Identificar e Solucionar Problemas do EVPN VxLAN TRM nos Switches Catalyst 9000

Contents

Introdução

Este documento descreve como solucionar problemas com o TRM (Tenant Routed Multicast) sobre EVPN VxLAN.

Pré-requisitos

- Recomenda-se que você esteja familiarizado com o recurso Unicast EVPN VxLAN, BGP e MVPN (Multicast Virtual Private Network).
- Além disso, você deve entender como o multicast opera e os conceitos de multicast

Requisitos

Este guia supõe que os peers BGP, NVE já estão corretos. Se houver problemas com a ativação básica da EVPN VxLAN (falha de ping unicast, BGP, peers NVE inativos e assim por diante) consulte os guias de solução de problemas de BGP, EVPN, rota/switch conforme necessário.

Disponibilidade de recursos em cada versão de código

Versão	Recurso
17.1.1	TRMv4 com RP Anycast
17.3.1	TRMv4 com RP Externo ou RP Único
17.3.1	TRMv6 com RP Anycast
17.3.1	TRMv6 com RP externo ou RP único
17.3.1	TRMv4 com entrelaçamento MVPN (perfil11) com RP único no lado da malha
17.6.2 e 17.7.1	TRMv4 Data mdt com RP Anycast, RP Externo ou RP Único

Componentes Utilizados

As informações neste documento são baseadas nestas versões de software e hardware:

- C9300
- C9400
- C9500
- C9600

As informações neste documento foram criadas a partir de dispositivos em um ambiente de laboratório específico. Todos os dispositivos utilizados neste documento foram iniciados com uma configuração (padrão) inicial. Se a rede estiver ativa, certifique-se de que você entenda o impacto potencial de qualquer comando.

Observação: consulte o guia de configuração apropriado para obter os comandos que são usados para habilitar esses recursos em outras plataformas Cisco.

Informações de Apoio

Para configurar o EVPN TRM, consulte: [Guia de Configuração do BGP EVPN VXLAN, Cisco IOS XE Amsterdam 17.3.x](#)

O Tenant Routed Multicast (TRM) é uma solução baseada em BGP-EVPN que permite o roteamento multicast entre fontes e receptores conectados em VTEPS na estrutura VxLAN [RFC7432]. O TRM depende das rotas presentes no EVPN unicast para descobrir o RP multicast de origem e multicast. Assim como com NG-MVPN, as informações de origem e receptor do Multicast são propagadas pelo protocolo BGP entre os VTEPs configurados com a família de endereços BGP MVPN. Nenhum pacote PIM/IGMP é enviado para a estrutura VxLAN a partir de um TRM VTEP.

O principal problema que o TRM resolve é a capacidade dos remetentes e receptores multicast que estão localizados em diferentes VLANs, mas no mesmo VRF, de se comunicarem uns com os outros. Sem o TRM, o tráfego multicast é enviado como parte da mesma infraestrutura BUM (Broadcast, Unicast e Multicast) na subjacência, que pode ser uma árvore multicast ou uma replicação de entrada. Essa infraestrutura é criada por VLAN e, como resultado, enquanto as origens de multicast e os receptores na mesma VLAN podem se comunicar, aqueles que estão em VLANs diferentes não podem. Com o TRM, o Multicast é removido do BUM e agrupado em conjunto sob o VRF pai. Devido a isso, a comunicação multicast é totalmente habilitada, independentemente das VLANs nas quais a origem ou o receptor residem.

O TRM fornece encaminhamento multicast com reconhecimento de multilocatário entre remetentes e receptores na mesma sub-rede ou em sub-redes diferentes locais ou através de VTEPs. Consulte o guia [Guia de Configuração de BGP EVPN VXLAN, Cisco IOS XE Amsterdam 17.3.x](#) para obter mais detalhes

Como se orientar neste guia:

- O guia é dividido em 4 cenários com base na localização do RP.
- Um cenário pode se referir a exemplos CLI não diretamente na seção em que você está. Por exemplo, o Cenário 2 do SSM consulta o Cenário 1 para entender como ler determinadas CLIs.
- **Somente o cenário 1 cobre IPv4 e IPv6,** pois os conceitos são fundamentalmente os mesmos para ambas as famílias de endereços.
- **Os requisitos listados nesses Cenários supõem que a Origem e o Destinatário estejam diretamente conectados aos VTEPs** (Consulte a Seção de Informações Relacionadas 'Origens e Destinatários Fora da Estrutura' para obter mais informações sobre isso).

