Configure and Verify vEdge Multicast Overlay Routing

Contents

Introduction

This document describes how to configure multicast in an SD-WAN environment and is specific for vEdge routers. All the configurations are based on Protocol Independent Multicast (PIM) Auto-Rendezvous Point (RP). It shows a sample network scenario, configuration, and verification outputs.

Prerequisites

Requirements

There are no specific requirements for this document. However, a basic understanding of multicast and working knowledge of SD-WAN can help.

Components Used

This document is not restricted to the specific software or hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Here you can find a list of acronyms used in this article.

- vEdge (VE)

- First Hop Router (FHR)
- Last Hop Router (LHR)
- Rendezvous Point (RP)
- Virtual Private Network (VPN)
- Overlay Management Protocol (OMP)
- Transport Location (TLOC)
- Internet Group Management Protocol (IGMP)
- Cloud Service Router (CSR)
- Protocol Independent Multicast (PIM)
- Multicast Routing Information Base (MRIB) or Multicast Routing Table
- Reverse Path Forwarding (RPF)
- Time To Live (TTL)

For a detailed description of SD-WAN terminology, refer to Cisco SD-WAN Terminology

Configure

For Cisco SD-WAN multicast general overview, refer to Multicast Overlay Routing Overview.

Network Diagram

> **Note**: In this topology, both BR1-VE-1 and BR3-VE-1 have GOLD TLOC in common. In real scenarios, sites can have same or different TLOCs.

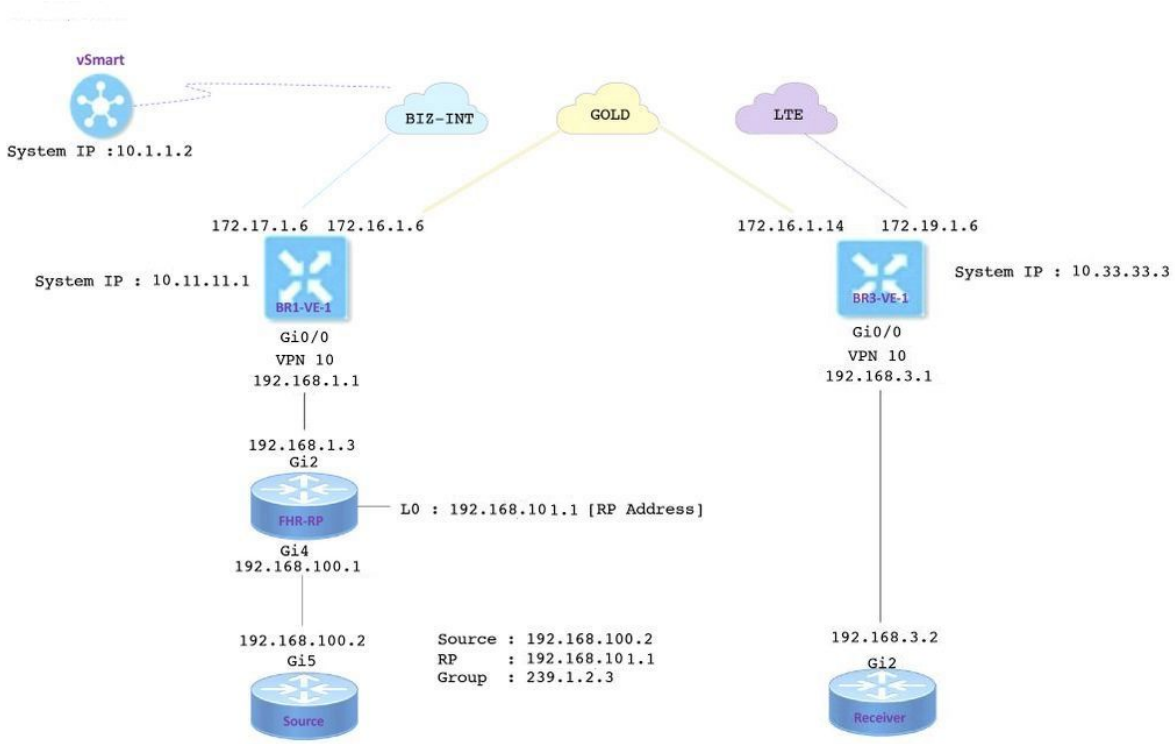

Configurations

BR1-VE-1 has SD-WAN overlay/underlay basic configuration with a default route. Besides this, local multicast replicator and PIM has been configured on Ge0/0 interface. The command **multicast-replicator local** configures the VE router as a multicast replicator.

```
vpn 10
 router
  multicast-replicator local
  pim
   auto-rp
   interface ge0/0
   exit
 !
 interface ge0/0
  ip address 192.168.1.1/24
  no shutdown
```

BR3-VE-1 has SD-WAN overlay/underlay basic configuration with a default route. Besides this, IGMP and PIM is configured on Ge0/0 interface.

```
vpn 10
 router
  pim
   auto-rp
   interface ge0/0
   exit
  !
  igmp
   interface ge0/0
   exit
 !
 interface ge0/0
  ip address 192.168.3.1/24
  no shutdown
```

RP router also has basic underlay configuration with a default route.

