Catalyst 9000交换机上的EVPN VxLAN TRM故障排除

目录

简介

本文档介绍如何对EVPN VxLAN上的TRM（租户路由组播）问题进行故障排除。

先决条件

- 建议您熟悉单播EVPN VxLAN功能、BGP和MVPN（组播虚拟专用网络）。
- 此外，您必须了解组播的运行方式和组播概念

要求

本指南假设BGP、NVE对等体已经正确。如果出现基本EVPN VxLAN问题（单播ping失败、BGP、NVE对等体关闭等），请根据需要参考BGP、EVPN、路由/交换机故障排除指南。

每个代码版本中的功能可用性

版本	功能
17.1.1	具有任播RP的TRMv4
17.3.1	带外部RP或单个RP的TRMv4
17.3.1	具有任播RP的TRMv6
17.3.1	带外部RP或单个RP的TRMv6
17.3.1	TRMv4，具有MVPN互通(profile11)，在交换矩阵端使用单个RP
17.6.2和17.7.1	具有任播RP、外部RP或单个RP的TRMv4数据模块

使用的组件

本文档中的信息基于以下软件和硬件版本：

- C9300
- C9400
- C9500
- C9600

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

注意：请参阅相应的配置指南，了解用于在其他Cisco平台上启用这些功能的命令。

背景信息

要配置EVPN TRM，请参阅：BGP EVPN VXLAN配置指南，Cisco IOS XE Amsterdam 17.3.x

租户路由组播(TRM)是基于BGP-EVPN的解决方案，可在VxLAN交换矩阵[RFC7432]中连接VTEPS的源和接收器之间实现组播路由。TRM依靠单播EVPN中存在的路由来发现组播源和组播RP。与NG-MVPN一样，组播源和接收器信息通过BGP协议在配置了BGP MVPN地址系列的

VTEP之间传播。没有从TRM VTEP向VxLAN交换矩阵发送PIM/IGMP数据包。

TRM解决的关键问题是位于不同VLAN但位于同一VRF中的组播发送器和接收器能够彼此通信。如果没有TRM，则组播流量将作为底层中相同BUM（广播、单播和组播）基础设施的一部分发送，该基础设施可以是组播树或入口复制。此基础设施是按VLAN构建的，因此，虽然同一VLAN上的组播源和接收器可以通信，但位于不同VLAN中的组播源和接收器无法通信。使用TRM时，组播将从BUM中移出，并在父VRF下汇聚在一起。因此，无论源或接收器所在的VLAN如何，组播通信都完全启用。

TRM在本地或跨VTEP的相同或不同子网内的发送方和接收方之间提供多租户感知组播转发。请参阅指南 BGP EVPN VXLAN配置指南，Cisco IOS XE Amsterdam 17.3.x 了解更多详情

如何在本指南中介绍自己：

- 本指南根据RP位置分为4个场景。
- 场景可以参考您所在部分未直接介绍的CLI示例。例如，SSM场景2引用场景1以了解如何读取某些CLI。
- 只有场景1涵盖IPv4和IPv6，因为这两个地址系列的概念在本质上是相同的。
- 这些方案中列出的要求假定源和接收器直接连接到VTEP(有关详细信息，请参阅相关信息部分“交换矩阵外的源和接收器”)。

场景	IPv4/v6	每个产品中都包含哪些内容
所有其他场景的常见详细信息	IPv4	始终要求：MVPN底层和NVE正常，RPF检查到任何TRM源都是L3VNI。
AnyCast RP（每个VTEP都是一个具有通用RP IP的RP）	IPv4/v6	BGP、PIM、IGMP、MFIB和FED命令对于IPv4/v6以及捕获示例都非常详细
无RP（SSM重叠）	IPv4	SSM特定信息。（有关常见信息，请参阅场景1）
交换矩阵内部的RP（交换矩阵的一个通用RP）	IPv4	IPv4的BGP、PIM、IGMP、MFIB和FED命令的全面详细信息
交换矩阵外部的RP（RP不在交换矩阵内）	IPv4	IP到交换矩阵边界特定信息。（有关常见信息，请参阅场景3）
交换矩阵内部的RP（交换矩阵的一个通用RP），具有对称L2VNI	IPv4	当VNI同时位于发送方和接收方VTEP上时，在交换矩阵中使用单个RP的注意事项。（有关常见信息，请参阅场景3）

在本故障排除文档中，在show命令输出的某些行末尾添加了注释。这样做是为了突出或解释输出行的特定方面。如果注释以新行开始，则它引用注释前面的输出行。在本文档中，此记法用于突出显示show命令输出中的注释：

<#root>

```
<-- Text highlighted in this format inside a command's output represents a comment.

    This is done for explanation purpose only and is not part of the command's output.
```

术语

EVPN	以太网虚拟专用网络	允许BGP传输第2层MAC和第3层IP信息的扩展是EVPN，它使用多协议边界网关协议(MP-BGP)作为协议来分发属于VXLAN重叠网络的可达性信息。
VXLAN	虚拟可扩展LAN（局域网）	VXLAN旨在克服VLAN和STP的固有局限性。推荐的IETF标准[RFC 7348]提供与VLAN相同的以太网第2层网络服务，但灵活性更高。功能上，它是MAC-in-UDP封装协议，在第3层底层网络上作为虚拟重叠运行。
VTEP	虚拟隧道终端	这是执行封装和解封的设备
NVE	网络虚拟接口	进行封装和解封的逻辑接口
VNI	VXLAN网络标识符	唯一标识每个第2层子网或网段。VNI有两种类型： 对称(L2VNI):VTEP具有相同的VNI 非对称(L3VNI):VTEP没有相同的VNI并通过单个公共VNI路由。
MDT	组播分布树	在VTEP之间构建的组播树，用于租户组播流量的封装和隧道。
BUM	广播、未知单播、组播	BUM流量通过与NVE配置下的VNI关联的Mcast组发送。

RP	汇集点	设备在PIM稀疏模式下执行的角色。组播源和接收器的公共会议点。
AnyCast(RP)	AnyCast集合点	两个或多个RP在环回接口上配置了相同的IP地址。FHR根据单播路由注册到最近的RP。
RPT（树）	根路径树	也称为共享或*,G树。此路径指向RP
SPT（树）	最短路径树	到源的最短路径，由单播路由表确定
FHR	第一跳路由器	与源直接连接的设备（ARP邻接）。FHR向RP注册源信息。
LHR	最后一跳路由器	连接接收器的设备
RPF	反向路径转发	返回到源的单播路径。 除非收到的组播数据包与单播路由表路径相同，否则不会接受/转发传入组播数据包。（“ip multicast multipath”使用案例已排除）。
MRIB	组播路由信息库	软件组播路由表，也称为mroute表
MFIB	组播转发信息库	组播等效的CEF。由MRIB的更新填充，用于由数据平面转发。
美联储	转发引擎驱动程序	对设备硬件进行编程的组件。
IIF	传入接口	启用PIM的接口，也是返回源的单播RPF上游路径。（在show ip mroute中显示）
OIF	传出接口	下游的启用PIM的接口。（在show ip mroute中显示）

验证

所有方案通用的验证

本节介绍任何场景所需的基本要求。

- 确保所需的NVE对等体已启动
- 确保租户VRF中指向源的RPF接口是L3VNI SVI。如果RPF接口不是L3VNI SVI，则BGP不会发送类型7加入路由。在任何情况下，RPF接口都必须指向此接口。
- 确保对等体之间的底层路径（MDT隧道）已完成。
- 确保BGP用于组播控制平面（使用MVPN对比PIM）

注意：此部分适用于IPv4和IPv6租户组播验证。

验证NVE对等

对于本指南中的任何场景，检查以确保NVE对等体在VTEP之间运行

- NVE对等体由从BGP获取的地址形成。

<#root>

```
Leaf-01#
sh nve peers

Interface VNI    Type Peer-IP       RMAC/Num_RTs   eVNI    state flags UP time
nve1      50901  L3CP 172.16.254.4  7c21.0dbd.9548 50901   UP    A/-/4 01:54:11 <-- IPv4 peering

with Leaf 02

nve1      50901  L3CP 172.16.254.4  7c21.0dbd.9548 50901   UP    A/M/6 17:48:36 <-- IPv6 peering with Le

Leaf-02#
sh nve peers

Interface VNI    Type Peer-IP       RMAC/Num_RTs   eVNI    state flags UP time
nve1      50901  L3CP 172.16.254.3  10b3.d56a.8fc8 50901   UP    A/-/4 01:55:44 <-- IPv4 peering with Le

nve1      50901  L3CP 172.16.254.3  10b3.d56a.8fc8 50901   UP    A/M/6 17:56:19 <-- IPv6 peering with Le
```

验证租户VRF中的RPF接口

如果此接口是除L3VNI SVI以外的任何接口，则BGP不会发起MVPN类型7连接。

- 如果您没有看到此接口，请确认使路由返回源接口的配置不存在问题，该接口不是L3VNI。

<#root>

```
Leaf-03#
sh ip rpf vrf green 10.1.101.11 <-- Multicast source IP

RPF information for ? (10.1.101.11)

 RPF interface: Vlan901                <-- RPF interface is the L3VNI SVI


RPF neighbor: ? (172.16.254.3)        <-- Underlay Next hop IP


RPF route/mask: 10.1.101.0/24         <-- Network prefix for the Source

  RPF type: unicast (bgp 65001)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

验证组播控制平面使用BGP

- mdt overlay use-bgp：通知设备使用BGP MVPN类型5/6/7作为信号协议（相对于PIM消息）
- spt-only：其他关键字通知设备仅使用AnyCast RP场景中的SPT树。由于每个VTEP都是RP，因此不使用MVPN Type-6路由。

<#root>

```
Leaf-01
!
vrf definition green
rd 1:1
!
address-family ipv4
mdt auto-discovery vxlan
mdt default vxlan 239.1.1.1       <-- Defines MDT default underlay group address


mdt overlay use-bgp [spt-only]    <-- Required for VTEP to use MVPN Type 5/6/7 versus PIM for multicast
```

验证MDT组

MDT组对所有方案都通用，因为这是封装TRM组的外部隧道组。

检查MDT组是否已在源端正确编程

- MDT组的传入接口是源端环回

- MDT组的传出接口是底层接口

验证Leaf-01:MRIB/MFIB中的MDT mroute正确

<#root>

Leaf-01#

sh ip mroute 239.1.1.1 172.16.254.3

(

172.16.254.3

,

239.1.1.1

), 00:46:35/00:02:05, flags: FTx
  Incoming interface:

Loopback1

, RPF nbr

0.0.0.0

<-- IIF is local loopback with 0.0.0.0 RPF indicating local

  Outgoing interface list:

GigabitEthernet1/0/2

, Forward/Sparse, 00:46:35/00:03:12

<-- OIF is the underlay uplink

Leaf-01#

sh ip mfib 239.1.1.1 172.16.254.3

(172.16.254.3,239.1.1.1) Flags: HW

  SW Forwarding: 2/0/150/0, Other: 1/1/0

HW Forwarding: 1458/0/156/0

, Other: 0/0/0

<-- Hardware counters indicate the entry is operating in hardware and forwarding packets

Null0 Flags: A NS **<--- Null0 (originated locally)**

GigabitEthernet1/0/2

Flags: F NS

<-- OIF is into the Underlay (Global route table)

Pkts: 0/0/1 Rate: 0 pps

验证MDT组的Leaf-01:FED条目

<#root>

Leaf-01#

sh platform software fed switch active ip mfib 239.1.1.1/32 172.16.254.3 detail <-- the detail option gi

MROUTE ENTRY

vrf 0

(

172.16.254.3, 239.1.1.1/32

)

<-- vrf 0 = global for this MDT S,G pair

HW Handle: 139738317079128 Flags:

RPF interface: Null0

(1)):

<-- Leaf-01 the Source (Null0)

HW Handle:139738317079128 Flags:A
Number of OIF: 2
Flags: 0x4

Pkts : 71 <-- packets that used this adjacency (similar to mfib command, but shown at the FED

OIF Details:

Null0 A

<-- The incoming interface is Local Loopback1 and A-Accept flag set

GigabitEthernet1/0/2

```
F

 NS

<-- The Underlay Outgoing Interface and F-Forward flag set


Htm: 0x7f175cc0beb8 Si: 0x7f175cc0a6b8

Di: 0x7f175cc09df8

 Rep_ri: 0x7f175cc0a1d8

<-- The DI (dest index) handle



DI details


----------
Handle:0x7f175cc09df8 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICA
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles:

index0:0x538d

 mtu_index/l3u_ri_index0:0x0

index1:0x538d

 mtu_index/l3u_ri_index1:0x0

Brief Resource Information (ASIC_INSTANCE# 1)
---------------------------------------

Destination index = 0x538d


pmap = 0x00000000 0x00000002

pmap_intf : [GigabitEthernet1/0/2]  <-- FED has the correct programming for the OIF


================================================================
```

检查MDT组是否已在接收器端正确编程

- MDT组的传入接口是返回源端环回的RPF接口
- MDT组的传出接口是Encap/Decap隧道接口

验证Leaf-02:MRIB/MFIB中的MDT mroute正确

```
<#root>

Leaf-02#

sh ip mroute 172.16.254.3 239.1.1.1         <-- This is the Global MDT group
```

(

172.16.254.3

,

239.1.1.1

), 00:23:35/00:01:09, flags: JTx

<-- Source is Leaf-01 Lo1 IP

  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:

Tunnel0

, Forward/Sparse, 00:23:35/00:00:24

<-- Decap Tunnel

Leaf-02#

sh ip mfib 239.1.1.1 172.16.254.3

Default **<-- Global routing table**

  (172.16.254.3,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 0/0/0

HW Forwarding: 5537/0/168/0, Other: 0/0/0 <-- Hardware counters indicate the entry is operating in hard

GigabitEthernet1/0/2 Flags: A **<-- Accept via Underlay (Global) interface**

Tunnel0, VXLAN Decap Flags: F NS **<-- Forward to VxLAN decap Tunnel**

      Pkts: 0/0/1 Rate: 0 pps

验证MDT组的枝叶02:FED条目

<#root>

Leaf-02#

sh platform software fed switch active ip mfib 239.1.1.1/32 172.16.254.3 detail

MROUTE ENTRY

 vrf 0

 (

172.16.254.3, 239.1.1.1/32

)

<-- vrf 0 = global for this MDT S,G pair

  HW Handle: 140397391831832 Flags:

RPF interface: GigabitEthernet1/0/2

(57)):

<-- RPF interface to 172.16.254.3

  HW Handle:140397391831832 Flags:A
Number of OIF: 2
Flags: 0x4

Pkts : 1585 <-- packets that used this adjacency (similar to mfib command, but shown at the FE

OIF Details:

Tunnel0 F NS <-- Send to decap tunnel to remove VxLAN header

(Adj: 0x73) <-- Tunnel0 Adjacency

GigabitEthernet1/0/2 A <-- Accept MDT packets from this interface

Htm: 0x7fb0d0f1f388 Si: 0x7fb0d0f1dc08 Di: 0x7fb0d0ed0438 Rep_ri: 0x7fb0d0ed07a8

RI details <-- Rewrite Index is used for VxLAN decapsulation

Handle:0x7fb0d0ed07a8 Res-Type:ASIC_RSC_RI_REP Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MUL
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x38 mtu_index/l3u_ri_index0:0x0 index1:0

Brief Resource Information (ASIC_INSTANCE# 0)

ASIC# 0
Replication list :

Total #ri : 6
Start_ri : 56
Common_ret : 0

Replication entry

rep_ri 0x38

 #elem = 1
0)

ri[0]=0xE803

 Dynamic port=88ri_ref_count:1 dirty=0

```
Leaf-02#
sh platform hardware fed sw active fwd-asic resource asic all rewrite-index range 0xE803 0xE803