Cenários	IPv4/v6	O que é abordado em cada
Detalhes comuns para todos os outros cenários	IPv4	Sempre obrigatório: a base de MVPN e o NVE estão íntegros, a verificação de RPF para qualquer fonte de TRM é o L3VNI.
AnyCast RP (cada VTEP é um RP com um RP IP comum)	IPv4/v6	Comandos BGP, PIM, IGMP, MFIB e FED em detalhes completos para IPv4/v6 e exemplos de captura
Sem RP (Sobreposição de SSM)	IPv4	Informações específicas do SSM. (Consulte o cenário 1 para obter informações comuns)
RP dentro da estrutura (um RP comum para a estrutura)	IPv4	Comandos BGP, PIM, IGMP, MFIB e FED detalhadamente para IPv4
RP fora da estrutura (o RP não está na estrutura)	IPv4	Informações específicas de IP para borda de malha. (Consulte o cenário 3 para obter informações comuns)
RP dentro da estrutura (um RP comum para a estrutura) com L2VNI simétrico	IPv4	Avisos para uso de um único RP na estrutura quando o VNI está nos VTEPs do remetente e do receptor. (Consulte o Cenário 3 para obter informações comuns)

Neste documento de Troubleshooting, comentários foram adicionados ao final de determinadas linhas das saídas dos comandos show. Isso foi feito para destacar ou explicar um aspecto específico dessa linha de saída. Se um comentário começa em uma nova linha, ele se refere à linha de saída que precede o comentário. Esta notação foi usada em todo o documento para destacar os comentários dentro das saídas dos comandos show:

<#root>

<-â€" 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.

Terminologia

EVPN	Rede Privada Virtual Ethernet	A extensão que permite que o BGP transporte informações de MAC de Camada 2 e IP de Camada 3 é o EVPN e usa o Protocolo de Gateway de Borda Multiprotocolo (MP-BGP - Multi-Protocol Border Gateway Protocol) como o protocolo para distribuir informações de alcance que pertencem à rede de sobreposição de VXLAN.
VXLAN	LAN virtual extensível (rede local)	A VXLAN foi projetada para superar as limitações inerentes de VLANs e STP. É um padrão IETF proposto [RFC 7348] para fornecer os mesmos serviços de rede Ethernet de Camada 2 que as VLANs, mas com maior flexibilidade. Funcionalmente, é um protocolo de encapsulamento MAC-em-UDP executado como uma sobreposição virtual em uma rede de camada 3 subjacente.
VTEP	Ponto de Extremidade de Túnel Virtual	Esse é o dispositivo que faz o encapsulamento e o desencapsulamento
NVE	Interface virtual de rede	Interface lógica onde ocorre o encapsulamento e o desencapsulamento
VNI	identificador de rede VXLAN	Identifica exclusivamente cada sub-rede ou segmento da Camada 2. Existem dois tipos de VNI: Simétrico (L2VNI): VTEPs têm o mesmo VNI. Assimétrico (L3VNI): os VTEPs não têm o mesmo VNI e são roteados por meio de um único VNI comum.
MDT	Árvore de distribuição multicast	A árvore multicast criada entre VTEPs para encapsulamento e tunelamento de Tráfego Multicast de Locatário.

