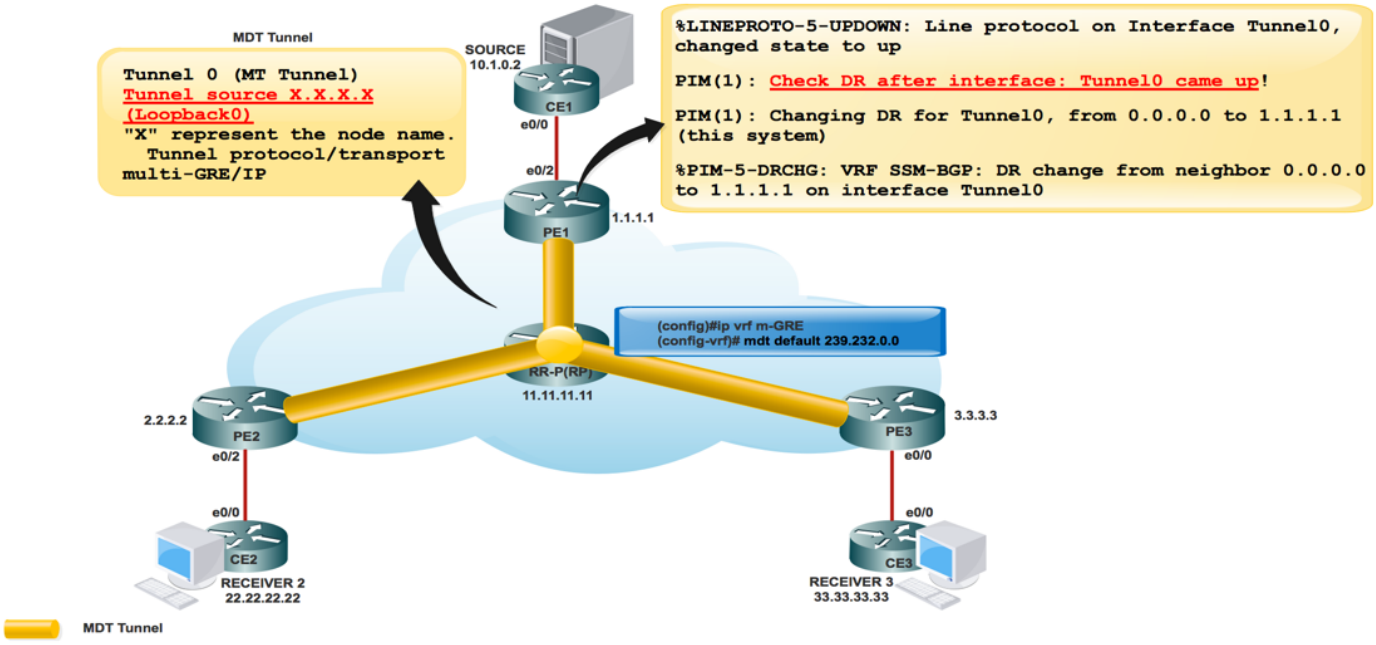


```
PE1#sh int tunnel 0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered. Using address of Loopback0 (1.1.1.1)
  MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 1.1.1.1 (Loopback0)
   Tunnel Subblocks:
      src-track:
         Tunnel0 source tracking subblock associated with Loopback0
          Set of tunnels with source Loopback0, 1 member (includes iterators), on interface <OK>
  Tunnel protocol/transport multi-GRE/IP
    Key disabled, sequencing disabled
    Checksumming of packets disabled
```

As soon as BGP MVPN comes UP, all the PE discover each other via Type 1 route. Multicast Tunnel formed. BGP carries all the Group and Source PE address in PMSI attribute.

This image shows Exchange of Type 1 route.