> **Note**: It is mandatory to use a non-viptela device as RP. In this example, CSR that runs Cisco IOS® XE software has been used for this purpose.

```
ip multicast-routing distributed
!
interface Loopback0 ip address 192.168.101.1 255.255.255.255 ip pim sparse-mode ! ! interface GigabitEthernet2 ip address 192.168.1.3 255.255.255.0 ip pim sparse-mode ! ! ! ip pim send-rp-announce Loopback0 scope 20 ip pim send-rp-discovery Loopback0 scope 20
```

When Auto-RP is used, these events happen:

1. The RP mapping agent listens on a well-known group address CISCO-RP-ANNOUNCE (224.0.1.39), which candidate RP announcements are sent to. When you use Auto-RP to distribute group-to-RP mappings, the **ip pim send-rp-announce** command causes the router to send an Auto-RP announcement message to the well-known group CISCO-RP-ANNOUNCE (224.0.1.39).
2. The RP mapping agent sends group-to-RP mappings in an Auto-RP discovery message to the well-known group CISCO-RP-DISCOVERY (224.0.1.40). The TTL value limits how many

hops the message can take.
3. PIM routers listen to this group and use the RPs they learn about from the discovery message.

Source router is a CSR that runs Cisco IOS® -XE software, which also has basic underlay configuration with a default route. Traffic is generated with the help of a **ping** command to the multicast address.

```
ip multicast-routing distributed
!
interface GigabitEthernet5 ip address 192.168.100.2 255.255.255.0 ip pim sparse-mode
```

Receiver is a CSR that runs Cisco IOS® -XE software as well and has been configured as an IGMP receiver with the help of **ip igmp join-group** command. It also has default route and basic underlay configuration.

```
ip multicast-routing distributed
!
interface GigabitEthernet2
 ip address 192.168.3.2 255.255.255.0
 ip igmp join-group 239.1.2.3
```

Verify

You can use this section in order to confirm that your configuration works properly.

Step 1. The receiver sends IGMP join message to the RP. **debug ip igmp 239.1.2.3** output from the receiver.

```
Oct  9 12:29:12.707: IGMP(0): v2 querier for GigabitEthernet2 is this system.
Oct  9 12:29:12.708: IGMP(0): Send v2 init  Query on GigabitEthernet2
Oct  9 12:29:12.708: IGMP(0): Set report delay time to 0.9 seconds for 239.1.2.3 on
GigabitEthernet2
Oct  9 12:29:13.669: IGMP(0): Send v2 Report for 239.1.2.3 on GigabitEthernet2
Oct  9 12:29:13.669: IGMP(0): Received v2 Report on GigabitEthernet2 from 192.168.3.2 for
239.1.2.3  <<<<<<<<<<<
Oct  9 12:29:13.670: IGMP(0): Received Group record for group 239.1.2.3, mode 2 from 192.168.3.2
for 0 sources
Oct  9 12:29:13.670: IGMP(0): Updating EXCLUDE group timer for 239.1.2.3
Oct  9 12:29:13.670: IGMP(0): MRT Add/Update GigabitEthernet2 for (*,239.1.2.3) by 0
Oct  9 12:29:17.377: IGMP(0): Received v2 Query on GigabitEthernet2 from 192.168.3.1
```

Step 2. BR3-VE-1 which acts as LHR. It receives IGMP join message, sends this information to RP. These IGMP join messages are carried as part of multicast routes in OMP updates.

```
BR3-VE-1# show igmp groups

                     V1
     IF              MEMBERS                                          V1
VPN  NAME   GROUP    PRESENT  STATE          UPTIME      EXPIRES     EXPIRES  EVENT
-----------------------------------------------------------------------------------------------
```

```
-----
10   ge0/0  239.1.2.3  false    members-present  1:11:00:11  0:00:02:41  -        membership-
report
```