BUM	Broadcast, unicast desconhecido, multicast	O tráfego BUM é enviado através do grupo Mcast vinculado ao VNI na configuração NVE.
RP	Ponto de reunião	Uma função que um dispositivo executa quando está no Modo PIM Esparso. O ponto de reunião comum para fontes Multicast e Receptores.
AnyCast (RP)	Ponto de encontro do AnyCast	Dois ou mais RPs são configurados com o mesmo endereço IP em interfaces de loopback. O FHR registra o RP mais próximo com base no roteamento unicast.
RPT (Árvore)	Árvore de Caminho Raiz	Também chamada de árvore Shared ou *,G. Esse caminho é em direção ao RP
SPT (Árvore)	Árvore de caminho mais curto	O caminho mais curto para a Origem, conforme determinado pela Tabela de Roteamento Unicast
FHR	Roteador First Hop	O dispositivo que está diretamente conectado (ARP adjacente) à origem. O FHR registra informações de origem com o RP.
LHR	Roteador do último salto	O dispositivo onde o receptor está conectado
RPF	Encaminhamento de caminho reverso	O caminho Unicast de volta para a Origem. Os pacotes multicast de entrada não são aceitos/encaminhados, a menos que sejam recebidos no mesmo caminho que a tabela de roteamento unicast. (casos de uso de 'multicast multicast ip' excluídos).
MRIB	Base de Informações de Roteamento Multicast	A tabela de roteamento multicast do software, também chamada de tabela mroute
MFIB	Base de Informações de Encaminhamento Multicast	O equivalente multicast de CEF. Preenchido por atualizações do MRIB e usado para encaminhar pelo plano de dados.
FED	Driver do mecanismo de encaminhamento	O componente que programa o hardware do dispositivo.
IIF	Interface de entrada	Interface PIM habilitada, que também é o caminho de upstream de RPF unicast de volta à Origem. (visto em show ip mroute)

OIF	Interface de saída	Interface PIM habilitada que é downstream em direção ao receptor. (visto em show ip mroute)

Verificar

Verificação comum a todos os cenários

Esta primeira seção abrange os requisitos básicos **necessários para qualquer um dos cenários**.

- Certifique-se de que os correspondentes NVE necessários estejam ativos
- Certifique-se de que a interface RPF em direção à Origem no VRF de Locatário seja a SVI L3VNI. Se a interface RPF não for a SVI L3VNI, o BGP não enviará uma rota de junção tipo 7. Em qualquer cenário, a interface RPF deve apontar para essa interface.
- Verifique se o caminho de subjacência (túnel MDT) entre os correspondentes está completo.
- Certifique-se de que o BGP seja usado para o plano de controle de multicast (use MVPN versus PIM)

Observação: esta seção se aplica à verificação multicast de locatário IPv4 e IPv6.

Verificar peering NVE

Verifique se os peers NVE estão ativos entre VTEPs para qualquer um dos cenários neste guia

- Os peers NVE são formados por endereços aprendidos do BGP.

<#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

Verifique a interface RPF no VRF do usuário

Se essa interface for qualquer interface diferente da SVI do L3VNI, o BGP não originará uma junção MVPN Tipo 7.

- Se você não vir essa interface, confirme se não há nenhum problema com a configuração que tornaria a rota de volta para a origem uma interface que não seja a 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
```

Verifique se o plano de controle multicast usa BGP

- **mdt overlay use-bgp:** informa os dispositivos para usar o BGP MVPN tipo 5/6/7 como o protocolo de sinal (versus mensagens PIM)
- **spt-only:** palavra-chave adicional informa ao dispositivo para usar somente árvores SPT no cenário RP do AnyCast. Como cada VTEP é um RP, **não é usada a rota MVPN Tipo-6.**

```
<#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
```

Verificar grupo MDT

O grupo MDT é comum a todos os cenários, pois este é o grupo de túneis externos no qual o grupo TRM está encapsulado.

Verifique se o grupo MDT está programado corretamente no lado da origem

- A interface de entrada do grupo MDT é o Loopback do lado da origem
- A interface de saída do grupo MDT é a Interface Subjacente

Verifique a Leaf-01: a mroute MDT está correta em MRIB/MFIB

```
<#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

Verifique Leaf-01: entradas FED para o grupo MDT

<#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 FEI**

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_MULTICAS
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

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

Verifique se o grupo MDT está programado corretamente no lado do receptor

- A interface de entrada do grupo MDT é a interface RPF de volta ao Loopback do lado da origem
- A interface de saída do grupo MDT é a interface Encap/Decap Tunnel

Verifique a Leaf-02: a mroute MDT está correta em MRIB/MFIB

<#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

Verifique Leaf-02: entradas FED para o grupo MDT

<#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 FP


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_MULT
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x38 mtu_index/l3u_ri_index0:0x0 index1:0x