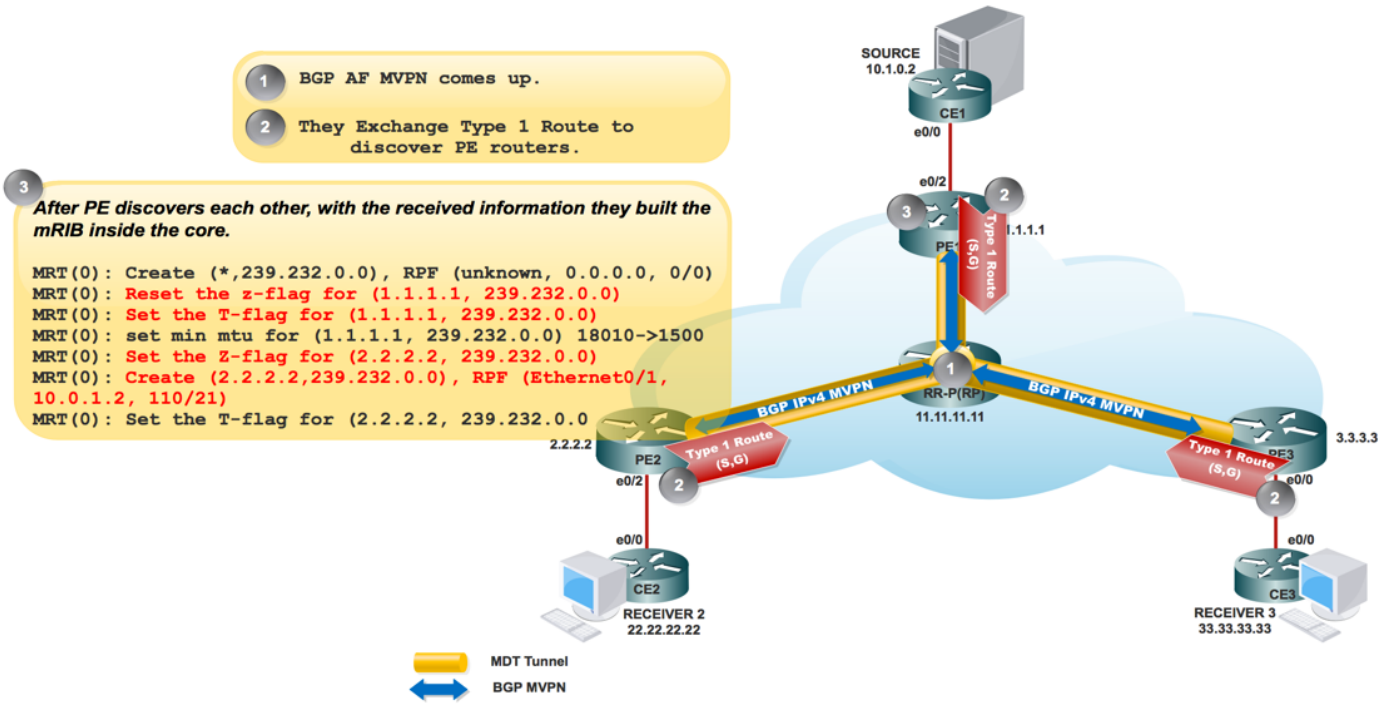

This image shows PCAP-1.

```
▼ Path attributes
  ▼ Path Attribute – MP_REACH_NLRI
    ▶ Flags: 0x80, Optional: Optional, Non-transitive, Complete
      Type Code: MP_REACH_NLRI (14)
      Length: 23
      Address family identifier (AFI): IPv4 (1)
      Subsequent address family identifier (SAFI): MCAST-VPN (5)
      Next hop network address (4 bytes)
      Number of Subnetwork points of attachment (SNPA): 0
    ▼ Network layer reachability information (14 bytes)
        Route Type: Intra-AS I-PMSI A-D route (1)  ──►  Type 1 Route
        Length: 12
  ▶ Path Attribute – ORIGIN: INCOMPLETE
  ▶ Path Attribute – AS_PATH: empty
  ▶ Path Attribute – MULTI_EXIT_DISC: 0
  ▶ Path Attribute – LOCAL_PREF: 100
  ▶ Path Attribute – COMMUNITIES: NO_EXPORT
  ▶ Path Attribute – EXTENDED_COMMUNITIES
  ▼ Path Attribute – PMSI_TUNNEL_ATTRIBUTE
    ▶ Flags: 0xc0, Optional, Transitive: Optional, Transitive, Complete
      Type Code: PMSI_TUNNEL_ATTRIBUTE (22)
      Length: 13
      Flags: 0
      Tunnel Type: PIM SSM Tree (3)  ──►  PIM SSM TREE (Tunnel Type)
    ▶ MPLS Label Stack:  (withdrawn)