ASIC#:0 RI:59395
Rewrite_type:
AL_RRM_REWRITE_L2_PAYLOAD_
IPV4_EVPN_DECAP
(118) Mapped_rii:LVX_EVPN_DECAP(246)
<...snip...>
```

场景 1.AnyCast RP（仅SPT树）IPv4和IPv6

在此模式中，每个VTEP上都有一个RP。这些VTEP无法通过MSDP同步已了解的源，并且没有共享树。相反，MDT模式使用BGP信息仅创建SPT组播树。 此模式可互换称为仅SPT模式或分布式任意播RP模式。在此模式下，每个VTEP都是PIM RP。因此，每个站点的(*,G)树在本地VTEP自身被截断。无需跨交换矩阵发送(*,G)加入或MVPN RT-6。

网络图

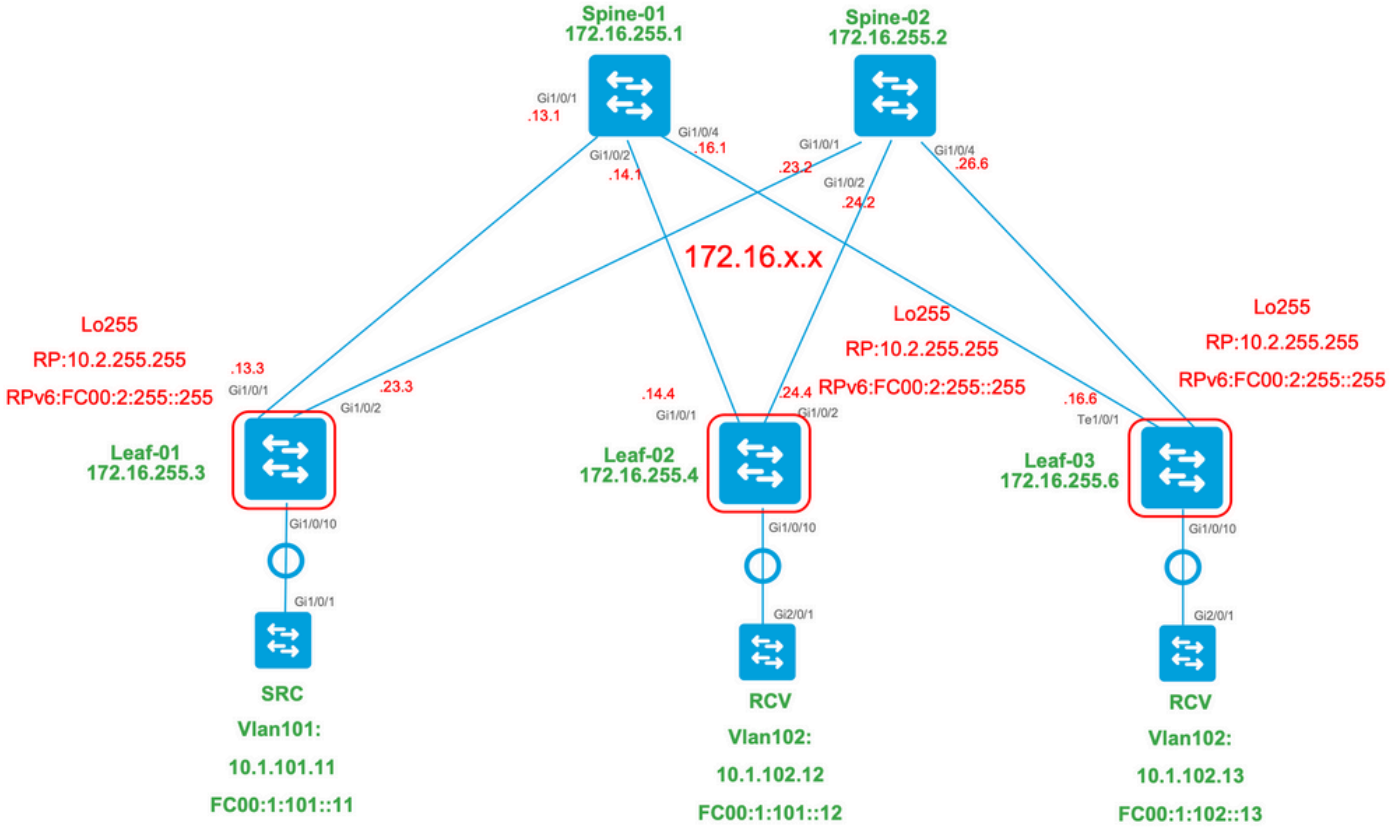

对于此模式，请考虑以下3种BGP路由类型：

1. EVPN路由类型2。这允许需要将C组播路由（MVPN类型6/7）构建回源PE的其他PE附加适当的C组播导入RT，以便发起方PE可以导入C组播路由(RFC 6514 11.1.3)[RFC6514]。 此

VRI的使用取决于命令“mdt overlay use-bgp” VRF命令。
2. MVPN路由类型5。这与MVPN中相同，并且是可用的组播源/组的通告
3. MVPN路由类型7。来自IGMP或MLD层以及来自EVPN类型2的信息用于创建此BGP类型联接。类型7驱动在源端创建MRIB OIF。

EVPN第2类要求：

1. 直连组播源已联机。
2. FHR（源VTEP）验证ARP（或ND）和CEF邻接（确认源为直连）。
3. FHR发出EVPN第2类BGP更新

MVPN类型5要求：

1. 源直接连接的要求已解决
2. RP是本地的，因此FHR向自身注册
3. FHR发起MVPN第5类BGP更新

MVPN类型7要求：

1. 存在EVPN类型2条目（使用正确的VRI构建C-Multicast路由类型7时需要该条目，并且是从源VTEP发送的）
2. 存在MVPN第5类条目（解析可用于SPT加入的源/组对时需要该条目）
3. LHR VTEP已接收并处理IGMP或MLD成员身份报告
4. LHR VTEP RPF接口是交换矩阵L3VNI接口


提示：在出口LHR VTEP PIM检查通往源的路径。PIM必须在RIB中找到作为RPF接口的L3VNI路由。如果L3VNI未正确配置、已关闭，依此类推。VTEP不会尝试创建类型7 BGP连接。

检验BGP EVPN和MVPN路由

验证Leaf-01:已创建EVPN类型2

<#root>

```
### IPv4 ###


Leaf-01#
sh bgp l2vpn evpn all route-type 2 0 F4CFE24334C5 10.1.101.11
...or you can also use:


Leaf-01#
sh bgp l2vpn evpn detail [2][172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24


BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24, version 6
Paths: (1 available, best #1,
```

table evi_101

)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local

:: (via default) from 0.0.0.0 (172.16.255.3) <-- Leaf-01 locally created

      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 MVPN AS:65001:0.0.0.0

        MVPN VRF:172.16.255.3:2

 ENCAP:8 Router MAC:10B3.D56A.8FC8

 <-- MVPN VRI RT is part of the EVPN Type-2

      Local irb vxlan vtep:

        vrf:green, l3-vni:50901 <-- Vrf and VxLAN tag

        local router mac:10B3.D56A.8FC8

        core-irb interface:Vlan901 <-- L3VNI SVI

        vtep-ip:172.16.254.3 <-- Leaf-01 VTEP

      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 17:40:29 UTC

IPv6

Leaf-01#

sh bgp l2vpn evpn all route-type 2 0 F4CFE24334C1 FC00:1:101::11

...or you can also use:
Leaf-01#

sh bgp l2vpn evpn detail [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36, version
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local

:: (via default) from 0.0.0.0 (172.16.255.3) <-- Leaf-01 locally created

      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 MVPN AS:65001:0.0.0.0

MVPN VRF:172.16.255.3:2

ENCAP:8 Router MAC:10B3.D56A.8FC8

<-- MVPN VRI RT is part of the EVPN Type-2

Local irb vxlan vtep:

vrf:green, l3-vni:50901

local router mac:10B3.D56A.8FC8

core-irb interface:Vlan901 <-- L3VNI SVI

vtep-ip:172.16.254.3 <-- Leaf-01 VTEP

rx pathid: 0, tx pathid: 0x0
Updated on Mar 22 2021 19:54:18 UTC

检验Leaf-01:ARP/IPv6 ND和EVPN调试显示ARP/ND已获知，然后创建并发送路由类型2

<#root>

IPv4

Leaf-01#

sh debugging

ARP:

ARP packet debugging is on

BGP L2VPN EVPN:

BGP updates debugging is on for address family: L2VPN E-VPN
BGP update events debugging is on for address family: L2VPN E-VPN

*Dec 17 17:00:06.480:

IP ARP: rcvd rep src 10.1.101.11 f4cf.e243.34c5

, dst 10.1.101.11 Vlan101

tableid 2 <-- Multicast Source ARP

*Dec 17 17:00:06.481:

BGP: EVPN Rcvd pfx: [2]

[172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24, net flags: 0

<-- BGP Triggered Type-2 creation

*Dec 17 17:00:06.481:

TRM communities added to sourced RT2 <-- TRM extended VRI communities being injected into EVPN Type-2

*Dec 17 17:00:06.481:

BGP(10): update modified for [2]

[172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/30

<-- Modifying the update

*Dec 17 17:00:06.481: BGP(10): 172.16.255.1 NEXT_HOP set to vxlan local vtep-ip 172.16.254.3 for net [2
*Dec 17 17:00:06.481: BGP(10): update modified for [2][172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.
*Dec 17 17:00:06.481: BGP(10): (base) 172.16.255.1

send UPDATE

(format)

[2]

[172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/30, next 172.16.254.3, metric 0, path Local, e

MVPN VRF:172.16.255.3:2

ENCAP:8 Router MAC:10B3.D56A.8FC8

<--- Final update sent to RR with standard EVPN community info and required MVPN community attributes

IPv6

Leaf-01#

debug ipv6 nd

ICMP Neighbor Discovery events debugging is on
ICMP ND HA events debugging is ON

IPv6 ND:

Mar 23 14:29:51.935:

ICMPv6-ND: (Vlan101,FC00:1:101::11) Resolution request

Mar 23 14:29:51.935: ICMPv6-ND: (Vlan101,FC00:1:101::11) DELETE -> INCMP
Mar 23 14:29:51.935: ICMPv6-ND HA: in Update Neighbor Cache: old state 6 new state 0
Mar 23 14:29:51.935: ICMPv6-ND HA: add or delete entry not synced as no peer detected
Mar 23 14:29:51.936: ICMPv6-ND: (Vlan101,FC00:1:101::11) Sending NS
Mar 23 14:29:51.936: ICMPv6-ND: (Vlan101,FC00:1:101::11) Queued data for resolution
Mar 23 14:29:51.953:

ICMPv6-ND: (Vlan101,FC00:1:101::11) Received NA from FC00:1:101::11

Mar 23 14:29:51.953:

ICMPv6-ND: Validating ND packet options: valid

Mar 23 14:29:51.953:

ICMPv6-ND: (Vlan101,FC00:1:101::11) LLA f4cf.e243.34c1

Mar 23 14:29:51.953: ICMPv6-ND HA: modify entry not synced as no peer detected
Mar 23 14:29:51.953:

ICMPv6-ND: (Vlan101,FC00:1:101::11) INCMP -> REACH <-- peer is reachable

Leaf-01#

debug bgp l2vpn evpn updates

Leaf-01#

debug bgp l2vpn evpn updates events

BGP L2VPN EVPN:

Mar 23 14:11:56.462:

BGP: EVPN Rcvd pfx: [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36,

net flags: 0

<-- BGP Triggered Type-2 creation

Mar 23 14:11:57.462:

TRM communities added to sourced RT2

ar 23 14:11:57.474:

BGP(10): update modified for [2]

[172.16.254.3:101][0][48][F4CFE24334C1][128]

[FC00:1:101::11]/42

Mar 23 14:11:57.474: BGP(10): 172.16.255.1 NEXT_HOP set to vxlan local vtep-ip 172.16.254.3 for net [2]
Mar 23 14:11:57.474: BGP(10): update modified for [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:
Mar 23 14:11:57.474: BGP(10): (base) 172.16.255.1

send UPDATE

 (format)

[2]

[172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/42, next 172.16.254.3, metric 0, path Loca

MVPN VRF:172.16.255.3:2

 ENCAP:8 Router MAC:10B3.D56A.8FC8

<--- Final update sent to RR with standard EVPN community info and required MVPN community attributes

验证Leaf-02：在接收端的BGP中获知了源端路由类型2

<#root>

IPv4

Leaf-02#

sh bgp l2vpn evpn all | b 10.1.101.11

 * i

[2]

[172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24

<-- Remote VTEP route-type 2

                     172.16.254.3   0   100   0 ?

 ***>i   172.16.254.3   0   100   0 ?   <-- IP of Leaf01 Lo1**

Leaf-02#

sh bgp l2vpn evpn route-type 2 0 F4CFE24334C5 10.1.101.11

...or you can also use:
Leaf-02#

sh bgp l2vpn evpn detail [2][172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24, version 175
Paths: (2 available, best #2, table

EVPN-BGP-Table) <-- In BGP EVPN table
  Flag: 0x100

  Not advertised to any peer
  Refresh Epoch 2
  Local

    172.16.254.3

 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 10101,

Label2 50901

      Extended Community: RT:1:1 RT:65001:101

MVPN AS:65001:0.0.0.0

        MVPN VRF:172.16.255.3:2

ENCAP:8

Router MAC:10B3.D56A.8FC8

Originator: 172.16.255.3, Cluster list: 172.16.255.2
rx pathid: 0, tx pathid: 0
Updated on Dec 14 2020 19:58:57 UTC

MVPN AS:65001:0.0.0.0 <-- MVPN Autonomous System
MVPN VRF:172.16.255.3:2 <-- VRI Extended Community to be used in MVPN Type-7
Router MAC:10B3.D56A.8FC8 <-- Leaf-01 RMAC
Label2 50901 <-- L3VNI 50901

IPv6

Leaf-02#

sh bgp l2vpn evpn all | b FC00:1:101::11

* i [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
172.16.254.3 0 100 0 ?

***>i 172.16.254.3 0 100 0 ? <-- IP of Leaf01 Lo1**

Leaf-02#

sh bgp l2vpn evpn route-type 2 0 F4CFE24334C1 FC00:1:101::11

...or you can also use:
Leaf-02#

sh bgp l2vpn evpn detail [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36

BGP routing table entry for

[2]

[172.16.254.3:101][0][48][

F4CFE24334C1

][128][

FC00:1:101::11

]/36, version 659
Paths: (2 available, best #2,

table EVPN-BGP-Table

)

<-- In BGP EVPN table

Flag: 0x100
Not advertised to any peer
Refresh Epoch 2
Local

172.16.254.3

```
 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 10101,
```

Label2 50901

```
      Extended Community: RT:1:1 RT:65001:101 MVPN
```

AS:65001:0.0.0.0

MVPN VRF:172.16.255.3:2

```
 ENCAP:8
```

Router MAC:10B3.D56A.8FC8

```
      Originator: 172.16.255.3, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Mar 23 2021 14:11:57 UTC
```

MVPN AS:65001:0.0.0.0 <-- MVPN Autonomous System
MVPN VRF:172.16.255.3:2 <-- VRI Extended Community to be used in MVPN Type-7
Router MAC:10B3.D56A.8FC8 <-- Leaf-01 RMAC
Label2 50901 <-- L3VNI 50901

验证Leaf-02：在接收器VTEP Leaf-02上的BGP中学习了源路由类型5

<#root>

IPv4

Leaf-02#

sh bgp ipv4 mvpn all route-type 5 10.1.101.11 226.1.1.1

...or you can also use:
Leaf-02#

sh bgp ipv4 mvpn detail [5][1:1][10.1.101.11][226.1.1.1]/18

BGP routing table entry for

[5]

[1:1]

[10.1.101.11][226.1.1.1]

/18, version 72

<-- Type-5 contains advertised S,G pair

Paths: (2 available, best #1,

table MVPNv4-BGP-Table

, not advertised to EBGP peer)

<-- In BGP IPv4 MVPN table

  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local

    172.16.255.3

 (metric 3) from 172.16.255.2 (172.16.255.2)

<-- Loopback0 of Leaf-01

      Origin incomplete, metric 0, localpref 100, valid, internal
      Community: no-export
      Extended Community: RT:1:1

Originator: 172.16.255.3

, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Dec 15 2020 16:54:53 UTC

IPv6

Leaf-02#

sh bgp ipv6 mvpn all route-type 5 FC00:1:101::11 FF06:1::1

...or you can also use:
Leaf-02#

sh bgp ipv6 mvpn detail [5][1:1][FC00:1:101::11][FF06:1::1]/42

BGP routing table entry for

[5]

[1:1]

[FC00:1:101::11][FF06:1::1]

/42, version 11

<-- Type-5 contains advertised S,G pair

Paths: (2 available, best #1,

table MVPNV6-BGP-Table

, not advertised to EBGP peer)

<-- In BGP IPv6 MVPN table

  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1

```
  Local

172.16.255.3

 (metric 3) from 172.16.255.2 (172.16.255.2)

<-- Loopback0 of Leaf-01


      Origin incomplete, metric 0, localpref 100, valid, internal
      Community: no-export
      Extended Community: RT:1:1


 Originator: 172.16.255.3

, Cluster list: 172.16.255.2


      rx pathid: 0, tx pathid: 0
      Updated on Mar 23 2021 15:13:06 UTC
```

验证Leaf-02:需要来自Leaf-01的BGP信息以创建类型7。最终要求是IGMP或MLD已处理成员身份报告，该报告通知VTEP有感兴趣的接收方。

```
<#root>

### IPv4 ###


Leaf-02#

sh ip igmp snooping groups vlan 102


Vlan     Group                Type            Version         Port List
------------------------------------------------------------------------
102      226.1.1.1

igmp

            v2

 Gi1/0/10

<-- Receiver joined on Gi1/0/10


### IPv6 ###

Leaf-02#

sh ipv6 mld vrf green groups detail
Interface:      Vlan102                 <-- Join on Vlan 102


Group:          FF06:1::1               <-- Group joined


Uptime:         06:38:25
Router mode:    EXCLUDE (Expires: 00:02:14)
```

```
Host mode:      INCLUDE
Last reporter:  FE80::46D3:CAFF:FE28:6CC1 <-- MLD join from Receiver link-local address


Source list is empty                      <-- ASM join, no sources listed


Leaf-02#
sh ipv6 neighbors vrf green
IPv6 Address
                                Age Link-layer Addr State Interface
FE80::46D3:CAFF:FE28:6CC1
                    0
44d3.ca28.6cc1
  REACH Vl102
<-- Receiver IP & MAC
Leaf-02#sh ipv6 mld snooping address vlan 102 <-- If MLD snooping is on, it can be checked as well
Vlan     Group               Type             Version         Port List
------------------------------------------------------------------------
102
FF06:1::1


  mld
             v2
Gi1/0/10        <-- Receiver joined on Gi1/0/10
```

验证Leaf-02: MVPN调试显示，当IGMP/MLD成员身份报告到达且已安装所需的EVPN类型2和类型5时，已创建路由类型7。