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...>
```

Cenário 1. AnyCast RP (árvores somente SPT) IPv4 e IPv6

Neste modo, existe um RP localizado em cada VTEP. Esses VTEPs não sincronizam fontes aprendidas via MSDP e não há nenhuma árvore compartilhada. Em vez disso, o modo MDT usa informações BGP para criar apenas árvores multicast SPT. Esse modo é alternadamente chamado de modo somente SPT ou modo Anycast-RP distribuído. Neste modo, cada VTEP é o RP PIM. Assim, a árvore (*,G) em cada site é truncada no próprio VTEP local. Não há necessidade de enviar junções (*,G) ou MVPN RT-6 através da estrutura.

Diagrama de Rede

Spine-01
172.16.255.1

Spine-02
172.16.255.2

Gi1/0/1
.13.1

Gi1/0/4
.16.1

Gi1/0/2
.14.1

Gi1/0/1
.23.2

Gi1/0/2
.24.2

Gi1/0/4
.26.6

172.16.x.x

Lo255
RP:10.2.255.255
RPv6:FC00:2:255::255

.13.3
Gi1/0/1

.23.3
Gi1/0/2

.14.4
Gi1/0/1

.24.4
Gi1/0/2

Lo255
RP:10.2.255.255
RPv6:FC00:2:255::255

Leaf-01
172.16.255.3

Leaf-02
172.16.255.4

Lo
172.

Gi1/0/10

Gi1/0/10

Gi1/0/1

Gi2/0/1

SRC
Vlan101:
10.1.101.11
FC00:1:101::11

RCV
Vlan102:
10.1.102.12
FC00:1:101::12
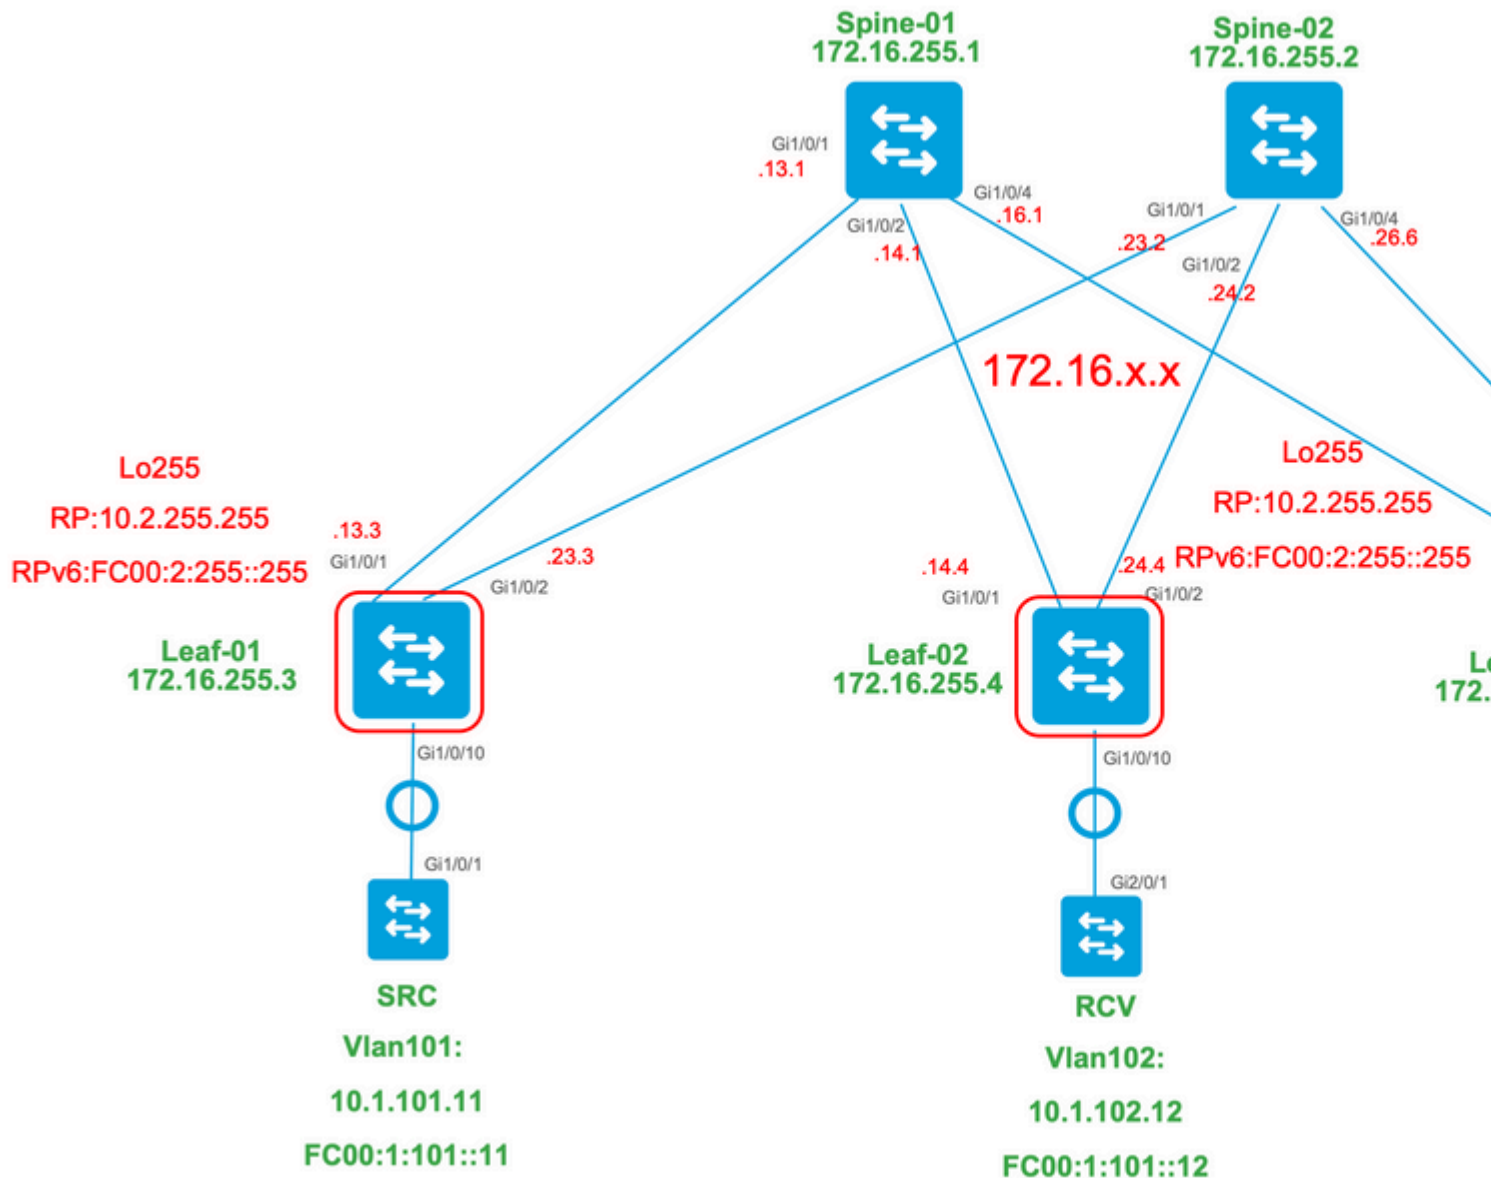
Para este modo, considere 3 tipos de rota BGP:

1. EVPN Route-type 2. Isso permite que os outros PEs que precisam construir uma rota C-Multicast (MVPN tipo 6/7) de volta ao PE de origem, anexem o RT de importação C-multicast apropriado para que o PE de origem possa importar a rota C-Multicast (RFC 6514 11.1.3) [RFC6514]. O uso desse VRI depende do comando **"mdt overlay use-bgp"** do comando VRF.