    ▼ Tunnel ID: < 1.1.1.1, 239.232.0.0 >
        PIM-SSM Tree tunnel Root Node: 1.1.1.1
        PIM-SSM Tree tunnel P-multicast group: 239.232.0.0  ◄──►  PIM SSM Tree Tunnel Root and Group
```

```
PE1#sh ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,

 (3.3.3.3, 239.232.0.0), 00:01:41/00:01:18, flags: sTIZ
  Incoming interface: Ethernet0/1, RPF nbr 10.0.1.2
  Outgoing interface list:
    MVRF SSM-BGP, Forward/Sparse, 00:01:41/00:01:18

(2.2.2.2, 239.232.0.0), 00:01:41/00:01:18, flags: sTIZ
  Incoming interface: Ethernet0/1, RPF nbr 10.0.1.2
```

```
  Outgoing interface list:
    MVRF SSM-BGP, Forward/Sparse, 00:01:41/00:01:18
```

“Z” Multicast Tunnel formed after BGP mVPN comes up, as it advertises the Source PE and Group Address in PMSI attribute.

PIM Neighborship

```
PE1#sh ip pim vrf SSM-BGP neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
10.1.0.2          Ethernet0/2           00:58:18/00:01:31 v2  1 / DR S P G
3.3.3.3           Tunnel0                00:27:44/00:01:32 v2    1 / S P G
2.2.2.2           Tunnel0                00:27:44/00:01:34 v2    1 / S P G
```