Step 3. vSmart receives an (*,G) entry via OMP and forwards this information to the replicator.

```
vsmart# show omp multicast-routes
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Stg -> staged
Inv -> invalid


ADDRESS SOURCE
FAMILY   TYPE   VPN  ORIGINATOR  DESTINATION    GROUP   SOURCE   FROM PEER     RP
STATUS
-----------------------------------------------------------------------------------------------
-----
ipv4    (*,G)   10   10.33.33.3   10.11.11.1  239.1.2.3 0.0.0.0  10.33.33.3 192.168.101.1  C,R
```

Step 4. In this topology, BR1-VE-1 acts as a replicator. BR1-VE-1 forwards this information to the RP.

```
BR1-VE-1# show omp multicast-routes
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Stg -> staged
Inv -> invalid


ADDRESS SOURCE FROM
FAMILY  TYPE  VPN  ORIGINATOR  DESTINATION   GROUP     SOURCE      PEER      RP          STATUS
-----------------------------------------------------------------------------------------------
--
ipv4    (*,G) 10   10.33.33.3   10.11.11.1  239.1.2.3  0.0.0.0   10.1.1.2 192.168.101.1   C,I,R
```

Step 5. The RP now has an (*,G) entry created.

```
FHR-RP#show 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,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.2.3), 1d12h/00:02:51, RP 192.168.101.1, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse, 1d12h/00:02:51
```

Step 6. Now, it is the turn of the source to register with the RP. In this example, multicast traffic is generated with the use of the **ping** command with multicast address as a destination.

```
Source#ping 239.1.2.3 repeat 10
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 239.1.2.3, timeout is 2 seconds:

<SNIP>
```

The source sends a register message to the RP.

```
FHR-RP#show 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,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.2.3), 00:00:12/00:03:27, RP 192.168.101.1, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse, 00:00:02/00:03:27

(192.168.100.2, 239.1.2.3), 00:00:12/00:02:47, flags: T
  Incoming interface: GigabitEthernet4, RPF nbr 192.168.100.2
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse, 00:00:02/00:03:29
<SNIP>
```

Step 7. BR1-VE-1 forwards PIM (S, G) join message to the vSmart. Like an IGMP join, PIM (S, G)

join messages are carried as part of multicast routers in OMP updates. vSmart now has (S, G) entry created in the MRIB. (S, G) information is then forwarded to the replicator as well as to LHR via OMP.

> **Note**: In a real scenario, the replicator can be at the same site or at a different site depends on your design preferences.

```
vsmart# show omp multicast-routes
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Stg -> staged
Inv -> invalid


ADDRESS SOURCE
FAMILY TYPE   VPN  ORIGINATOR  DESTINATION  GROUP      SOURCE        FROM PEER      RP
STATUS
-----------------------------------------------------------------------------------------------
-----------
ipv4   (*,G)  10   10.33.33.3  10.11.11.1  239.1.2.3   0.0.0.0       10.33.33.3  192.168.101.1
C,R
       (S,G)  10   10.33.33.3  10.11.11.1  239.1.2.3   192.168.100.2 10.33.33.3      -
C,R


BR1-VE-1# show omp multicast-routes
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Stg -> staged
Inv -> invalid


ADDRESS SOURCE FROM
FAMILY TYPE    VPN   ORIGINATOR  DESTINATION  GROUP      SOURCE          PEER       RP
STATUS
-----------------------------------------------------------------------------------------------
--------
ipv4    (*,G)  10    10.33.33.3  10.11.11.1   239.1.2.3  0.0.0.0         10.1.1.2  192.168.101.1
C,I,R
        (S,G)  10    10.33.33.3  10.11.11.1   239.1.2.3  192.168.100.2   10.1.1.2      -
C,I,R
```

Step 8. Last hop router now has (S, G) entry. LHR now sends an (S, G) join to a source.

Note: Here in the output you can see that for both (*, G) entry and (S, G) entry originator is shown as 10.33.33.3 and destination is 10.11.11.1 for the group. This is because LHR BR3-VE-1 is responsible for creating (*, G) entry as well as for (S, G) join to build the multicast control plane.

```
BR3-VE-1# show omp multicast-routes
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Stg -> staged
Inv -> invalid


ADDRESS SOURCE FROM
FAMILY TYPE   VPN  ORIGINATOR   DESTINATION   GROUP     SOURCE        PEER     RP
STATUS
-----------------------------------------------------------------------------------------------
---------
ipv4   (*,G)  10   10.33.33.3   10.11.11.1   239.1.2.3 0.0.0.0       0.0.0.0 192.168.101.1
C,Red,R
       (S,G)  10   10.33.33.3   10.11.11.1   239.1.2.3 192.168.100.2 0.0.0.0     -
C,Red,R
```

Data Plane Verification:

Ideal traffic flow must be (from, to):

1. Source to the FHR-RP
2. FHR-RP to the VE
3. VE to the replicator
4. Replicator to the LHR
5. LHR to the receiver

Note: This document does not cover details of PIM RPT and SPT switchover.