2. MVPN Route-type 5. É o mesmo que no MVPN e é o anúncio de uma fonte/grupo multicast disponível
3. MVPN Route-type 7. As informações da camada IGMP ou MLD e do EVPN Tipo 2 são usadas para criar essa junção do tipo BGP. O Type-7 orienta a criação do MRIB OIF no lado da origem.

Requisitos de EVPN tipo 2:

1. A origem de multicast diretamente conectada fica on-line.
2. O FHR (VTEP de origem) verifica o ARP (ou ND) e a adjacência de CEF (confirma que a origem está diretamente conectada).
3. O FHR origina a atualização do BGP EVPN Tipo 2

Requisitos de MVPN Tipo-5:

1. O requisito da conexão direta de origem foi resolvido
2. O RP é local, portanto, o FHR registra-se nele mesmo
3. O FHR origina a atualização de BGP Tipo 5 de MVPN

Requisitos de MVPN Tipo-7:

1. A entrada EVPN Tipo-2 está presente (necessária para construir a rota C-Multicast tipo-7 com VRI correto e enviada do VTEP de origem)
2. A entrada MVPN Tipo-5 está presente (necessária para resolver qual par de Origem/Grupo está disponível para união SPT)
3. O relatório de associação IGMP ou MLD foi recebido e processado pelo VTEP LHR
4. A interface RPF LHR VTEP é a interface L3VNI da estrutura

Dica: na saída, o PIM de VTEP de LHR verifica o caminho em direção à origem. O PIM deve encontrar uma rota no RIB que seja o L3VNI como a interface RPF. Se o L3VNI não estiver configurado corretamente, estiver inoperante e assim por diante. o VTEP não tenta criar a junção BGP tipo 7.

Verificar rotas BGP EVPN e MVPN

Verificar Leaf-01: o EVPN Type-2 é criado