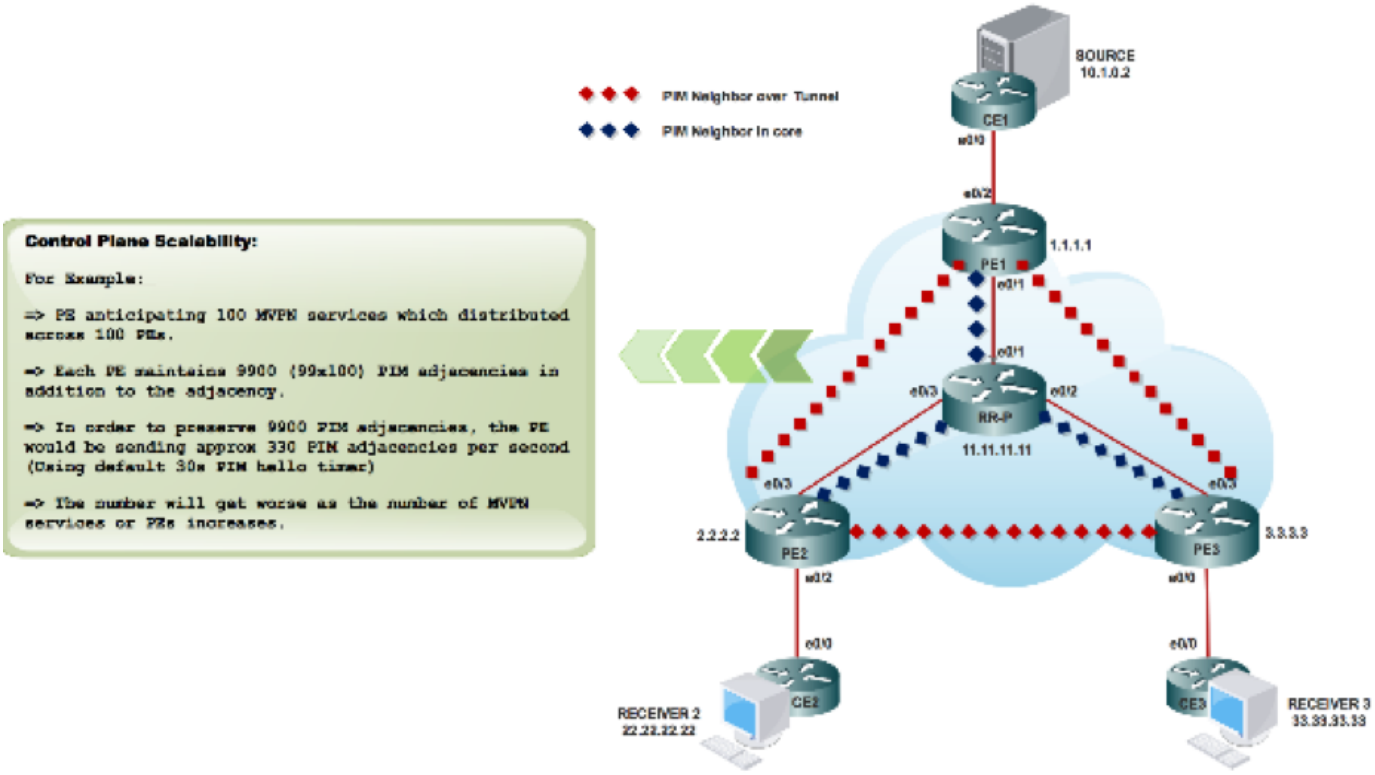

As soon as you configure RP information:

%LINEPROTO-5-UPDOWN: **Line protocol on Interface Tunnel1, changed state to up**

The bootstrap message exchange via MDT tunnel

```
PIM(1): Received v2 Bootstrap on Tunnel0 from 2.2.2.2
PIM(1): pim_add_prm:: 224.0.0.0/240.0.0.0, rp=22.22.22.22, repl = 0, ver =2, is_neg =0, bidir =
0, crp = 0
PIM(1): Update
prm_rp->bidir_mode = 0 vs bidir = 0 (224.0.0.0/4, RP:22.22.22.22), PIMv2
*May 18 10:28:42.764: PIM(1): Received RP-Reachable on Tunnel0 from 22.22.22.22
```

This image shows bootstrap message exchange via MDT tunnel.