```
<#root>
### IPv4 ###


Leaf-02#
debug bgp ipv4 mvpn updates


Leaf-02#
debug bgp ipv4 mvpn updates events


*Dec 14 19:41:57.645: BGP[15] MVPN:
add c-route, type 7
```

, bs len 0 asn=0,

rd=1:1

,
*Dec 14 19:41:57.645:

source=10.1.101.11/4,

*Dec 14 19:41:57.645:

group=226.1.1.1/4,

*Dec 14 19:41:57.645:

nexthop=172.16.254.3

,

<-- Source is via Leaf-01 IP

*Dec 14 19:41:57.645: len left = 0
*Dec 14 19:41:57.645: BGP[14] MVPN umh lookup: vrfid 2, source 10.1.101.11
*Dec 14 19:41:57.645: BGP[4] MVPN umh lookup: vrfid 2, source 10.1.101.11, net 1:1:10.1.101.11/32, 1:1:

0x10B:172.16.255.3:2

,
*Dec 14 19:41:57.646:

BGP: MVPN(15) create local route [7][172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22

*Dec 14 19:41:57.646:

BGP[15] MVPN: add c-route, type 7, bs len 0 asn=65001, rd=1:1,

IPv6

Leaf-02#

debug bgp ipv6 mvpn updates

Leaf-02#

debug bgp ipv6 mvpn updates events

Mar 23 15:46:11.171: BGP[16] MVPN:

add c-route, type 7

, bs len 0 asn=0, rd=1:1,
Mar 23 15:46:11.171:

source=FC00:1:101::11/16,

Mar 23 15:46:11.171:

group=FF06:1::1/16,

Mar 23 15:46:11.171:

nexthop=::FFFF:172.16.254.3

,

<-- IPv4 next hop of Leaf-01

Mar 23 15:46:11.171: len left = 0
Mar 23 15:46:11.171: BGP[19] MVPN umh lookup: vrfid 2, source FC00:1:101::11
Mar 23 15:46:11.171: BGP[5] MVPN umh lookup: vrfid 2, source FC00:1:101::11, net [1:1]FC00:1:101::11/12

0x10B:172.16.255.3:2

,

Mar 23 15:46:11.172: BGP: MVPN(16) create local route [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:
Mar 23 15:46:11.172: BGP[16] MVPN: add c-route, type 7, bs len 0 asn=65001, rd=1:1,

验证Leaf-01:从Leaf-02收到的MVPN类型7

<#root>

IPv4

Leaf-01#

sh bgp ipv4 mvpn all route-type 7 172.16.254.3:101 65001 10.1.101.11 226.1.1.1

...or you can also use:
Leaf-01#

sh bgp ipv4 mvpn detail [7][172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22

BGP routing table entry for

[7][172.16.254.3:101]

[65001][10.1.101.11/32][226.1.1.1/32]/22, version 76
Paths: (2 available, best #1, table

MVPNv4-BGP-Table

)

<-- In BGP IPv4 MVPN table

  Not advertised to any peer
  Refresh Epoch 1
  Local

    172.16.255.4

 (metric 3) from 172.16.255.2 (172.16.255.2)

```
<-- loopback of Leaf-02 Receiver VTEP


      Origin incomplete, metric 0, localpref 100, valid, internal

      Extended Community: RT:172.16.255.3:2                  <-- The VRI derived from EVPN Type-2 and ad


      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Dec 15 2020 14:14:38 UTC



### IPv6 ###


Leaf-01#

sh bgp ipv6 mvpn all route-type 7 172.16.254.3:101 65001 FC00:1:101::11 FF06:1::1

...or you can also use:
Leaf-01#

sh bgp ipv6 mvpn detail [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46


BGP routing table entry for

[7][172.16.254.3:101]

[65001][FC00:1:101::11][FF06:1::1]/46, version 45
Paths: (2 available, best #1, table

MVPNV6-BGP-Table

)

<-- In BGP IPv6 MVPN table


  Not advertised to any peer
  Refresh Epoch 1
  Local

    172.16.255.4

 (metric 3) from 172.16.255.1 (172.16.255.1)

<-- loopback of Leaf-02 Receiver VTEP


      Origin incomplete, metric 0, localpref 100, valid, internal, best

      Extended Community: RT:172.16.255.3:2    <-- The VRI derived from EVPN Type-2 and added to the MVP


      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Mar 23 2021 15:46:11 UTC
```

验证Leaf-01:MVPN调试显示通过MVPN VRI Route-Target接收的路由类型7

<#root>

*Dec 17 16:16:31.923: BGP(15): 172.16.255.2

rcvd UPDATE w/ attr: nexthop 172.16.255.4

, origin ?, localpref 100, metric 0, originator 172.16.255.4, clusterlist 172.16.255.2,

extended community RT:172.16.255.3:2 <-- VRI RT

*Dec 17 16:16:31.923: BGP(15): 172.16.255.2

rcvd [7]

[172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22

<-- Received MVPN Type-7

<...only update from Spine-02 172.16.255.2 ...>

*Dec 17 16:16:31.923: BGP(15): skip vrf default table RIB route [7][172.16.254.3:101][65001][10.1.101.1
*Dec 17 16:16:31.924: BGP(15): add RIB route (0:0)[7][1:1][65001][10.1.101.11/32][226.1.1.1/32]/22

(Skipping IPv6, see the debugs demonstrated in previous steps)

验证Leaf-02：完整BGP表包含Leaf-01 EVPN类型2和MVPN类型5，以及由接收方Leaf-02生成的类型7

<#root>

IPv4

Leaf-02#

sh bgp l2vpn evpn all | b 10.1.101.11

* i

[2]

[172.16.254.3:101][0][48][F4CFE24334C5][32][10.1.101.11]/24

<-- Remote VTEP route-type 2

172.16.254.3 0 100 0 ?

***>i 172.16.254.3 0 100 0 ? <-- IP of Leaf01 Lo1**

Leaf-02#

sh bgp ipv4 mvpn all

Network Next Hop Metric LocPrf Weight Path

Route Distinguisher: 1:1

(default for vrf green)

<-- default RD for vrf green

*>i

[5][1:1][10.1.101.11][226.1.1.1]

/18

<-- Type-5, source & group

172.16.255.3

0 100 0 ?

<-- Next hop Leaf-01 IP

* i 172.16.255.3 0 100 0 ?

Route Distinguisher: 172.16.254.3:101 <-- MVPN RD sent from Source Leaf-01

*>

[7]

[172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22

<-- Type-7 BGP Join Entry

0.0.0.0

32768

?

<-- Locally created (0.0.0.0) by Leaf-02

IPv6

Leaf-02#

sh bgp l2vpn evpn all | b FC00:1:101::11

* i

[2]

[172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36

<-- Remote VTEP route-type 2

172.16.254.3 0 100 0 ?

***>i 172.16.254.3 0 100 0 ? <-- IP of Leaf-01 Lo1**

```
Leaf-02#
sh bgp ipv6 mvpn all
     Network          Next Hop              Metric LocPrf Weight Path

Route Distinguisher: 1:1
 (default for vrf green)
<-- default RD for vrf green

 *>i
[5][1:1][FC00:1:101::11][FF06:1::1]
/42
<-- Type-5, source & group



    172.16.255.3
                0 100        0 ?
<-- IPv4 Next hop Leaf-01 IP

 * i                   172.16.255.3               0 100        0 ?
Route Distinguisher: 172.16.254.3:101                              <-- MVPN RD sent from Source Leaf-0

 *>
[7]
[172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
<-- Type-7 BGP Join Entry



::                                    32768
 ?
<-- Locally created (::) by Leaf-02
```

验证TRM组枝叶–01(FHR)

检查MDT和TRM组是否在源端正确形成。

- TRM组的传入接口是与客户端VRF关联的SVI
- TRM组的传出接口是L3VNI SVI

验证Leaf-01:TRM组MRIB/MFIB

<#root>

IPv4

Leaf-01#

sh ip mroute vrf green 226.1.1.1 10.1.101.11

(10.1.101.11, 226.1.1.1), 02:57:56/00:03:14,

flags: FTGqx <-- Flags: BGP S-A Route

  Incoming interface:

Vlan101

, RPF

nbr 0.0.0.0   <-- Local to Vlan101 Direct connected source

  Outgoing interface list:

Vlan901

, Forward/Sparse, 02:57:56/stopped

<-- OIF is VxLAN L3VNI

Leaf-01#

sh ip mfib vrf green 226.1.1.1 10.1.101.11

VRF green   <-- Tenant VRF

(10.1.101.11,226.1.1.1) Flags: HW
  SW Forwarding: 1/0/100/0, Other: 0/0/0

HW Forwarding: 5166/0/118/0, Other: 0/0/0 <-- Hardware counters indicate the entry is operating in hardw

Vlan101 Flags: A   <-- Accept flag set on Connected Source SVI

  Vlan102 Flags: F NS
    Pkts: 0/0/1 Rate: 0 pps

Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F <-- Forward via Vlan 901. Use MDT group 239.1.1.1, v

    Pkts: 0/0/0 Rate: 0 pps

```
### IPv6 ###
Leaf-01#
sh ipv6 mroute vrf green
(FC00:1:101::11, FF06:1::1), 01:01:00/00:01:08,
flags: SFTGq <-- Flags: q - BGP S-A Route, G - BGP Signal Received
  Incoming interface:
Vlan101

 RPF nbr: FE80::F6CF:E2FF:FE43:34C1  <-- link local address of Source
  Immediate Outgoing interface list:


Vlan901
, Forward, 01:01:00/never
<-- OIF is VxLAN L3VNI

Leaf-01#
sh ipv6 mfib vrf green FF06:1::1
VRF green   <-- Tenant VRF

(FC00:1:101::11,FF06:1::1) Flags: HW

   SW Forwarding: 0/0/0/0, Other: 1/0/1


HW Forwarding: 1968/0/118/0, Other: 0/0/0 <-- Hardware counters indicate the entry is operating in hardw

 Vlan101 Flags: A NS                        <-- Accept flag set on Connected Source SVI

 Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F <-- Forward via Vlan 901. Use MDT group 239.1.1.1,
     Pkts: 0/0/0 Rate: 0 pps
```

验证Leaf-01:FED中的TRM组

<#root>

IPv4

Leaf-01#

sh platform software fed switch active ip mfib vrf green 226.1.1.1/32 10.1.101.11

Multicast (S,G) Information

VRF   : 2   <-- VRF ID 2 = vrf green (from "show vrf detail")

Source Address : 10.1.101.11
HTM Handler   : 0x7f175cc08578
SI Handler   : 0x7f175cc06ea8
DI Handler   : 0x7f175cc067c8
REP RI handler : 0x7f175cc06b38
Flags   : {Svl}

Packet count   : 39140   <-- packets that used this adjacency (similar to mfib command, but shown at

State   : 4

RPF

**   :**

**   Vlan101   A   <-- Accept on Vlan 101 in Tenant vrf green**

OIF   :
     Vlan102   F NS
       Vlan101   A

       Vlan901   F {Remote}

<-- Forward via L3VNI interface

**   (Adj: 0x6a)   <-- Adjacency for this entry**

IPv6

Leaf-01#

sh plat soft fed switch active ipv6 mfib vrf green FF06:1::1/128 FC00:1:101::11

Multicast (S,G) Information

VRF   : 2   <-- VRF ID 2 = vrf green (from "show vrf detail")

Source Address : fc00:1:101::11
HTM Handler   : 0x7fba88d911b8
SI Handler   : 0x7fba88fc4348
DI Handler   : 0x7fba88fc8dc8
REP RI handler : 0x7fba88fc8fd8
Flags   : {Svl}

Packet count   : 2113

<-- packets that used this adjacency (similar to mfib command, but shown at the FED layer)

State   : 4

RPF   :

**   Vlan101   A {Remote}   <-- Accept on Vlan 101 in Tenant vrf green (says remote, but this is a loca**

OIF   :
     Vlan101   A {Remote}

** Vlan901   F {Remote}**

<-- Forward via L3VNI interface

** (Adj: 0x7c)   <-- Adjacency for this entry**

验证Leaf-01：邻接关系正确

<#root>

IPv4

Leaf-01#

sh platform software fed switch active ip adj

IPV4 Adj entries

dest   if_name   dst_mac   si_hdl   ri_hdl

adj_id

  Last-modified
----   -------   -------   ------   ------

239.1.1.1

nve1.VNI50901

  4500.0000.0000 0x7f175ccd8c38 0x7f175ccd8de8 0x60

 0x6a

    2020/12/16 17:39:55.747

```
*** Adjacency 0x6a details ***


Destination =

 the MDT tunnel multicast group 239.1.1.1

Interface   =

 nve1.VNI50901 (the L3VNI 50901)


### IPv6 ###

Leaf-01#

sh platform software fed switch active ipv6 adj
IPV6 Adj entries

dest                                          if_name        dst_mac         si_hdl          ri_hdl

adj_id

  Last-modified
----                                          -------        -------         ------          ------

239.1.1.1


nve1.VNI50901

  4500.0000.0000  0x7fba88cf9fc8  0x7fba88cfa248 0x60

  0x7c

    2021/03/22 19:54:09.831

*** Adjacency 0x7c details ***
Destination =

the MDT tunnel multicast group 239.1.1.1

Interface   =

nve1.VNI50901 (the L3VNI 50901)
```

验证TRM组枝叶02(LHR)

检查MDT和TRM组是否在接收器端正确形成。

- TRM组的传入接口是与L3VNI关联的SVI
- TRM组的传出接口是处理IGMP加入的客户端SVI。

验证Leaf-02:MRIB/MFIB中的TRM（租户组播路由）路由

<#root>

```
Leaf-02#
sh ip mroute vrf green 226.1.1.1 10.1.101.11      <-- The TRM Client group

(10.1.101.11, 226.1.1.1), 00:26:03/00:02:37, flags: TgQ
  Incoming interface: Vlan901, RPF nbr 172.16.254.3       <-- Via L3VNI, RPF to Leaf-01

  Outgoing interface list:
    Vlan102,
 Forward/Sparse, 00:26:03/00:03:10
<-- Client Receiver Vlan


Leaf-02#
sh ip mfib vrf green 226.1.1.1 10.1.101.11


VRF green                                          <--- The Tenant VRF

  (10.1.101.11,226.1.1.1) Flags: HW
    SW Forwarding: 1/0/100/0, Other: 0/0/0

HW Forwarding: 39013/0/126/0, Other: 0/0/0     <-- Hardware counters indicate the entry is operating in



 Vlan901, VXLAN Decap Flags: A                  <-- L3VNI Accept and decapsulate from VxLAN



Vlan102 Flags: F NS                            <-- Forward to the Tenant Vlan

      Pkts: 0/0/1 Rate: 0 pps
```

验证Leaf-02:FED中的TRM组

<#root>

```
### IPv4 ###


Leaf-02#
sh platform software fed switch active ip mfib vrf green 226.1.1.1/32 10.1.101.11 detail <-- Use detail


MROUTE ENTRY vrf 2 (10.1.101.11, 226.1.1.1/32)
  HW Handle: 140397391947768 Flags: {Svl}
RPF interface: Vlan901
```

(60)):

SVI   <-- RPF interface = L3VNI SVI Vlan901

  HW Handle:140397391947768 Flags:A {Remote}
Number of OIF: 2
Flags: 0x4

Pkts : 39387   <-- packets that used this adjacency (similar to mfib command, but shown at the FED la

OIF Details:

     Vlan102 F NS   <-- Client Vlan

     Vlan901 A {Remote}   <-- Accept interface is RPF to source via Remote EVPN next hop

          (Adj: 0xf80003c1) <-- Adj for vlan 901(show plat soft fed sw active ipv4 adj)

Htm: 0x7fb0d0edfb48 Si: 0x7fb0d0ee9158 Di: 0x7fb0d0eca8f8 Rep_ri: 0x7fb0d0ef2b98

DI details <-- Dest index (egress interface) details

Handle:0x7fb0d0eca8f8 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICA
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x538b mtu_index/l3u_ri_index0:0x0 index1

Brief Resource Information

(ASIC_INSTANCE# 1)

 <-- Gi1/0/10 is mapped to instance 1

Destination index = 0x538b

pmap = 0x00000000 0x00000200

pmap_intf : [GigabitEthernet1/0/10]   <-- Maps to Gi1/0/10, the port toward the client

==

IPv6

Leaf-02#

sh platform software fed switch active ipv6 mfib vrf green FF06:1::1/128 FC00:1:101::11 detail

MROUTE ENTRY

vrf 2

 (fc00:1:101::11, ff06:1::1/128)
  HW Handle: 139852137577736 Flags: {Svl}

RPF interface: Vlan901

(62)): SVI

<-- RPF to Source L3VNI SVI 901

HW Handle:139852137577736

Flags:A {Remote}

Number of OIF: 2

Flags: 0x4 Pkts : 7445 <-- Packets use this Entry

OIF Details:

Vlan102 F NS <-- F - Forward. The OIF Vlan SVI 901

Vlan901 A {Remote}

(Adj: 0xf80003e2) <-- Adj for vlan 901 (show plat soft fed sw active ipv6 adj)

Htm: 0x7f31dcfee238 Si: 0x7f31dcfba5d8 Di: 0x7f31dcfc2358 Rep_ri: 0x7f31dcfcb1a8

DI details

Handle:0x7f31dcfc2358 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICA
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x5381 mtu_index/l3u_ri_index0:0x0 index1

Brief Resource Information

(ASIC_INSTANCE# 1) <-- Gig1/0/10 is mapped to Instance 1

Destination index = 0x5381

pmap = 0x00000000 0x00000200

pmap_intf : [GigabitEthernet1/0/10] <-- Maps to Gig1/0/10, the port toward the client

==========

Leaf-02#

sh platform software fed switch active ifm mappings

Interface IF_ID

Inst

```
Asic
 Core Port SubPort Mac Cntx LPN GPN Type Active
GigabitEthernet1/0/10
     0x12
 1


0
   1    9      0      5   15   10  10  NIF  Y
 <-- Instance 1 of ASIC 0
```

Verify Leaf-02：捕获的数据包显示具有内部客户端流量的外部MDT隧道组

<#root>