In this example, traffic flow is as:

1. From the source to the FHR-RP
2. FHR-RP to BR1-VE-1
3. BR1-VE-1 to BR3-VE-1 via IPSec data plane tunnel
4. BR3-VE-1 to the Receiver

Note: Multicast traffic flows between BR1-VE-1 and BR3-VE-1 via data plane IPsec tunnel. vSmart controller never participates in the actual traffic forwarding.

In this topology, BR1-VE-1 is configured as a replicator and located close to the source. There can be scenarios when replicators are located at a different site from the source. In any case, ensure

data plane tunnels are up between particular site and site where replicator resides.

```
BR1-VE-1# show multicast topology
Flags:
  S: SPT switchover
OIF-Flags:
  A: Assert winner

                                JOIN                                          UPSTREAM     UPSTREAM
UPSTREAM                                      OIF   OIF
VPN  GROUP       SOURCE         TYPE     FLAGS  RP ADDRESS     REPLICATOR  NEIGHBOR     STATE
INTERFACE  UP TIME     EXPIRES     INDEX  NAME  FLAGS  OIF TUNNEL
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------
10   224.0.1.39  192.168.101.1  Auto-RP  -      -              -           192.168.1.3  joined
ge0/0      0:00:41:29  0:00:02:33  513    -     -      10.33.33.3
10   224.0.1.40  192.168.101.1  Auto-RP  -      -              -           192.168.1.3  joined
ge0/0      0:00:41:26  0:00:02:17  513    -     -      10.33.33.3
10   239.1.2.3   0.0.0.0        (*,G)    -      192.168.101.1  -           192.168.1.3  joined
ge0/0      0:00:03:47  0:00:00:53  513    -     -      10.33.33.3
10   239.1.2.3   192.168.100.2  (S,G)    -      -              -           192.168.1.3  joined
ge0/0      0:00:00:10  0:00:00:52  513    -     -      10.33.33.3


BR1-VE-1# show bfd sessions system-ip 10.33.33.3
                                     SOURCE TLOC      REMOTE TLOC
DST PUBLIC                      DST PUBLIC         DETECT      TX
SYSTEM IP        SITE ID  STATE       COLOR            COLOR            SOURCE IP
IP                              PORT        ENCAP  MULTIPLIER  INTERVAL(msec) UPTIME
TRANSITIONS
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------
-------------
10.33.33.3       30       up          gold             gold             172.16.1.6
172.16.1.14                     12406       ipsec  7           1000           3:21:24:02      0
10.33.33.3       30       up          gold             lte              172.16.1.6
172.19.1.6                      12426       ipsec  7           1000           3:21:24:02      0
10.33.33.3       30       up          biz-internet     gold             172.17.1.6
172.16.1.14                     12406       ipsec  7           1000           3:21:24:59      0
10.33.33.3       30       up          biz-internet     lte              172.17.1.6
172.19.1.6                      12426       ipsec  7           1000           3:21:24:59      0


BR1-VE-1# show multicast topology vpn 10 239.1.2.3 topology-oil
Flags:
  S: SPT switchover
OIF-Flags:
  A: Assert winner

                              JOIN          OIF   OIF
VPN  GROUP      SOURCE         TYPE   INDEX  NAME  FLAGS  OIF TUNNEL
-------------------------------------------------------------------
10   239.1.2.3  0.0.0.0        (*,G)  513    -     -      10.33.33.3
10   239.1.2.3  192.168.100.2  (S,G)  513    -     -      10.33.33.3


BR3-VE-1# show bfd sessions system-ip 10.11.11.1
                                     SOURCE TLOC      REMOTE TLOC
DST PUBLIC                      DST PUBLIC         DETECT      TX
SYSTEM IP        SITE ID  STATE       COLOR            COLOR            SOURCE IP
IP                              PORT        ENCAP  MULTIPLIER  INTERVAL(msec) UPTIME
```

```
TRANSITIONS
-----------------------------------------------------------------------------------------------
-----------------------------------------------------------------------------------------------
-------------
10.11.11.1       10       up          gold             gold             172.16.1.14
172.16.1.6                        12406       ipsec  7          1000           3:21:25:16      0
10.11.11.1       10       up          gold             biz-internet     172.16.1.14
172.17.1.6                        12406       ipsec  7          1000           3:21:26:13      0
10.11.11.1       10       up          lte              gold             172.19.1.6
172.16.1.6                        12406       ipsec  7          1000           3:21:25:16      0
10.11.11.1       10       up          lte              biz-internet     172.19.1.6
172.17.1.6                        12406       ipsec  7          1000           3:21:26:13      0
```