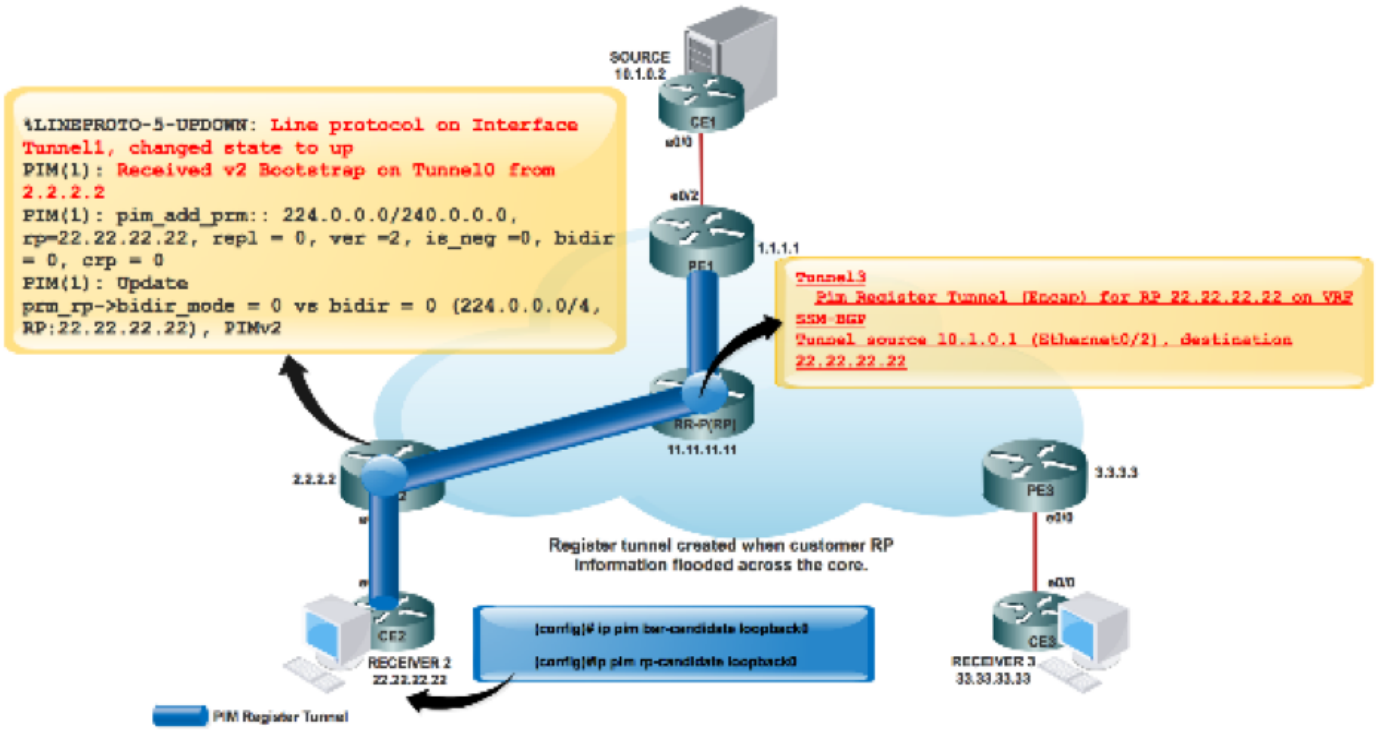


```
PE2#sh int tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  Description: Pim Register Tunnel (Encap) for RP 22.22.22.22 on VRF SSM-BGP
  Interface is unnumbered. Using address of Ethernet0/2 (10.2.0.1)
  MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.2.0.1 (Ethernet0/2), destination 22.22.22.22
   Tunnel Subblocks:
      src-track:
         Tunnel1 source tracking subblock associated with Ethernet0/2
          Set of tunnels with source Ethernet0/2, 1 member (includes iterators), on interface
<OK>
  Tunnel protocol/transport PIM/IPv4
  Tunnel TOS/Traffic Class 0xC0,  Tunnel TTL 255
  Tunnel transport MTU 1472 bytes
  Tunnel is transmit only
```

Two Tunnel formed PIM register tunnel and MDT Tunnel.

- Tunnel 0 is used to send the PIM Join and low bandwidth multicast traffic.
- Tunnel 1 is used to send the PIM encapsulated Register message.

Command to check :

**MDT BGP:

PE1#sh ip pim vrf m-SSM mdt bgp

** send data FHR:

PE1#sh ip pim vrf m-SSM mdt

Flag	Name	Description
Z	Multicast Tunnel	Indicates that this entry is an IP multicast group that belongs to the Default or Data MDT tunnel. All packets received for this IP multicast state are sent to the MDT tunnel for decapsulation . Set on receiving PE. Global mulitcast routing table
Y	Joined MDT-data group	Indicates that the traffic was received through a Data MDT tunnel that was set up specifically for this source and group. MVRF multicast routing table
z	MDT-data group sender	Set on sending PE. Global mulitcast routing table
y	Sending to MDT-data group	Indicates that the traffic was sent through a Data MDT tunnel that was set up specifically for this source and group. MVRF multicast routing table
V	RD & Vector	
v	Vecor	
E	Extranet source mroute entry	Indicates that a (*, G) or (S, G) entry in the VRF routing table is a source Multicast VRF (MVRF) entry and has extranet receiver MVRF entries linked to it

Related Information

- https://tools.ietf.org/html/rfc4760
- https://tools.ietf.org/html/rfc5110
- https://tools.ietf.org/html/rfc6513
- Technical Support & Documentation - Cisco Systems