```
Leaf-02#
sh mon ca 1 parameter


monitor capture 1 interface GigabitEthernet1/0/2 IN
monitor capture 1 match any
monitor capture 1 buffer size 10
monitor capture 1 limit pps 1000


### IPv4 ###


Leaf-02#
sh mon capture 1 buffer detailed


Ethernet II, Src: 7c:21:0d:bd:2c:d6 (7c:21:0d:bd:2c:d6),
Dst: 01:00:5e:01:01:01
 (01:00:5e:01:01:01)
<-- MAC is matching 239.1.1.1



Type: IPv4 (0x0800)  <-- IPv4 outer packet


Internet Protocol Version 4,
Src: 172.16.254.3, Dst: 239.1.1.1 <- Leaf-01 Source IP and MDT outer tunnel Group


    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Time to live: 253
```

User Datagram Protocol

, Src Port: 65287,

Dst Port: 4789 <-- VxLAN UDP port 4789

Virtual eXtensible Local Area Network

Flags: 0x0800,

VXLAN Network ID (VNI)

Group Policy ID: 0

VXLAN Network Identifier (VNI): 50901 <-- L3VNI value
Type: IPv4

(0x0800)

<-- IPv4

inner packet

Internet Protocol Version 4

,

Src: 10.1.101.11, Dst: 226.1.1.1 <-- Encapsulated IPv4 TRM group

0100 .... = Version: 4
Time to live: 254
Protocol: ICMP (1)

(multiple lines removed from this example capture)

IPv6

Leaf-02#

sh mon capture 1 buffer detailed

Ethernet II,

Src: 7c:21:0d:bd:2c:d6

(7c:21:0d:bd:2c:d6),

Dst: 01:00:5e:01:01:01

(01:00:5e:01:01:01)

<-- DMAC is matching 239.1.1.1

Type: IPv4 (0x0800) **<-- IPv4 outer packet**

Internet Protocol Version 4, Src: 172.16.254.3, Dst: 239.1.1.1

    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 150
    Identification: 0x4e4b (20043)
    Flags: 0x4000, Don't fragment
        0... .... = Reserved bit: Not set

**   .1.. .... = Don't fragment: Set <-- DF flag=1. MTU can be an issue if too low in path**

        ..0. .... = More fragments: Not set
    ...0 0000 0000 0000 = Fragment offset: 0
    Time to live: 253

**   Protocol: UDP (17)**

    Header checksum: 0x94f4 [validation disabled]
    [Header checksum status: Unverified]
    Source: 172.16.254.3
    Destination: 239.1.1.1
User Datagram Protocol,

Src Port: 65418, Dst Port: 4789 <-- VxLAN UDP port 4789

    Source Port: 65418

** Destination Port: 4789**

<...snip...>

Virtual eXtensible Local Area Network

    Flags: 0x0800,

VXLAN Network ID (VNI)

        0... .... = GBP Extension: Not defined
        .... .0.. .... = Don't Learn: False

** .... 1... .... = VXLAN Network ID (VNI): True**

        .... 0... = Policy Applied: False
        .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 0

**   VXLAN Network Identifier (VNI): 50901 <-- L3VNID 50901**

    Reserved: 0
Ethernet II, Src: 10:b3:d5:6a:00:00 (10:b3:d5:6a:00:00), Dst:

```
33:33:00:00:00:01
 (33:33:00:00:00:01)
<-- DMAC matches ff06:1::1

    Type: IPv6 (0x86dd)                  <-- IPv6 inner packet

Internet Protocol Version 6
,
Src: fc00:1:101::11, Dst: ff06:1::1 <-- Encapsulated IPv6 TRM group

    0110 .... = Version: 6
<...snip...>

    Source: fc00:1:101::11

    Destination: ff06:1::1

Internet Control Message Protocol v6
    Type: Echo (ping) request (128)
<...snip...>
```

场景2：交换矩阵中的PIM SSM

在此模式下，重叠中没有RP，并且没有使用MVPN类型5或类型7（底层继续作为PIM ASM运行）。 在SSM中，接收方向LHR VTEP发送和IGMPv3 S,G加入。此VTEP对RIB中的源执行RPF查找。如果发现L3VNI SVI作为RPF接口，LHR VTEP会将MVPN RT-7发送到接收和安装此路由的FHR VTEP。然后，FHR VTEP通知PIM添加L3VNI SVI作为S，G mroute的传出接口。

本节显示与方案1的差异。相同的步骤和方法仅在场景1中说明。

- 请参阅场景1中BGP和PIM的验证和调试步骤，因为BGP和PIM操作相同

网络图

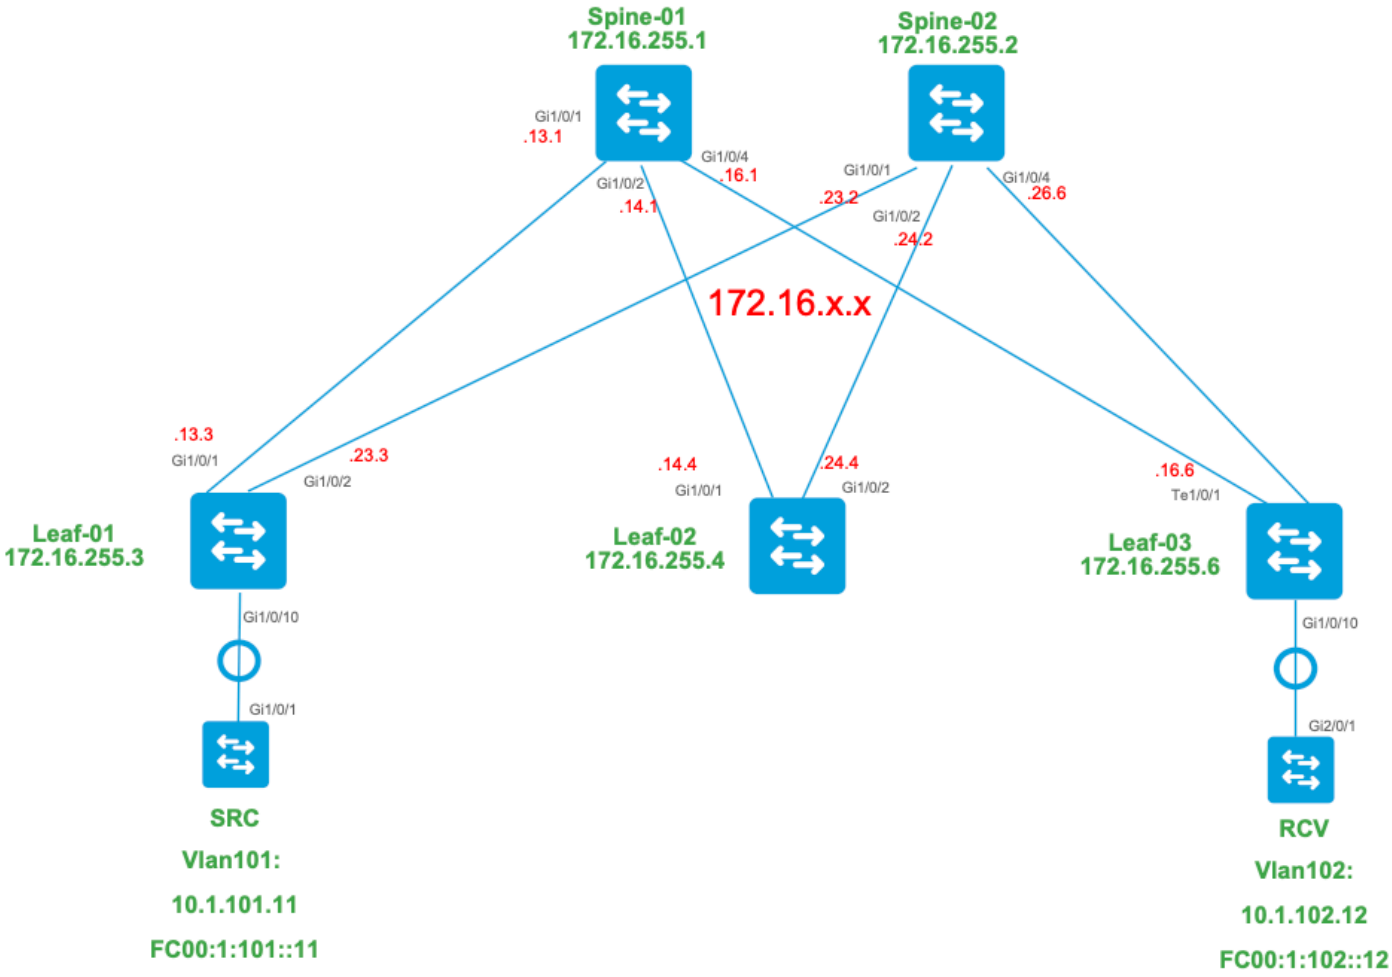

对于此模式，请考虑以下BGP路由类型及其来源

创建者：源VTEP

- EVPN路由类型2。用于获取源的单播和VRI信息，并在VTEP加入STP树时添加到C组播路由（MVPN类型7）。

创建者：接收方VTEP

- MVPN路由类型7。来自IGMP或MLD层以及来自EVPN类型2的信息用于创建此BGP类型联接。类型7驱动在源端创建MRIB OIF。

EVPN第2类要求：

1. FHR（源VTEP）验证ARP（或ND）和CEF邻接（确认源为直连）。
2. FHR发出EVPN第2类BGP更新

MVPN类型7要求：

1. 存在EVPN类型2条目（使用正确的VRI构建C-Multicast路由类型7时需要该条目，并且是从源VTEP发送的）
2. 接收方VTEP:LHR VTEP已接收并处理源特定成员身份报告
3. LHR VTEP RPF接口是交换矩阵L3VNI接口

对于此模式，LHR VTEP上需要添加配置以启用SSM范围，并处理IGMPv3成员身份报告

配置Leaf-03：在租户SVI下将IGMP查询器设置为版本3

<#root>

interface Vlan102

vrf forwarding green
ip address 10.1.102.1 255.255.255.0
ip pim sparse-mode

ip igmp version 3 <-- Sets the version to V3

end

验证Leaf-03:IGMP查询器设置为版本3

<#root>

Leaf-03#

sh ip igmp snooping querier vlan 102

IP address : 10.1.102.1   <-- IP is that of the Vlan102 SVI

IGMP version : v3   <-- Querier is now version 3

Port : Router   <-- Mrouter port is "Router" meaning querier is local to this VTEP

Max response time : 10s
Query interval : 60s
Robustness variable : 2

启用枝叶03:租户VRF所需的SSM范围

<#root>

Leaf-03(config)#

ip pim vrf green ssm

?

```
default
Use 232/8 group range for SSM          <-- Set to the normally defined SSM range

range
ACL for group range
 to be used for SSM
<-- use an ACL to define a non-default SSM range
```

提示:SSM组不创建*,G mroute。如果您看到组的*,G，请验证您的配置是否正确。

验证此方案所需事件的顺序

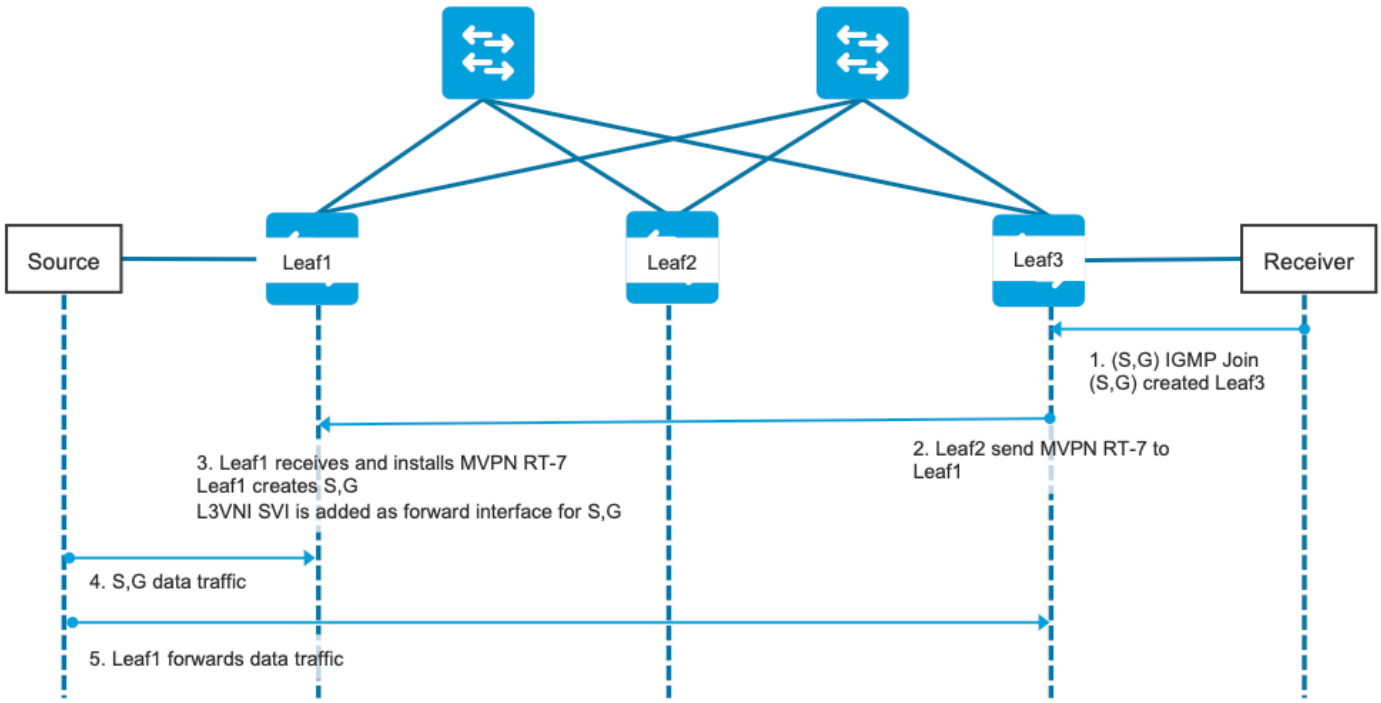

第0步EVPN(Leaf-03)：验证BGP可以找到VRI在MVPN类型7中使用的EVPN前缀。

<#root>

```
Leaf-03#
sh bgp l2vpn evpn all

BGP table version is 16, local router ID is 172.16.255.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
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 green)
* i

[2]

[172.16.254.3:101][0][48][F4CFE24334C1][32]

[10.1.101.11]

/24
                    172.16.254.3 0 100 0 ?

*>i                  172.16.254.3 0 100 0 ? <-- From Leaf-01



Leaf-03#

sh bgp l2vpn evpn all route-type 2 0 F4CFE24334C1 10.1.101.11     <-- Detailed view of the EVPN type-2 e


BGP routing table entry for

[2]

[172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24, version 283
Paths: (2 available, best #2,

table EVPN-BGP-Table

)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.3 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 MVPN AS:65001:0.0.0.0


 MVPN VRF:172.16.255.3:4

 ENCAP:8 Router MAC:10B3.D56A.8FC8

<-- BGP finds the VRI in this entry


      Originator: 172.16.255.3, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2021 16:17:06 UTC
```

第1步（枝叶–03）：已接收并包含源的IGMPv3成员身份报告

```
<#root>

Leaf-03#

show ip igmp snooping groups vlan 102 226.1.1.1
```

```
Vlan
Group
                    Type
Version
      Port List
------------------------------------------------------------------------
102
226.1.1.1
                igmp
 v3
           Gi1/0/10


Leaf-03#
show ip igmp snooping groups vlan 102 226.1.1.1 sources <-- Specify "sources" to see Source information

Vlan      Group                   Type        Version      Port List
------------------------------------------------------------------------


Source information for group 226.1.1.1

:
Timers: Expired sources are deleted on next IGMP General Query


SourceIP
       Expires   Uptime
Inc Hosts
 Exc Hosts
-----------------------------------------------------
10.1.101.11
    00:01:20  00:02:58
1
         0
<-- Source specified in IGMP includes one source
```

第2步（枝叶–03）：BGP会收到此加入通知，创建并发送第7类MVPN加入。