Step 9. The receiver is now getting traffic.

```
Receiver#show 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,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.2.3), 1d13h/stopped, RP 192.168.101.1, flags: SJPCL
  Incoming interface: GigabitEthernet2, RPF nbr 192.168.3.1
  Outgoing interface list: Null

(192.168.100.2, 239.1.2.3), 00:01:08/00:01:51, flags: PLTX
  Incoming interface: GigabitEthernet2, RPF nbr 192.168.3.1
  Outgoing interface list: Null



Receiver#show ip mroute count
Use "show ip mfib count" to get better response time for a large number of mroutes.

IP Multicast Statistics
6 routes using 3668 bytes of memory
3 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 239.1.2.3, Source count: 1, Packets forwarded: 0, Packets received: 16
  RP-tree: Forwarding: 0/0/0/0, Other: 7/0/7
  Source: 192.168.100.2/32, Forwarding: 0/0/0/0, Other: 9/0/9



Source#ping 239.1.2.3 repeat 10
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 239.1.2.3, timeout is 2 seconds:
```

```
Reply to request 0 from 192.168.3.2, 221 ms
Reply to request 1 from 192.168.3.2, 238 ms
Reply to request 2 from 192.168.3.2, 135 ms
Reply to request 3 from 192.168.3.2, 229 ms
Reply to request 4 from 192.168.3.2, 327 ms
Reply to request 5 from 192.168.3.2, 530 ms
<SNIP>
```

Troubleshoot

This section provides information you can use in order to troubleshoot your configuration.

1. Verify that (*, G) and (S,G) are present on the RP.

2. Ensure that you have data plane tunnels and BFD sessions are up between VE and site where replicator configured with the help of **show bfd sessions** command.

3. Check that BR3-VE-1 learned about replicator on BR1-VE-1.

```
BR3-VE-1# show multicast replicator

     REPLICATOR  REPLICATOR  LOAD
VPN  ADDRESS     STATUS      PERCENT
-------------------------------------
10   10.11.11.1  UP          -
```

4. Ensure a multicast tunnel is established with BR3-VE-1.

```
BR3-VE-1# show multicast tunnel

     TUNNEL     TUNNEL
VPN  ADDRESS    STATUS  REPLICATOR
-----------------------------------
10   10.11.11.1   UP      yes
```

5. Ensure that the group-to-RP mapping is distributed and correct.

```
BR3-VE-1#show pim rp-mapping

VPN  TYPE     GROUP        RP ADDRESS
----------------------------------------
10   Auto-RP  224.0.0.0/4  192.168.101.1
```

6. Ensure that multicast routes (*, G) and (S, G) are propagated correctly to the vEdge, the Replicator router, and the vSmart. Use **show multicast topology** and **show omp multicast-routes** commands.

7. Check for RPF table on LHR.

```
BR3-VE-1#  show multicast rpf | tab

                                                    RPF                              RPF
                    RPF       NEXTHOP         RPF NBR     IF                RPF TUNNEL    TUNNEL
VPN  RPF ADDRESS    STATUS    COUNT    INDEX  ADDR        NAME  RPF TUNNEL  COLOR         ENCAP
-------------------------------------------------------------------------------------------------
--
```

```
10   192.168.101.1  resolved  2        0      10.11.11.1  -     10.11.11.1  biz-internet  ipsec
                                       1      10.11.11.1  -     10.11.11.1  gold          ipsec
10   192.168.100.2  resolved  2        0      10.11.11.1  -     10.11.11.1  biz-internet  ipsec
                                       1      10.11.11.1  -     10.11.11.1  gold          ipsec
```

8. Check that LHR learned all required information about Auto-RP and data multicast groups with help of **show ip mfib summary** command.

9. Check that **show ip mfib oil** command output on the LHR contains egress interface pointing to the Receiver router.

10. Check that traffic flows with help of the **show ip mfib stats** command.

Other useful debug commands:

- **debug pim auto-rp level high** – Enables auto-rp debug.
- **debug pim events level high vpn <vpn number>** – Enables PIM events debug.
- **debug ftm mcast** – Enables multicast programming debug.

Conclusion

These scenarios have been tested successfully in this topology.

- The multicast source is connected directly to the RP at the same site and the receiver is located at the remote site (test scenario).
- The multicast receiver is connected directly to RP at the same site, while source is at a remote site.
- The multicast source is connected directly to the VE, while receiver and RP are at the remote site.