```
<#root>
debug mvpn
```

debug ip igmp vrf green 226.1.1.1

May 6 17:11:08.500:

IGMP(6): Received v3 Report for 1 group on Vlan102 from 10.1.102.12

May 6 17:11:08.500:

IGMP(6): Received Group record for group 226.1.1.1, mode 5 from 10.1.102.12 for 1 sources <-- IGMPv3 typ

May 6 17:11:08.500: IGMP(6): WAVL Insert group: 226.1.1.1 interface: Vlan102 Successful
May 6 17:11:08.500: IGMP(6): Create source 10.1.101.11
May 6 17:11:08.500: IGMP(6): Updating expiration time on (10.1.101.11,226.1.1.1) to 180 secs
May 6 17:11:08.500: IGMP(6): Setting source flags 4 on (10.1.101.11,226.1.1.1)
May 6 17:11:08.500: IGMP(6): MRT Add/Update Vlan102 for (10.1.101.11,226.1.1.1) by 0

May 6 17:11:08.501:

MVPN: Received local route update for (10.1.101.11, 226.1.1.1) with RD: 1:1, Route Type: 7, flags: 0x00

May 6 17:11:08.501: MVPN: Route Type 7 added [(10.1.101.11, 226.1.1.1)] rd:1:1 send:1
May 6 17:11:08.501:

MVPN: Sending BGP prefix=[7:0 1:1 : (10.1.101.11,226.1.1.1)] len=23, nh 172.16.254.3, Originate route

May 6 17:11:08.501:

MVPN: Originate C-route, BGP remote RD 1:1

Leaf-03#

sh bgp ipv4 mvpn all

BGP table version is 10, local router ID is 172.16.255.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
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 green)
*>

[7][1:1][65001][10.1.101.11/32][226.1.1.1/32]/22   <-- Locally created Type-7

                      0.0.0.0   32768 ?

Leaf-03#

sh ip mroute vrf green 226.1.1.1   <-- for SSM you only see S,G and no *,G

IP Multicast Routing Table

<...snip...>

(10.1.101.11, 226.1.1.1), 00:29:12/00:02:46, flags: sTIg <-- s = SSM, I = Source Specific Join received,

Incoming interface: Vlan901

, RPF nbr 172.16.254.3

<-- RPF interface is the L3VNI

  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:29:12/00:02:46

第3步（枝叶–01）：源枝叶接收并安装MVPN 7类加入路由，并通知PIM安装L3VNI OIF

<#root>

debug mvpn

debug ip pim vrf green 226.1.1.1

May 6 18:16:07.260: MVPN: Received BGP prefix=[7:65001 1:1 : (10.1.101.11,226.1.1.1)] len=23, nexthop: 1

May 6 18:16:07.260: MVPN: Received BGP route update for (10.1.101.11, 226.1.1.1) with RD: 1:1, Route Ty
May 6 18:16:07.260: MVPN:

Route Type 7 added [(10.1.101.11, 226.1.1.1), nh 172.16.255.6] rd:1:1 send:0, to us   <-- add type-7 rou

May 6 18:16:07.260: PIM(4)[green]: Join-list: (10.1.101.11/32, 226.1.1.1), S-bit set, BGP C-Route
May 6 18:16:07.263:

PIM(4)[green]: Add Vlan901/0.0.0.0 to (10.1.101.11, 226.1.1.1), Forward state, by BGP SG Join <-- PIM a

May 6 18:16:07.264: PIM(4)[green]: Insert (10.1.101.11,226.1.1.1) join in nbr 10.1.101.11's queue
May 6 18:16:07.264:

MVPN(green[AF_IPv4]): Add (10.1.101.11, 226.1.1.1) intf Vlan901 olist Join state for BGP C-Rt type 7 Acc

Leaf-01#

sh bgp ipv4 mvpn all

<...snip...>
    Network   Next Hop   Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)

***>i [7][1:1][65001][10.1.101.11/32][226.1.1.1/32]/22**

172.16.255.6

0 100 0 ?

<-- Recieved from Reciever Leaf-03

* i 172.16.255.6 0 100 0 ?

Leaf-01#

sh ip mroute vrf green 226.1.1.1

<...snip...>

(10.1.101.11, 226.1.1.1), 00:42:41/stopped, flags: sTGx <-- s = SSM Group, G = Received BGP

Incoming interface: Vlan101, RPF nbr 10.1.101.11

Outgoing interface list:

Vlan901, Forward/Sparse, 00:42:41/stopped <-- L3VNI installed as OIF interface

第4步和第5步（枝叶–01和枝叶–03）：组播到达FHR枝叶并通过交换矩阵发送到LHR枝叶。此处提供的验证命令摘要。您可以在场景1中检查这些命令的详细验证。

<#root>

show ip mroute vrf green 226.1.1.1 count <-- software mr

show ip mfib vrf green 226.1.1.1

<-- hardware mroute details & counters

sh platform software fed switch active ip mfib vrf green 226.1.1.1/32 10.1.101.11 detail <-- ASIC entry

场景3：交换矩阵内的单个RP（常规稀疏模式）

此模式可互换称为非任意播RP或外部RP模式。在此模式下，重叠中只有一个RP。因此，重叠中的(*,G)树可以跨越多个站点。BGP使用MVPN RT-6在交换矩阵中通告(*,G)成员身份。如果RP和FHR位于不同的站点，PIM寄存器将通过交换矩阵发送。这是重叠中PIM SM的默认操作模式。

网络图

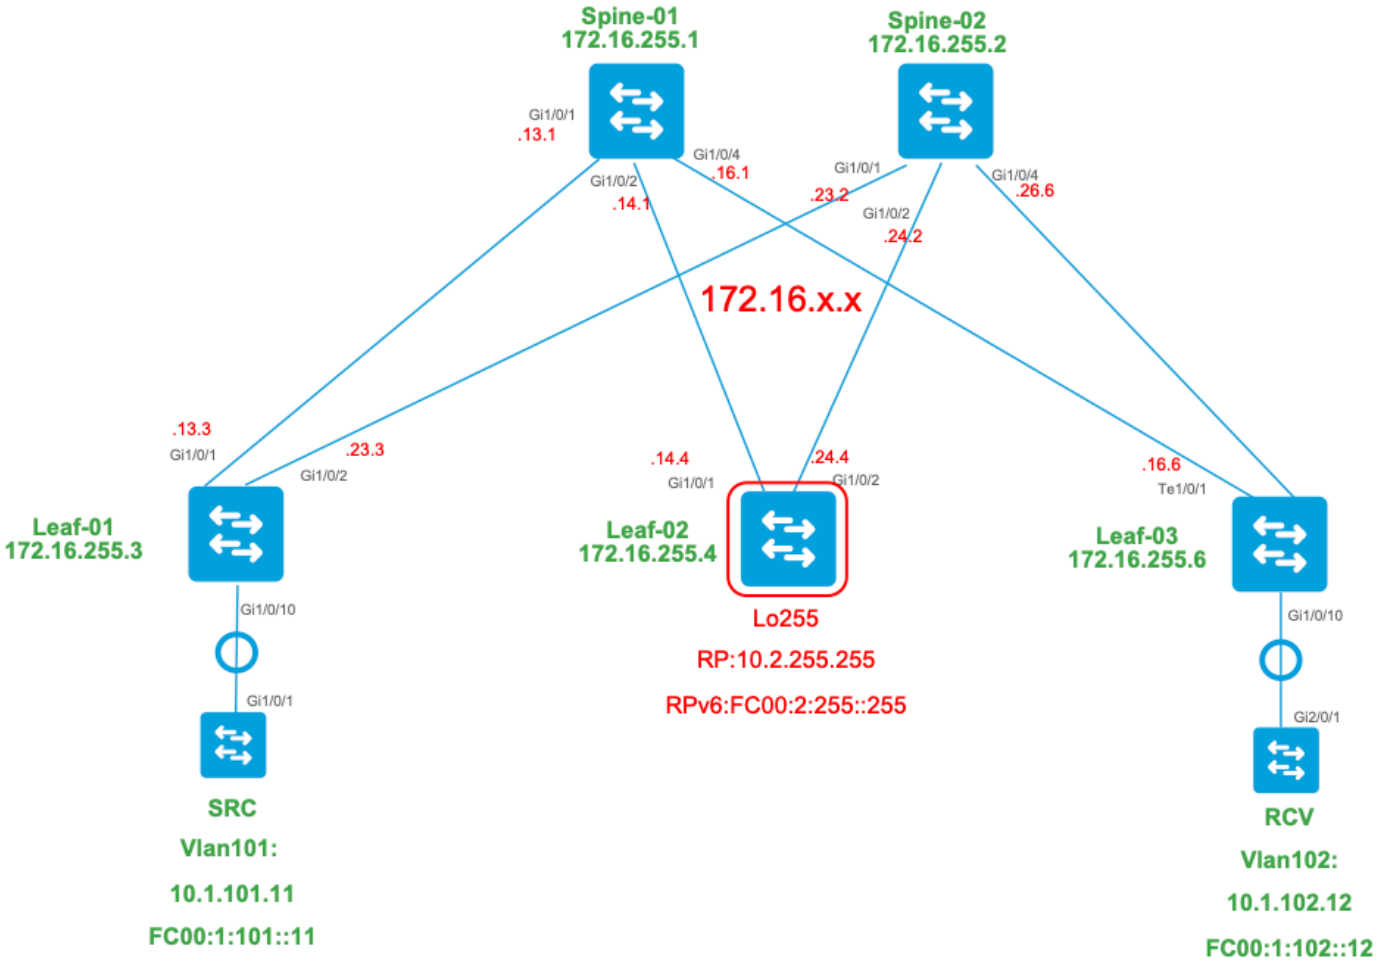

对于此模式，请考虑以下BGP路由类型及其来源

创建者：源VTEP

- EVPN路由类型2。用于获取源的单播和VRI信息，并在VTEP加入STP树时添加到C组播路由（MVPN类型7）。
- MVPN路由类型5。发送到S、G的VTEP的源A-D路由

创建者：RP VTEP

- EVPN路由类型5。用于获取RP环回的单播和VRI信息。环回不会创建路由类型2，因此使用类型5。
- MVPN路由类型7。这是从EVPN类型2获取并发送到源VTEP的IGMP加入+ RT VRI详细信息，并且它驱动创建MRIB OIF。

创建者：接收方VTEP

- MVPN路由类型6。 由接收方VTEP创建的路由类型，用于将共享树*,G（RPT树）连接到RP。
- MVPN路由类型7。来自IGMP或MLD层以及来自EVPN类型2的信息用于创建此BGP类型联接。类型7驱动在源端创建MRIB OIF。

EVPN第2类要求：

1. FHR（源VTEP）验证ARP（或ND）和CEF邻接（确认源为直连）。
2. FHR发出EVPN第2类BGP更新

EVPN第5类要求：

1. RP环回已配置并通告到BGP

MVPN类型5要求：

在此模式下，源站点的枝叶只有在满足这两个条件时才会通告(S，G)的源活动A-D消息。

1. 它在RPF接口上接收流向源的流量。（源将组播发送到FHR）
2. L3VNI SVI接口作为(S，G)条目的转发接口添加，作为从RP加入S，G作为PIM注册过程的一部分。（L3VNI SVI安装在OIF列表中）

MVPN类型6要求：

1. RP通告包含其VRI和单播可达性详细信息的其EVPN类型5路由。
2. 在LHR上收到的IGMP加入触发向RP的BGP更新

MVPN类型7要求：

1. 存在EVPN类型2条目（使用正确的VRI构建C-Multicast路由类型7时需要该条目，并且是从源VTEP发送的）
2. 存在MVPN类型5条目（解析可用于STP加入的源/组对时需要该条目）
3. 接收方VTEP:LHR VTEP已接收并处理IGMP成员身份报告
4. RP VTEP:RP已收到组播注册数据包，具有EVPN路由，并具有S、G的接收器（通过类型6学习）
5. LHR VTEP RPF接口是交换矩阵L3VNI接口


提示：在出口LHR VTEP PIM检查通往源的路径。PIM必须在RIB中找到作为RPF接口的L3VNI路由。如果L3VNI未正确配置、已关闭，依此类推。VTEP不会创建类型7 BGP连接。

验证此方案所需事件的顺序

验证接收方VTEP最初加入共享树所需的步骤，然后将其剪切到最短路径树。这涉及检查BGP表、IGMP和MRIB创建状态。

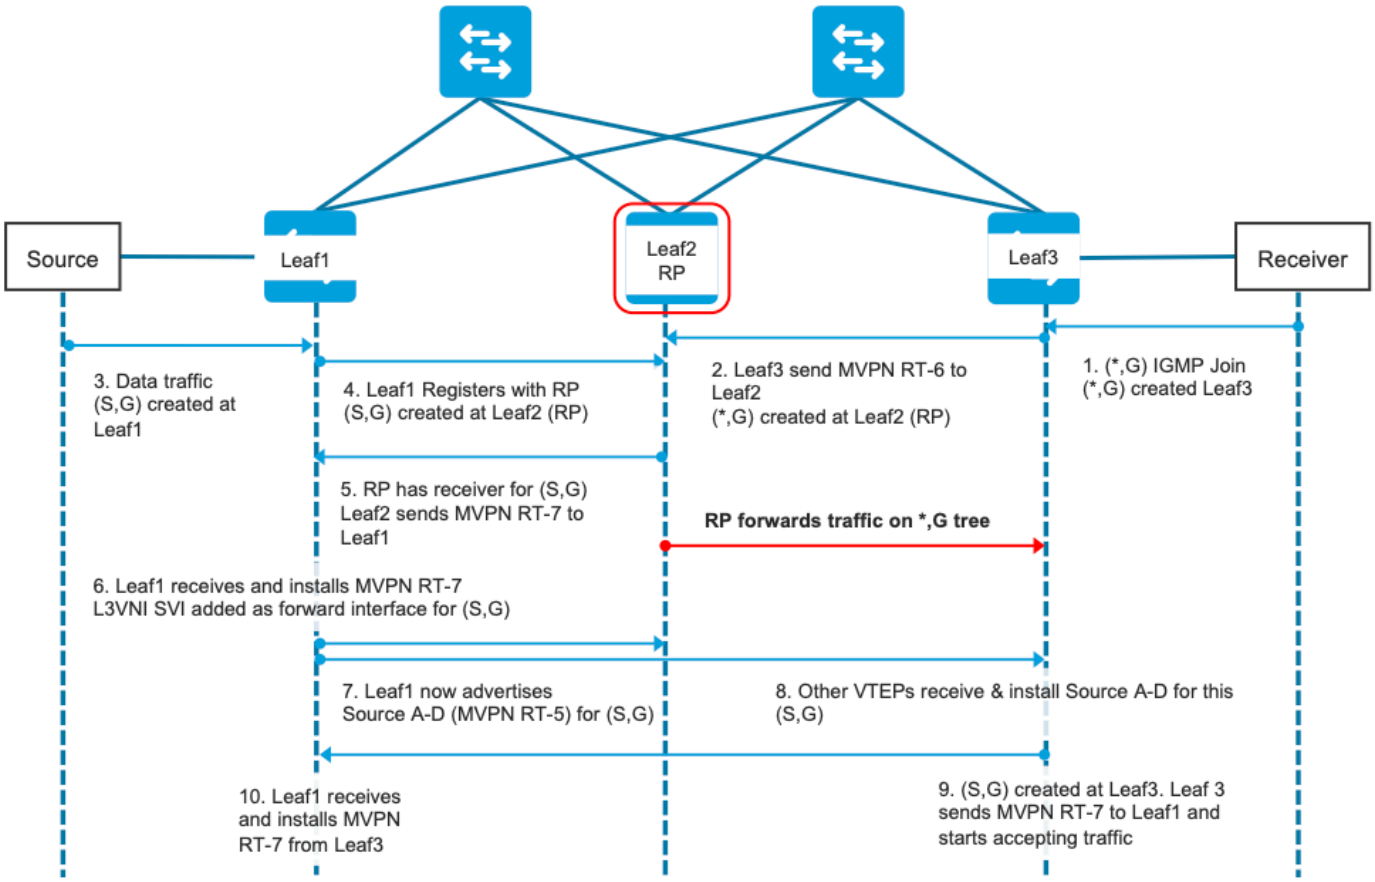

第EVPN步（枝叶–03）：从RP获取EVPN第5类LHR。接收方VTEP创建MVPN类型6路由需要此步骤

<#root>

Leaf-03#

sh bgp l2vpn evpn all route-type 5 0 10.2.255.255 32

...or you can also use:
Leaf-03#

sh bgp l2vpn evpn detail [5][1:1][0][32][10.2.255.255]/17

```
BGP routing table entry for [5][1:1][0][32][10.2.255.255]/17, version 25
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
```

172.16.254.4

```
 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
```

<-- RP's global next hop IP

```
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 MVPN AS:65001:0.0.0.0
```

```
        MVPN VRF:172.16.255.4:2

 ENCAP:8

Router MAC:7C21.0DBD.9548


      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 13 2021 19:09:31 UTC
  Refresh Epoch 2
  Local


MVPN VRF:172.16.255.4:2


<-- MVPN VRI


Router MAC:7C21.0DBD.9548 <-- Leaf-02 RMAC
```

第1步(Leaf-03)：收到IGMP成员报告

<#root>

```
Leaf-03#

sh ip igmp snooping groups


Vlan     Group            Type                Version        Port List
------------------------------------------------------------------------
102      224.0.1.40       igmp                v2             Gi1/0/10

102      226.1.1.1        igmp                v2             Gi1/0/10   <-- Client has joined
```

第2步（枝叶–03）：创建MVPN类型6，发送到RP，由RP(枝叶–02)接收

<#root>

```
#### Type-6 from the Receiver VTEP perspective ###


Leaf-03#

sh bgp ipv4 mvpn all route-type 6 1:1 65001 10.2.255.255 226.1.1.1  <-- Source is RP Loopback


...or you can also use:
Leaf-03#

sh bgp ipv4 mvpn


detail [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
```

```
BGP routing table entry for [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22, version 13
Paths: (1 available, best #1, table MVPNv4-BGP-Table)
  Advertised to update-groups:
      1
  Refresh Epoch 1
  Local
```

0.0.0.0 from 0.0.0.0 (172.16.255.6) <-- Generated locally

```
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
```

Extended Community: RT:172.16.255.4:2 <-- VRI Ext Comm added from EVPN Type-5

```
      rx pathid: 2, tx pathid: 0x0
      Updated on Jan 14 2021 14:51:29 UTC
```

Type-6 from the RP perspective

```
Leaf-02#
```

sh bgp ipv4 mvpn all route-type 6 1:1 65001 10.2.255.255 226.1.1.1 <-- type-6, RD 1:1, AS 65001, Source/

```
...or you can also use:
Leaf-02#
```

sh bgp ipv4 mvpn detail [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22

```
BGP routing table entry for
```

[6]

```
[1:1][65001][10.2.255.255/32][226.1.1.1/32]/22, version 25
Paths: (2 available, best #1, table MVPNv4-BGP-Table)
   Flag: 0x100
   Not advertised to any peer
   Refresh Epoch 2
   Local
     172.16.255.6 (metric 3) from 172.16.255.1 (172.16.255.1)
       Origin incomplete, metric 0, localpref 100, valid, internal, best
```

Extended Community: RT:172.16.255.4:2 <-- Contains VRI learned from EVPN Type-5

Originator: 172.16.255.6

```
, Cluster list: 172.16.255.1
```

<-- Sent from Leaf03 IP to RP

```
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 14 2021 14:54:29 UTC
```

第1步和第2步调试（枝叶–01）：IGMP报告、EVPN源查找和MVPN类型–6创建

<#root>

debug ip igmp vrf green 226.1.1.1

debug bgp ipv4 mvpn updates

debug bgp ipv4 mvpn updates events

Client sends IGMP membership report

IGMP processes this IGMP report

*Feb 1 21:13:19.029: IGMP(2): Received v2 Report on Vlan102 from 10.1.102.12 for 226.1.1.1

<--- IGMP processes received report

*Feb 1 21:13:19.029: IGMP(2): Received Group record for group 226.1.1.1, mode 2 from 10.1.102.12 for 0
*Feb 1 21:13:19.029: IGMP(2): WAVL Insert group: 226.1.1.1 interface: Vlan102 Successful
*Feb 1 21:13:19.029: IGMP(2): Switching to EXCLUDE mode for 226.1.1.1 on Vlan102
*Feb 1 21:13:19.029: IGMP(2): Updating EXCLUDE group timer for 226.1.1.1
Feb 1 21:13:19.029: IGMP(2): MRT Add/Update Vlan102 for (,226.1.1.1) by 0

<--- Notify MRT to add Vlan 102 into Outgoing interface list

BGP is informed by IGMP, does an EVPN source lookup, creates the MVPN Type-6 route, sends to RR

(

Without the EVPN Type-5 prefix already in BGP you see IGMP debugs trigger, but no subsequent BGP debugs

*Feb 1 21:13:19.033: BGP[15] MVPN:

add c-route, type 6

, bs len 0 asn=0, rd=1:1,

<-- Start creation of Type-6 C-multicast Shared Tree Join

*Feb 1 21:13:19.033:

source=10.2.255.255

/4,

<-- RP loopback255

*Feb 1 21:13:19.033: group=226.1.1.1/4,

<-- Group IP

```
*Feb 1 21:13:19.033:
nexthop=172.16.254.4
,
<-- Global Next-Hop learned from EVPN VRI

*Feb 1 21:13:19.033: len left = 0
*Feb 1 21:13:19.033: BGP[14]
MVPN umh lookup:
 vrfid 2, source 10.2.255.255
<-- UMH (upstream multicast hop) as found in the RT of the EVPN type-5

*Feb 1 21:13:19.033: BGP[4] MVPN umh lookup: vrfid 2, source 10.2.255.255, net 1:1:10.2.255.255/32, 1:1
<-- EVPN info adding to MVPN

*Feb 1 21:13:19.033: BGP: MVPN(15) create local route [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
<--- MVPN creating type-6

*Feb 1 21:13:19.033: BGP[15] MVPN: add c-route, type 6, bs len 0 asn=65001, rd=1:1,
*Feb 1 21:13:19.033: source=10.2.255.255/4,
*Feb 1 21:13:19.033: group=226.1.1.1/4,
*Feb 1 21:13:19.033: nexthop=172.16.254.4,
*Feb 1 21:13:19.033: len left = 0
*Feb 1 21:13:19.033: BGP[14] MVPN umh lookup: vrfid 2, source 10.2.255.255
*Feb 1 21:13:19.033: BGP[4] MVPN umh lookup: vrfid 2, source 10.2.255.255, net 1:1:10.2.255.255/32, 1:1
*Feb 1 21:13:19.034: BGP(15): skip vrf default table RIB route [6][1:1][65001][10.2.255.255/32][226.1.1
*Feb 1 21:13:19.034: BGP(15): 172.16.255.1 NEXT_HOP self is set for sourced RT Filter for net [6][1:1][
*Feb 1 21:13:19.034: BGP(15): (base)
172.16.255.1 send UPDATE
 (format) [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22, next 172.16.255.6, metric 0, path Local, e
<-- Advertise to RR
 (
172.16.255.1)
```

第3步和第4步(Leaf-01)：从FHR角度，验证S、G创建和注册事件（S、G创建和注册几乎同时发生）

3.数据流量开始于FHR VTEP并创建S、G。此处适用“未检测到的组播源”部分中记录的要求。

4.枝叶–01通过其PIM隧道向RP执行源注册

```
<#root>
Leaf-01#
```

debug ip pim vrf green 226.1.1.1

PIM debugging is on

Leaf-01#

debug ip mrouting vrf green 226.1.1.1

IP multicast routing debugging is on

Debugs for PIM and Mroute show creation of S,G and PIM register encap event

Jan 29 18:18:37.602: PIM(2): Building Periodic (,G) Join / (S,G,RP-bit) Prune message for 226.1.1.1
*Jan 29 18:18:58.426:

MRT(2): (10.1.101.11,226.1.1.1), RPF install from /0.0.0.0 to Vlan101/10.1.101.11<-- S,G is creation me

*Jan 29 18:18:58.427:

PIM(2): Adding register encap tunnel (Tunnel4) as forwarding interface of (10.1.101.11, 226.1.1.1). <--

Jan 29 18:18:58.427: MRT(2): Set the F-flag for (, 226.1.1.1)
*Jan 29 18:18:58.427: MRT(2): Set the F-flag for (10.1.101.11, 226.1.1.1)
*Jan 29 18:18:58.428:

MRT(2): Create (10.1.101.11,226.1.1.1), RPF (Vlan101, 10.1.101.11, 0/0) <-- S,G is creation message (MR

*Jan 29 18:18:58.428: MRT(2): Set the T-flag for (10.1.101.11, 226.1.1.1)

Tunnel 4 is PIM Register tunnel (Encap: encapsulate in tunnel to RP)

Leaf-01#

sh int tunnel4

Tunnel4 is up, line protocol is up
  Hardware is Tunnel
  Description:

Pim Register Tunnel (Encap) for RP 10.2.255.255 on VRF green <-- VRF green for Leaf-02 RP

  Interface is unnumbered.

Using address of Loopback901 (10.1.255.1) <-- Local Loopback

S,G is created when Source sends data traffic

Leaf-01#

sh ip mroute vrf green 226.1.1.1

IP Multicast Routing Table
<...snip...>
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires

```
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 00:00:16/stopped, RP 10.2.255.255, flags: SPF
  Incoming interface: Vlan901, RPF nbr 172.16.254.4
  Outgoing interface list: Null


(10.1.101.11, 226.1.1.1)

, 00:00:16/00:02:47, flags: FTGqx


Incoming interface: Vlan101

,

RPF nbr 10.1.101.11

,

Registering <-- S,G created, in Register state, RPF IP is the /32 host prefix for this source



 Outgoing interface list:



Vlan901

, Forward/Sparse, 00:00:16/00:02:43

<-- OIF is the L3VNI SVI


#### Checking S,G in Hardware ###

Leaf-01#

sh platform software fed switch active ip mfib vrf green 226.1.1.1/32 10.1.101.11 de

MROUTE ENTRY

vrf 2

 (10.1.101.11, 226.1.1.1/32)

<-- VRF 2 is the ID for vrf green


  HW Handle: 140213987784872 Flags: {Svl}

RPF interface: Vlan101

(59)): SVI

<-- RPF is Direct connected on a Local Subnet


  HW Handle:140213987784872 Flags:A
Number of OIF: 2
Flags: 0x4

Pkts : 336         <-- packets that used this adjacency (similar to mfib command, but shown at the FED l


OIF Details:
```

Vlan101 A <-- Accept interface is programmed correctly

Vlan901 F {Remote} <-- Forward interface is L3VNI SVI

(Adj: 0x5f) <-- Validate this Adj

Htm: 0x7f861cf071b8 Si: 0x7f861cf04838 Di: 0x7f861cf097a8 Rep_ri: 0x7f861ceecb38

Check ADJ 0x5f for next hop details

Leaf-01#

sh platform software fed switch active ip adj

IPV4 Adj entries
dest if_name dst_mac si_hdl ri_hdl pd_flags

adj_id

Last-modified
---- ------- ------- ------ ------ -------- ----- ---

239.1.1.1

nve1.VNI50901

4500.0000.0000 0x7f861ce659b8 0x7f861ce65b68 0x60

0x5f

2021/01/29 17:07:06.568

Dest = MDT default group 239.1.1.1

Outgoing Interface = Nve1 using L3 VNI 50901

第4步（枝叶–02）：从RP角度，确认源注册到达RP，并创建S、G。

<#root>

PIM debugs showing PIM register event

Leaf-02#

debug ip pim vrf green 226.1.1.1

PIM debugging is on

Jan 29 18:21:35.500: PIM(2): Building Periodic (,G) Join / (S,G,RP-bit) Prune message for 226.1.1.1

***Jan 29 18:21:35.500: PIM: rp our address <-- Leaf-02 is the RP**

***Jan 29 18:21:41.005: PIM(2): Received v2 Register on Vlan901 from 10.1.255.1 <--- IP of Lo901 on Leaf-0**

***Jan 29 18:21:41.005: for 10.1.101.11, group 226.1.1.1**

***Jan 29 18:21:41.006: PIM(2): Adding register decap tunnel (Tunnel4) as accepting interface of (10.1.101**

*Jan 29 18:21:41.008: PIM(2): Upstream mode for (10.1.101.11, 226.1.1.1) changed from 1 to 2

Tunnel 4 is PIM Register tunnel (decap)

Leaf-02#

sh int tunnel 4

Tunnel4 is up, line protocol is up
  Hardware is Tunnel
  Description:

Pim Register Tunnel (Decap) for RP 10.2.255.255 on VRF green <-- decap side of register tunnel

  Interface is unnumbered.

Using address of Loopback255 (10.2.255.255) <-- RP IP

Mroute debugs show pim Register triggering S,G

Leaf-02#

debug ip mrouting vrf green 226.1.1.1

IP multicast routing debugging is on

*Jan 29 20:44:31.483: MRT(2):

(10.1.101.11,226.1.1.1)

,

RPF install from /0.0.0.0 to Vlan901/172.16.254.3 <-- RPF is to Leaf-01

*Jan 29 20:44:31.485: MRT(2):

Create (10.1.101.11,226.1.1.1), RPF (Vlan901, 172.16.254.3, 200/0) <-- Create the S,G

*Jan 29 20:44:33.458: MRT(2):

Set the T-flag for (10.1.101.11, 226.1.1.1) **<-- Set SPT bit for S,G**

S,G is created and traffic is now sent along the *,G shared tree
Leaf-02#sh ip mroute vrf green

```
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, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 00:05:49/stopped, RP 10.2.255.255, flags:
```

SGx <-- Sparse, Received BGP C-Mroute

Incoming interface: Null, RPF nbr 0.0.0.0 **<-- RP is us (Incoming Interface Null with**

```
  Outgoing interface list:
    Vlan901, Forward/Sparse, 00:05:49/stopped

(
```

10.1.101.11, 226.1.1.1

```
), 00:01:22/00:01:41, flags:
```

PTXgx <-- Pruned, SPT bit, Sent BGP C-Mroute

```
  Incoming interface: Vlan901,
```

RPF nbr 172.16.254.3 **<-- Leaf-01 is RPF next hop**

```
  Outgoing interface list: Null
```

第5步（枝叶–02）：RP有接收器，因此立即创建了第7类MVPN源树加入路由

<#root>

Leaf-02#

sh ip mroute vrf green 226.1.1.1

<...snip...>

(*, 226.1.1.1)

, 00:02:22/00:00:37, RP 10.2.255.255, flags: SGx
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:

    Vlan901, Forward/Sparse, 00:02:22/00:00:37   <-- L3 VNI is populated from Receiver BGP Type-6 join

Debugs showing Type-7 creation from RP

Leaf-02#

debug bgp ipv4 mvpn updates

BGP updates debugging is on for address family: MVPNv4 Unicast

Leaf-02#

debug bgp ipv4 mvpn updates events

BGP update events debugging is on for address family: MVPNv4 Unicast

*Jan 29 18:21:41.008: BGP[15]

MVPN: add c-route, type 7

, bs len 0 asn=0, rd=1:1,
*Jan 29 18:21:41.008:

source=10.1.101.11/4,

*Jan 29 18:21:41.008:

group=226.1.1.1/4,

*Jan 29 18:21:41.008:

nexthop=172.16.254.3

,

<-- Leaf-01 Global next hop

*Jan 29 18:21:41.008: len left = 0
*Jan 29 18:21:41.008: BGP[14] MVPN umh lookup: vrfid 2, source 10.1.101.11
*Jan 29 18:21:41.008: BGP[4] MVPN umh lookup: vrfid 2, source 10.1.101.11, net 1:1:10.1.101.11/32, 1:1:

0x10B:172.16.255.3:2

,

<-- This is the VRI picked up from the EVPN Type-2

*Jan 29 18:21:41.009: BGP:

MVPN(15) create local route [7][172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22

*Jan 29 18:21:41.009:

BGP[15] MVPN: add c-route, type 7, bs len 0 asn=65001, rd=1:1,

*Jan 29 18:21:41.009: source=10.1.101.11/4,
*Jan 29 18:21:41.009: group=226.1.1.1/4,
*Jan 29 18:21:41.009: nexthop=172.16.254.3,
*Jan 29 18:21:41.009: len left = 0
*Jan 29 18:21:41.009: BGP[14] MVPN umh lookup: vrfid 2, source 10.1.101.11
*Jan 29 18:21:41.009: BGP[4] MVPN umh lookup: vrfid 2, source 10.1.101.11, net 1:1:10.1.101.11/32, 1:1:

Type-7 Locally created on RP and sent to Source Leaf-01

Leaf-02#

sh bgp ipv4 mvpn all

BGP table version is 81, local router ID is 172.16.255.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path, L long-lived-stale,
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: 172.16.254.3:101 <-- Note the VRI is learnt from Leaf-01

 *>

[7][172.16.254.3:101]

[65001]

[10.1.101.11/32][226.1.1.1/32]

/22

<-- [7] = type-7 for this S,G / VRI 172.16.254.3:101 learned from Leaf-01

  0.0.0.0

32768

 ?

<-- 0.0.0.0 locally originated

with local Weight

第6步(Leaf-01):Source Leaf-01接收并安装MVPN Route-Type 7。（将L3 VNI SVI安装为S、G的转发接口）

<#root>

Received Type-7 from Leaf-02 RP

Leaf-01#

debug bgp ipv4 mvpn updates

BGP updates debugging is on for address family: MVPNv4 Unicast
Leaf-01#

debug bgp ipv4 mvpn updates events

BGP update events debugging is on for address family: MVPNv4 Unicast

*Jan 29 18:18:58.457:

BGP(15): 172.16.255.1 rcvd UPDATE w/ attr: nexthop 172.16.255.4, origin ?, localpref 100, metric 0, orig

*Jan 29 18:18:58.457: BGP(15): 172.16.255.1

rcvd [7][172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22 <-- Received |

*Jan 29 18:18:58.457: BGP(15): skip vrf default table RIB route [7][172.16.254.3:101][65001][10.1.101.1
*Jan 29 18:18:58.458: BGP(15): add RIB route (0:0)[7][1:1][65001][10.1.101.11/32][226.1.1.1/32]/22

PIM updated by MVPN to install L3 VNI in Outgoing Interface List

Leaf-01#

debug ip pim vrf green 226.1.1.1

PIM debugging is on
Leaf-01#

debug ip mrouting vrf green 226.1.1.1

IP multicast routing debugging is on

*Jan 29 18:18:58.458: PIM(2):

Join-list: (10.1.101.11/32, 226.1.1.1), S-bit set, BGP C-Route

```
*Jan 29 18:18:58.459: MRT(2):
WAVL Insert VxLAN interface: Vlan901 in (10.1.101.11,226.1.1.1) Next-hop: 239.1.1.1 VNI 50901 Successful

*Jan 29 18:18:58.459: MRT(2): set min mtu for (10.1.101.11, 226.1.1.1) 18010->9198
*Jan 29 18:18:58.460:
MRT(2): Add Vlan901/239.1.1.1/50901 to the olist of (10.1.101.11, 226.1.1.1), Forward state - MAC not bu

*Jan 29 18:18:58.460: PIM(2): Add Vlan901/0.0.0.0 to (10.1.101.11, 226.1.1.1), Forward state, by BGP SG
*Jan 29 18:18:58.460: MRT(2): Add Vlan901/239.1.1.1/50901to the olist of (10.1.101.11, 226.1.1.1), Forwa
```

第7步（枝叶–01）：枝叶–01为S、G通告MVPN源A-D类型5

```
<#root>

Leaf-01#
debug bgp ipv4 mvpn updates

BGP updates debugging is on for address family: MVPNv4 Unicast
Leaf-01#
debug bgp ipv4 mvpn updates events

BGP update events debugging is on for address family: MVPNv4 Unicast

*Jan 29 18:18:58.461: BGP(15): nettable_walker
[5][1:1][10.1.101.11][226.1.1.1]/18 route sourced locally <-- BGP determines route is local to Leaf-01

*Jan 29 18:18:58.461: BGP(15): delete RIB route (0:0)[5][1:1][10.1.101.11][226.1.1.1]/18
*Jan 29 18:18:58.461: BGP(15): 172.16.255.1 NEXT_HOP self is set for sourced RT Filter for net [5][1:1]
*Jan 29 18:18:58.461: BGP(15): (base) 172.16.255.1
send UPDATE (format) [5][1:1][10.1.101.11][226.1.1.1]/18, next 172.16.255.3, metric 0, path Local, exter
```

第8步（枝叶–03）：接收方VTEP获取类型5并安装源A-D路由用于S、G

```
<#root>

Leaf-03#
debug bgp ipv4 mvpn updates

BGP updates debugging is on for address family: MVPNv4 Unicast
Leaf-03#
debug bgp ipv4 mvpn updates events

BGP update events debugging is on for address family: MVPNv4 Unicast
```

```
*Jan 29 19:18:53.318: BGP(15): 172.16.255.1 rcvd UPDATE w/ attr: nexthop 172.16.255.3, origin ?, localp
*Jan 29 19:18:53.319: BGP(15): 172.16.255.1 rcvd [5][1:1][10.1.101.11][226.1.1.1]/18   <-- Type-5 Receiv

*Jan 29 19:18:53.319: BGP(15): skip vrf default table RIB route [5][1:1][10.1.101.11][226.1.1.1]/18


Leaf-03#
sh bgp ipv4 mvpn all route-type 5 10.1.101.11 226.1.1.1
...or you can also use:
Leaf-03#
sh bgp ipv4 mvpn detail [5][1:1][10.1.101.11][226.1.1.1]/18


BGP routing table entry for
[5][1:1][10.1.101.11][226.1.1.1]/18
, version 41
<-- Type-5 A-D route from Leaf-01


Paths: (2 available, best #2, table MVPNv4-BGP-Table, not advertised to EBGP peer)
Flag: 0x100
Not advertised to any peer
Refresh Epoch 1
Local


 172.16.255.3
 (metric 3) from 172.16.255.1 (172.16.255.1)
<-- Leaf-01 IP


    Origin incomplete, metric 0, localpref 100, valid, internal, best
    Community: no-export
     Extended Community: RT:1:1


Originator: 172.16.255.3
, Cluster list: 172.16.255.1
     rx pathid: 0, tx pathid: 0x0
     Updated on Jan 29 2021 19:18:53 UTC
```

第9步（枝叶–03）：创建S，G，枝叶–03发送MVPN类型–7以加入SPT树，并开始接受流量

<#root>

```
debug ip mrouting vrf green 226.1.1.1
debug bgp ipv4 mvpn updates
debug bgp ipv4 mvpn updates events


### Debug of Mrouting shows S,G create and call to BGP to create Type-7 BGP S,G join ###
```

*Feb 12 19:34:26.045:

MRT(2):

(10.1.101.11,226.1.1.1), RPF install from /0.0.0.0 to Vlan901/172.16.254.3 <-- RPF check done as first c

*Feb 12 19:34:26.046:

MRT(2):

Create (10.1.101.11,226.1.1.1), RPF (Vlan901, 172.16.254.3, 200/0) <-- RPF successful Creating S,G

*Feb 12 19:34:26.047: MRT(2): WAVL Insert interface: Vlan102 in (10.1.101.11,226.1.1.1) Successful
*Feb 12 19:34:26.047: MRT(2): set min mtu for (10.1.101.11, 226.1.1.1) 18010->9198
*Feb 12 19:34:26.047: MRT(2): Set the T-flag for (10.1.101.11, 226.1.1.1)
*Feb 12 19:34:26.048:

MRT(2):

Add Vlan102/226.1.1.1 to the olist of (10.1.101.11, 226.1.1.1)

, Forward state - MAC not built

<-- Adding Vlan102 Receiver SVI into OIF list

*Feb 12 19:34:26.048:

MRT(2): Set BGP Src-Active for (10.1.101.11, 226.1.1.1) <-- Signaling to BGP that this Source is seen a

BGP Type-7 created

Leaf-03#

sh bgp ipv4 mvpn all

Route Distinguisher:

172.16.254.3:101   <-- VRI Route Distinguisher

*>

[7]

[

172.16.254.3:101]

[65001]

[10.1.101.11/32][226.1.1.1/32]

/22

<-- Type [7], VRI, S,G info

0.0.0.0

```
                              32768 ?
```

<-- created locally

```
Leaf-03#
```

sh ip mroute vrf green 226.1.1.1 10.1.101.11

```
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, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.101.11, 226.1.1.1), 00:08:41/00:02:13,
```

flags: TgQ <-- SPT bit, Sent MVPN type-7, Received MVPN type-5

Incoming interface: Vlan901, RPF nbr 172.16.254.3   <-- Receive from L3VNI via Leaf-01 IP next hop

```
  Outgoing interface list:
```

Vlan102, Forward/Sparse, 00:08:41/00:02:22   <-- Send to host in Vlan 102

第10步(Leaf-01):Leaf-01从Leaf-03接收并安装MVPN Type-7

<#root>

debug bgp ipv4 mvpn updates

```
debug bgp ipv4 mvpn updates events


### Type-7 Received from Leaf-03 VTEP and installed into RIB ###

*Feb 12 19:55:29.000: BGP(15): 172.16.255.1
rcvd [7][172.16.254.3:101][65001][10.1.101.11/32][226.1.1.1/32]/22   <-- Type-7 from Leaf-03

*Feb 12 19:55:29.000: BGP(15): skip vrf default table RIB route [7][172.16.254.3:101][65001][10.1.101.1
*Feb 12 19:55:29.000: BGP(15): add RIB route (0:0)[7][1:1][65001][10.1.101.11/32][226.1.1.1/32]/22
```

场景4：交换矩阵外部的RP（从IP空间的边界枝叶02导入RP）

此场景与场景2基本相同。交换矩阵整体使用单个RP。区别在于RP IP必须从非交换矩阵IP空间导入到交换矩阵并通告到BGP。

本节显示与方案3的差异。相同的步骤和方法仅在场景3中说明

- 由于BGP和PIM操作相同，请参阅从方案3验证此方案所需的事件顺序

网络图

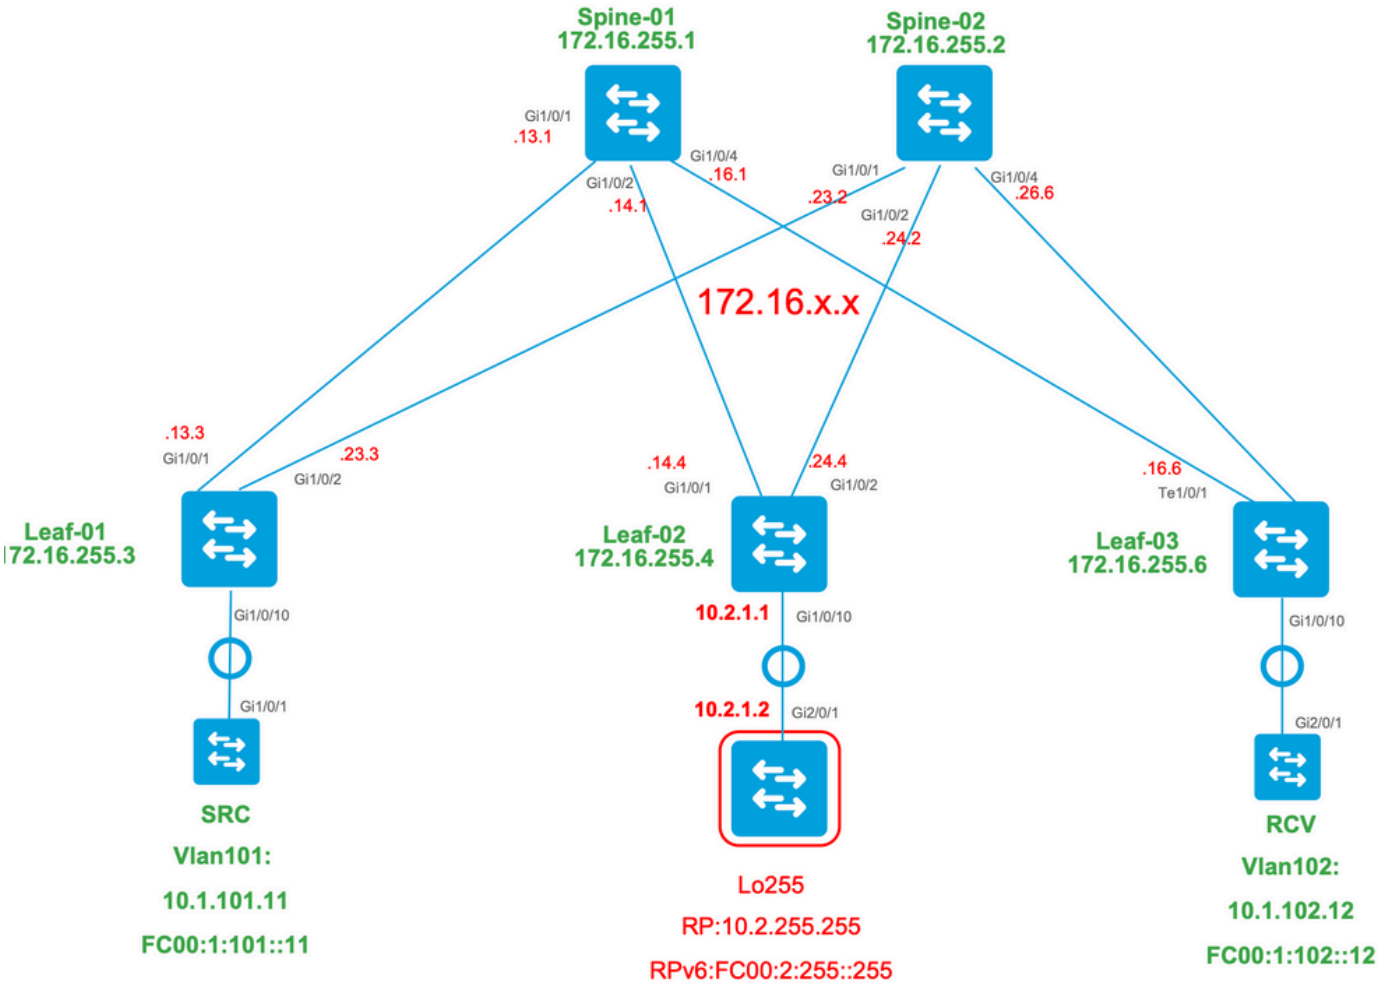

验证边界交换机从IP导入到交换矩阵

此设计与场景3的主要区别是需要首先将RP IP从IP空间导入EVPN。

Border需要包含特定命令来导入/导出交换矩阵和IP空间：

- route-target <value> spitching命令在VRF配置部分
- 在BGP vrf地址系列下通告l2vpn evpn

验证(Leaf-02)：配置

<#root>

```
Leaf-02#

sh run vrf green
Building configuration...

Current configuration : 1533 bytes

vrf definition green
rd 1:1
!
address-family ipv4
mdt auto-discovery vxlan
mdt default vxlan 239.1.1.1
```

```
mdt overlay use-bgp
route-target export 1:1

route-target import 1:1

route-target export 1:1 stitching    <-- BGP-EVPN fabric redistributes the stitching routes between the

route-target import 1:1 stitching

exit-address-family
Leaf-02#
sh run | sec router bgp

address-family ipv4 vrf green       <--- BGP VRF green address-family

advertise l2vpn evpn                <--- Use the 'advertise l2vpn evpn' command and 'export stitching' F

redistribute connected
redistribute static
redistribute ospf 2 match internal external 1 external 2   <-- Learning via external OSPF neighbor in VR

exit-address-family
```

验证(Leaf-02)：前缀导入和通告

<#root>

```
debug bgp vpnv4 unicast updates

debug bgp vpnv4 unicast updates events

debug bgp l2vpn evpn updates

debug bgp l2vpn evpn updates events

*Feb 15 15:30:54.407: BGP(4): redist event (1) request for 1:1:10.2.255.255/32
```

```
*Feb 15 15:30:54.407: BGP(4) route 1:1:10.2.255.255/32 gw-1 10.2.1.2 src_proto (ospf) path-limit 1
*Feb 15 15:30:54.407: BGP(4): route 1:1:10.2.255.255/32 up
*Feb 15 15:30:54.407: bgp_ipv4set_origin: redist 1, opaque 0x0, net 10.2.255.255
*Feb 15 15:30:54.407: BGP(4): sourced route for 1:1:10.2.255.255/32 path 0x7FF8065EB9C0 id 0 gw 10.2.1.
*Feb 15 15:30:54.408: BGP(4): redistributed route 1:1:10.2.255.255/32 added gw 10.2.1.2
*Feb 15 15:30:54.408: BGP: topo green:VPNv4 Unicast:base Remove_fwdroute for 1:1:10.2.255.255/32

*Feb 15 15:30:54.408: BGP(4): 1:1:10.2.255.255/32 import vpn re-orig or locally sourced or learnt from C



*Feb 15 15:30:54.409: BGP(10): update modified for [5][1:1][0][32][10.2.255.255]/17


*Feb 15 15:30:54.409: BGP(10): 172.16.255.1

NEXT_HOP set to vxlan local vtep-ip 172.16.254.4



for net [5][1:1][0][32][10.2.255.255]/17   <-- Set NH to Leaf-02 loopback


*Feb 15 15:30:54.409: BGP(10): update modified for [5][1:1][0][32][10.2.255.255]/17

*Feb 15 15:30:54.409: BGP(10): (base) 172.16.255.1 send UPDATE (format) [5][1:1][0][32][10.2.255.255]/17
<-- BGP EVPN Type update created from Non-fabric Imported prefix and sent to RR


### Verify the NLRI is learned and Imported on Border Leaf-02 ###

Leaf-02#

sh bgp vpnv4 unicast all


BGP table version is 39, local router ID is 172.16.255.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
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 green)

AF-Private Import to Address-Family: L2VPN E-VPN, Pfx Count/Limit: 3/1000  <-- Prefix Import details. (N


 *>

10.2.255.255/32  10.2.1.2                2         32768 ?          <-- Locally redistributed, Next hop


Leaf-02#

sh bgp l2vpn evpn all route-type 5 0 10.2.255.255 32

...or you can also use:
Leaf-02#

sh bgp l2vpn evpn detail [5][1:1][0][32][10.2.255.255]/17

BGP routing table entry for

[5][1:1][0][32][10.2.255.255]
```

```
/17, version 69
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     2
  Refresh Epoch 1


 Local, imported path from base



10.2.1.2 (via vrf green) from 0.0.0.0 (172.16.255.4)                <-- Imported to EVPN Fabric table fr


      Origin incomplete, metric 2, localpref 100, weight 32768, valid, external, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0,

local vtep: 172.16.254.4, VNI Label 50901,

MPLS VPN Label 17

<-- VTEP IP of Leaf-02, L3VNI label


      Extended Community: RT:1:1 OSPF DOMAIN ID:0x0005:0x000000020200
        MVPN AS:65001:0.0.0.0

MVPN VRF:172.16.255.4:2

 ENCAP:8

<-- MVPN VRI created


        Router MAC:7C21.0DBD.9548 OSPF RT:0.0.0.0:2:0
        OSPF ROUTER ID:10.2.255.255:0
      rx pathid: 0, tx pathid: 0x0
      Updated on Feb 15 2021 15:30:54 UTC
```

验证（枝叶–02）：到RP的边界路径

<#root>

```
Leaf-02#sh ip mroute vrf green

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, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
```

```
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode


(*, 226.1.1.1)

, 2d21h/stopped,

RP 10.2.255.255

, flags: SJGx

<-- *,G for group and Non-fabric RP IP



 Incoming interface: Vlan2001

,

RPF nbr 10.2.1.2            <-- RPF neighbor is populated for IP next hop outside VxLAN


  Outgoing interface list:

Vlan901, Forward/Sparse, 01:28:47/stopped              <-- Outgoing is L3VNI SVI
```

场景5：数据MDT

验证MDT数据组

MDT数据组与其他MDT默认组类似，其中要封装的TRM的外部隧道组。但是，与MDT默认设置不同，此组将只有VTEP的加入此树，如果他们对TRM组有感兴趣的接收器。

所需配置

<#root>

```
vrf definition green
rd 1:1
!
address-family ipv4
mdt auto-discovery vxlan
mdt default vxlan 239.1.1.1

mdt data vxlan 239.1.2.0 0.0.0.255  <-- Defines MDT Data underlay group address range



mdt data threshold 1



<-- Defines the threshold before cutting over to the Data group (In Kilobits per second)
```

```
mdt overlay use-bgp spt-only
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
!
```

检查MDT组是否已在源端正确编程

- MDT组的传入接口是源端环回
- MDT组的传出接口是底层接口

验证Leaf-01:MRIB/MFIB中的MDT mroute正确

<#root>

Leaf-01#

sh ip mroute 239.1.2.0 172.16.254.3

<snip>

(172.16.254.3, 239.1.2.0)

, 00:01:19/00:02:10, flags: FT
  Incoming interface:

Loopback1

, RPF nbr

0.0.0.0

<-- IIF is local loopback with 0.0.0.0 RPF indicating local

  Outgoing interface list:

    TenGigabitEthernet1/0/1

, Forward/Sparse, 00:01:19/00:03:10

<-- OIF is the underlay uplink

Leaf-01#

sh ip mfib 239.1.2.0 172.16.254.3

<snip>
(172.16.254.3,239.1.2.0) Flags: HW
  SW Forwarding: 2/0/828/0, Other: 0/0/0

  HW Forwarding: 450/2/834/13

, Other: 0/0/0

<-- Hardware counters indicate the entry is operating in hardware and forwarding packets

Null0 Flags: A <-- Null0 (Originated locally)

TenGigabitEthernet1/0/1

Flags: F NS

<-- OIF is into the Underlay (Global routing table)

Pkts: 0/0/0 Rate: 0 pps

验证MDT组的Leaf-01:FED条目

<#root>

Leaf-01#

show platform software fed switch active ip mfib 239.1.2.0/32 172.16.254.3 detail <-- The detail option

MROUTE ENTRY

vrf 0 (172.16.254.3, 239.1.2.0/32) <-- vrf 0 = global for this MDT Data S,G pair

HW Handle: 140028029798744 Flags:

RPF interface: Null0

(1)):

<-- Leaf-01 is the Source(Null0)

HW Handle:140028029798744 Flags:A
Number of OIF: 2
Flags: 0x4 Pkts : 570

<-- Packets that used this adjacency (similar to the mfib command, but shown at the FED layer)

OIF Details:

TenGigabitEthernet1/0/1 F NS <-- The Underlay Outgoing Interface and F-Forward flag

Null0 A <-- The Incoming Interface is local loopback1 and A-Acc

Htm: 0x7f5ad0fa48b8 Si: 0x7f5ad0fa4258

Di: 0x7f5ad0fa8948

Rep_ri: 0x7f5ad0fa8e28

<--The DI (dest index) handle

DI details

```
----------
Handle:0x7f5ad0fa8948 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICA
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles:
```

index0:0x536e

```
 mtu_index/l3u_ri_index0:0x0
```

index1:0x536e

```
 mtu_index/l3u_ri_index1:0x0 index2:0x536e mtu_index/l3u_ri_index2:0x0 index3:0x536e mtu_index/l3u_ri_i
```

```
<snip>

Brief Resource Information (ASIC_INSTANCE# 3)
---------------------------------------
```

Destination index = 0x536e

```
pmap = 0x00000000 0x00000001
```

pmap_intf : [TenGigabitEthernet1/0/1] <--FED has the correct programing of the OIF

```
==============================================================
```

检查MDT组是否已在接收器端正确编程

- MDT组的传入接口是返回源端环回的RPF接口
- MDT组的传出接口是Encap/Decap隧道接口

验证Leaf-02:MRIB/MFIB中的MDT mroute正确

<#root>

```
Leaf-03#
```

sh ip mroute 239.1.2.0 172.16.254.3   <-- This is the Global MDT Data Group

```
<snip>
(
```

172.16.254.3, 239.1.2.0

```
), 00:06:12/00:02:50, flags: JTx
```

<-- Source is Leaf-01 Loopback1 IP

```
  Incoming interface: TenGigabitEthernet1/0/1, RPF nbr 172.16.26.2
  Outgoing interface list:
```

Tunnel0

```
, Forward/Sparse, 00:06:12/00:02:47
<-- Decap Tunnel

Leaf-03#
sh ip mfib 239.1.2.0 172.16.254.3

<snip>
Default                                                    <-- Global Routing Table

  (
172.16.254.3,239.1.2.0
) Flags: HW
    SW Forwarding: 2/0/828/0, Other: 0/0/0
    HW Forwarding: 760/2/846/13
, Other: 0/0/0
<-- Hardware counters indicate the entry is operating in hardware and forwarding packets

    TenGigabitEthernet1/0/1 Flags: A                       <-- Accept via Underlay (Global) interface

    Tunnel0, VXLAN Decap Flags: F NS                       <-- Forward to VxLAN Decap Tunnel
      Pkts: 0/0/2 Rate: 0 pps
```

验证MDT组的枝叶02:FED条目

<#root>

```
Leaf-03#
show platform software fed switch active ip mfib 239.1.2.0/32 172.16.254.3 detail

MROUTE ENTRY
vrf 0 (172.16.254.3, 239.1.2.0/32)  <-- vrf 0 = global for this MDT Data S,G pair
  HW Handle: 140592885196696 Flags:
RPF interface: TenGigabitEthernet1/0/1
(55)):
<-- RPF Interface to 172.16.254.3

  HW Handle:140592885196696 Flags:A
```

```
Number of OIF: 2
Flags: 0x4
```

Pkts : 800 <-- packets that used this adjacency (similar to mfib command, but

```
OIF Details:
```

TenGigabitEthernet1/0/1 A <-- Accept MDT packets from this interface

Tunnel0 F NS <-- Forward to Decap Tunnel to remove VxLAN header

(Adj: 0x3c) <-- Tunnel0 Adjacency

```
Htm: 0x7fde54fb7d68 Si: 0x7fde54fb50d8 Di: 0x7fde54fb4948 Rep_ri: 0x7fde54fb4c58

<snip>
```

RI details <-- Rewrite Index is used for VxLAN decapsulation

```
----------
Handle:0x7fde54fb4c58 Res-Type:ASIC_RSC_RI_REP Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MUL
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x1a mtu_index/l3u_ri_index0:0x0 index1:0

Brief Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
ASIC# 0
Replication list :
------------------

Total #ri : 6
Start_ri : 26
Common_ret : 0

Replication entry
```

rep_ri 0x1A

```
 #elem = 1
0)
```

ri[0]=0xE803

```
 Dynamic port=88ri_ref_count:1 dirty=0
<snip>

Leaf-03#
```

show platfomr software fed switch active fwd-asic resource asic all rewrite-index range 0xE803 0XE803

```
ASIC#:0 RI:59395
```

Rewrite_type

```
:AL_RRM_REWRITE_L2_PAYLOAD_
```

IPV4_EVPN_DECAP

```
(118) Mapped_rii:LVX_EVPN_DECAP(143)
<snip>
```

Debug MDT Data Group

使用MVPN调试检查数据MDT转换事件

源端VTEP

<#root>

```
Leaf#
debug mvpn

<snip>
*Mar 27 12:12:11.115: MVPN: Received local withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Ty
*Mar 27 12:12:11.115: MVPN: Sending BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nh 0.0.0.0, W
*Mar 27 12:12:11.115: MVPN: Route Type 5 deleted [(10.1.101.11, 239.1.1.1), nh 0.0.0.0] rd:1:1 send:1
*Mar 27 12:12:11.115: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNK
*Mar 27 12:12:11.115: MVPN: Received BGP withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Type
*Mar 27 12:13:00.430: MVPN: Received local route update for (10.1.101.11, 239.1.1.1) with RD: 1:1, Rout
*Mar 27 12:13:00.431: MVPN: Route Type 5 added [(10.1.101.11, 239.1.1.1), nh 0.0.0.0] rd:1:1 send:1
*Mar 27 12:13:00.431: MVPN: RP 10.2.255.255 updated in newly created route
*Mar 27 12:13:00.431: MVPN: Sending BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nh 0.0.0.0, O
*Mar 27 12:13:00.431: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNK
*Mar 27 12:13:00.431: MVPN: Received BGP withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Type
*Mar 27 12:13:17.151:
MVPN(green[AF_IPv4]): Successfully notified nve fordatamdt adjacency create 239.1.2.0


<-- Notify NVE about creating DATA MDT

*Mar 27 12:13:17.151:
MVPN: Received local update <104:0x00:0>(172.16.254.3, 239.1.2.0) next_hop:0.0.0.0 router_id:172.16.255.

*Mar 27 12:13:17.151:
MVPN: LSM AD route added [(10.1.101.11,239.1.1.1) : <104:0x00:0>(172.16.254.3, 239.1.2.0)] orig:172.16.2

*Mar 27 12:13:17.151:
MVPN(green[AF_IPv4]): Sending VxLAN BGP AD prefix=[3:172.16.255.3 1:1 : (10.1.101.11,239.1.1.1)] len=23,

*Mar 27 12:13:17.151:
MVPN(green[AF_IPv4]): Originate VxLAN BGP AD rt:3

*Mar 27 12:13:17.151:
MVPN(green[AF_IPv4]): VXLAN MDT-Data, node added for (10.1.101.11,239.1.1.1) MDT: 239.1.2.0
```

```
Leaf-01#
```

接收端VTEP

<#root>

```
Leaf#
```

debug mvpn

```
<snip>
*Mar 27 12:27:54.920: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: 172
*Mar 27 12:27:54.920: MVPN: Received BGP route update for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route
*Mar 27 12:27:54.920: MVPN: Route Type 5 found [(10.1.101.11, 239.1.1.1), nh 172.16.255.3]rd:1:1 send:0
*Mar 27 12:27:54.920: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNK
*Mar 27 12:27:54.920: MVPN: Received BGP withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Type
*Mar 27 12:27:54.920: MVPN: Route Type 5 deleted [(10.1.101.11, 239.1.1.1), nh 172.16.255.3] rd:1:1 sen
*Mar 27 12:28:27.648: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNK
*Mar 27 12:28:27.657: MVPN: Received BGP withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Type
*Mar 27 12:28:44.235: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: 172
*Mar 27 12:28:44.235: MVPN: Received BGP route update for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route
*Mar 27 12:28:44.235: MVPN: Route Type 5 added [(10.1.101.11, 239.1.1.1), nh 172.16.255.3] rd:1:1 send:
*Mar 27 12:29:00.956: MVPN: Received BGP prefix=[3:172.16.255.3 1:1 : (10.1.101.11,239.1.1.1)] len=23,
*Mar 27 12:29:00.956: MVPN: Received BGP prefix=[3:172.16.255.3 1:1 : (10.1.101.11,239.1.1.1)] len=23,
*Mar 27 12:29:00.956:
```

MVPN: Received BGP update <104:0x00:50901>(172.16.254.3, 239.1.2.0) next_hop:172.16.255.3 router_id:172.

```
*Mar 27 12:29:00.956:
```

MVPN: LSM AD route added [(10.1.101.11,239.1.1.1) : <104:0x00:50901>(172.16.254.3, 239.1.2.0)] orig:172.

```
*Mar 27 12:29:00.957:
```

MVPN(green[AF_IPv4]): Activating PE (172.16.255.3, 1:1) ad route refcnt:1 control plane refcnt: 0

```
*Mar 27 12:29:00.958:
```

MVPN(green[AF_IPv4]): Successfully notified datamdt group for NVE (239.1.2.0, TRUE, FALSE)

```
*Mar 27 12:29:00.958: MVPN: Received BGP update <104:0x00:50901>(172.16.254.3, 239.1.2.0) next_hop:172.
Leaf-03#
```

故障排除

未检测到的组播源
在您了解组播流无法正常运行的原因之前，了解ARP和组播转发之间的关系非常重要

通常，当主机变为活动状态并发送流量时，ARP条目会通过常规的源检测程序完成。但是，对于组播源，
，FHR的L2平面处理此组播流量，而不解析源的ARP。

ARP完成在TRM功能中起着重要作用，原因有二。

1. 第一跳路由器的“直连”检查会调用FIB API，而后者又取决于ARP完成情况才能成功检查。如果通向组
向源的CEF邻接关系保持不完整，并且直接连接的check返回FALSE。
2. 源检测触发EVPN交换矩阵中的EVPN RT-2通告。此EVPN路由安装在接收器枝叶的L3RIB中，用作
，如果未检测到源，则无法找到(S，G)条目的RPF。在这种情况下，RPF将保持NULL或者在RIB中
果有）。

请确保ARP已解析且源在EVPN交换矩阵中可访问。

其他有用的调试

本节介绍有助于隔离TRM问题的其他调试

- debug mvpn（所有MVPN事件，请参阅场景2）
- debug ip|ipv6 pim <vrf>（PIM协议活动）
- debug ip mrib <vrf> trans（MRIB，经典PIM转换）
- debug ip mfib <vrf> pak|ps|fs(数据包转发| 进程交换| 快速交换)

交换矩阵外部的源和接收器

在某些情况下，源和/或接收器可以远离交换矩阵VTEP驻留一个或多个L3跳。

这是一个有效的设计，但会更改哪种EVPN路由类型承载VRI，以及负责在接收器VTEP创建连接的过程。

- 如果源位于交换矩阵外部，则入口VTEP通过PIM邻居（而不是直接连接）看到源，并向接收器VTEP
含在此类型5中。
- 如果接收器位于交换矩阵外部，则加入通过PIM加入IGMP进行。PIM加入中的信息用于创建MVPN类

eBGP多个AS（主干到主干）拓扑

在某些情况下，拓扑可能需要BGP将更新信息发送到另一个AS/交换矩阵。

BGP控制平面信息会经过最多30秒，然后组播开始工作。

- 这是因为默认eBGP通告间隔为30秒。
- 如果由于BGP更新的延迟而存在收敛时间较长的问题，则可以缩短eBGP通告间隔，以便更频繁
- 有关此计时器的详细信息，请参阅本文参考部分中的BGP配置指南。

eBGP inter-as需要一个额外命令

对MVPN地址系列路由使用inter-as关键字以跨BGP自治系统(AS)边界。

<#root>

```
Border-Leaf(config-vrf-af)#
mdt auto-discovery vxlan inter-as
```

对称L2VNI的注册隧道（FHR停滞在PIM注册状态）

如果VNI存在于FHR和其他VTEP上，则可能使FHR停滞在Register状态

这是因为PIM注册隧道源IP是AnyCast网关。当RP收到PIM寄存器时，它不知道哪个是发送寄存器停止的正
设备通用。

PIM注册隧道路由问题

(Leaf-01)这是实际FHR：向RP发送注册消息

<#root>

```
Leaf-01#sh ip pim vrf green tunnel
Tunnel5*
Type : PIM Encap
RP : 10.2.255.255
Source : 10.1.101.1 <-- Source of Register Tunnel
State : UP
Last event : Created (00:33:28)
```

(Leaf-03)：此VTEP（可能还有其他）包含与FHR相同的SVI和IP地址

<#root>

```
Leaf-03#sh ip pim vrf green tunnel
Tunnel4
Type : PIM Encap
RP : 10.2.255.255
Source : 10.1.101.1 <-- Source of Register Tunnel
State : UP
```

```
Last event : Created (00:11:53)
```

(Leaf-01):FHR在注册中仍然卡住（它不会收到来自RP的注册停止）

<#root>

```
Leaf-01#
```

show ip mroute vrf green 226.1.1.1 10.1.101.11

```
(10.1.101.11, 226.1.1.1), 02:02:19/00:02:22, flags: PFT
Incoming interface: Vlan101, RPF nbr 10.1.101.11,
```

Registering <-- Leaf-01 is stuck in register state

```
Outgoing interface list: Null
```

(Leaf-02)这是RP：在这种情况下，它还拥有与FHR相同的AnyCast IP，因此会向自身发送注册停止。

如果RP没有l2vni，但是有2个或3个其他vtep，注册停止可能会发送到错误的VTEP，因为RP无法选择正确

<#root>

```
Leaf-02#
```

sh ip route vrf green 10.1.101.1

```
Routing Table: green
Routing entry for 10.1.101.1/32
```

Known via "connected"

```
, distance 0, metric 0 (connected)
Routing Descriptor Blocks:
*
```

directly connected, via Vlan101 <-- Leaf-02 sees IP as Connected, and sends the Register-stop to itself

```
Route metric is 0, traffic share count is 1
```

(Leaf-02):RP上的调试显示RP将此路由作为本地连接路由的问题

<#root>

```
Leaf-02#
```

debug ip pim vrf green 226.1.1.1

```
PIM debugging is on
*May 26 17:33:15.797: PIM(2)[green]:
Received v2 Register on Vlan901 from 10.1.101.1 <-- Received from Leaf-01 with Source of 10.1.101.1

*May 26 17:33:15.797: PIM(2)[green]:
Send v2 Register-Stop to 10.1.101.1 for 10.1.101.11, group 226.1.1.1 <-- Sending Register-stop to FHR

*May 26 17:33:15.797: PIM(2)[green]:
Received v2 Register-Stop on Vlan101 from 10.2.255.255 <-- Leaf-02 receives its own Register-stop as th

*May 26 17:33:15.797: PIM(2)[green]:
for source 10.1.101.11, group 226.1.1.1 <-- S,G the Stop is for

*May 26 17:33:15.797: PIM(2)[green]:
Clear Registering flag to 10.2.255.255 for (10.1.101.11/32, 226.1.1.1) <-- Done with Register event

*May 26 17:33:17.801: PIM(2)[green]:
Received v2 Register on Vlan901 from 10.1.101.1 <-- Another Register messages from Leaf-01 and the even

*May 26 17:33:17.801: PIM(2)[green]: Send v2 Register-Stop to 10.1.101.1 for 10.1.101.11, group 226.1.1
*May 26 17:33:17.802: PIM(2)[green]: Received v2 Register-Stop on Vlan101 from 10.2.255.255
*May 26 17:33:17.802: PIM(2)[green]: for source 10.1.101.11, group 226.1.1.1
*May 26 17:33:17.802: PIM(2)[green]: Clear Registering flag to 10.2.255.255 for (10.1.101.11/32, 226.1.
```

PIM注册隧道路由问题解决方案

解决方案是在所有VTEP上使用唯一的环回IP，并使用本节中介绍的配置。

```
<#root>
Leaf-01#
sh run int lo 901

interface Loopback901
vrf forwarding green <-- Loopback is in the Tenant VRF

ip address 10.1.255.1
 255.255.255.255
<-- IP is unique to the VTEP

ip pim sparse-mode
```

```
Leaf-02(config)#
ip pim vrf green register-source loopback 901 <-- force the Register Source to use the Loopback


Leaf-01#
sh ip pim vrf green tunnel

Tunnel5
Type : PIM Encap     <-- Register Encapsulation tunnel


RP : 10.2.255.255    <-- RP IP is the Tunnel destination


Source : 10.1.255.1  <-- Loopback 901 is the Tunnel source

State : UP
Last event : Created (02:45:58)
Leaf-02#
show bgp l2vpn evpn all | beg 10.1.255.1

*>i
[5]
[1:1][0][32]
[10.1.255.1]
/17
          172.16.254.3
       0          100    0 ?
<-- Only one entry and next hop


to Leaf-01
```

相关信息

EVPN VxLAN TRM配置指南

EVPN VxLAN单播故障排除

MVPN配置指南17.3.x（Catalyst 9300交换机）

MVPN配置指南17.3.x（Catalyst 9500交换机）

BGP配置指南

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。