```
<#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
```

Verifique se Leaf-01: ARP/IPv6 ND e as depurações de EVPN mostram que o ARP/ND foi aprendido e, em seguida, o Route-type 2 foi criado e enviado

<#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][10.1.1
*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, ex

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][FC00:1
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 Local
```

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

Verifique se Leaf-02: Source side **Route-type 2** é aprendido no BGP no lado do receptor

<#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

Verificar Leaf-02: Source Route-type 5 é aprendido no BGP no Receiver VTEP Leaf-02

<#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
```

Verifique se Leaf-02: precisou de informações de BGP de Leaf-01 para criar o Type-7. O requisito final é que o IGMP ou MLD tenha processado um relatório de associação que informe ao VTEP que há um Receptor interessado.

<#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
```

Verifique se Leaf-02: MVPN Debugs show Route-type 7 é criado quando o relatório de associação IGMP/MLD chega e o EVPN Type-2 e Type-5 necessários já estão instalados.

<#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, 1:1:1

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::11/128

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,
```

Verificar Leaf-01: O MVPN Type-7 recebido da Leaf-02

<#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 MVF

      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

Verificar Leaf-01: Depurações de MVPN mostram Route-type 7 recebido com o Route-Target de VRI de MVPN

<#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][10.1.101.11
*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)

Verificar Leaf-02: a tabela completa de BGP contém o EVPN Tipo 2 e MVPN Tipo 5 Leaf-01 e o Tipo 7 gerado pelo Receiver Leaf-02

<#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

Verifique o grupo TRM Leaf-01 (FHR)

Verifique se os grupos MDT e TRM estão formados corretamente no lado da Origem.

- A interface de entrada do grupo TRM é o SVI associado ao VRF do cliente
- A interface de saída do grupo TRM é a SVI L3VNI

Verificar Leaf-01: o grupo TRM 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 hard

 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

Verificar Leaf-01: o grupo TRM no FED

<#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

Verifique Leaf-01: a adjacência está correta

<#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)

Verifique o grupo TRM Leaf-02 (LHR)

Verifique se os grupos MDT e TRM estão formados corretamente no lado do receptor.

- A interface de entrada do grupo TRM é o SVI associado ao L3VNI
- A interface de saída do grupo TRM é a SVI do cliente em que a junção IGMP foi processada.

Verifique a Leaf-02: a rota TRM (rota multicast do usuário) em MRIB/MFIB

<#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
```

Verificar Leaf-02: o grupo TRM no FED

<#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_MULTICAS
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_MULTICAS
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
```

Verificar Leaf-02: A captura de pacote feita mostra o grupo de túneis MDT externo com tráfego de cliente interno

<#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)
```
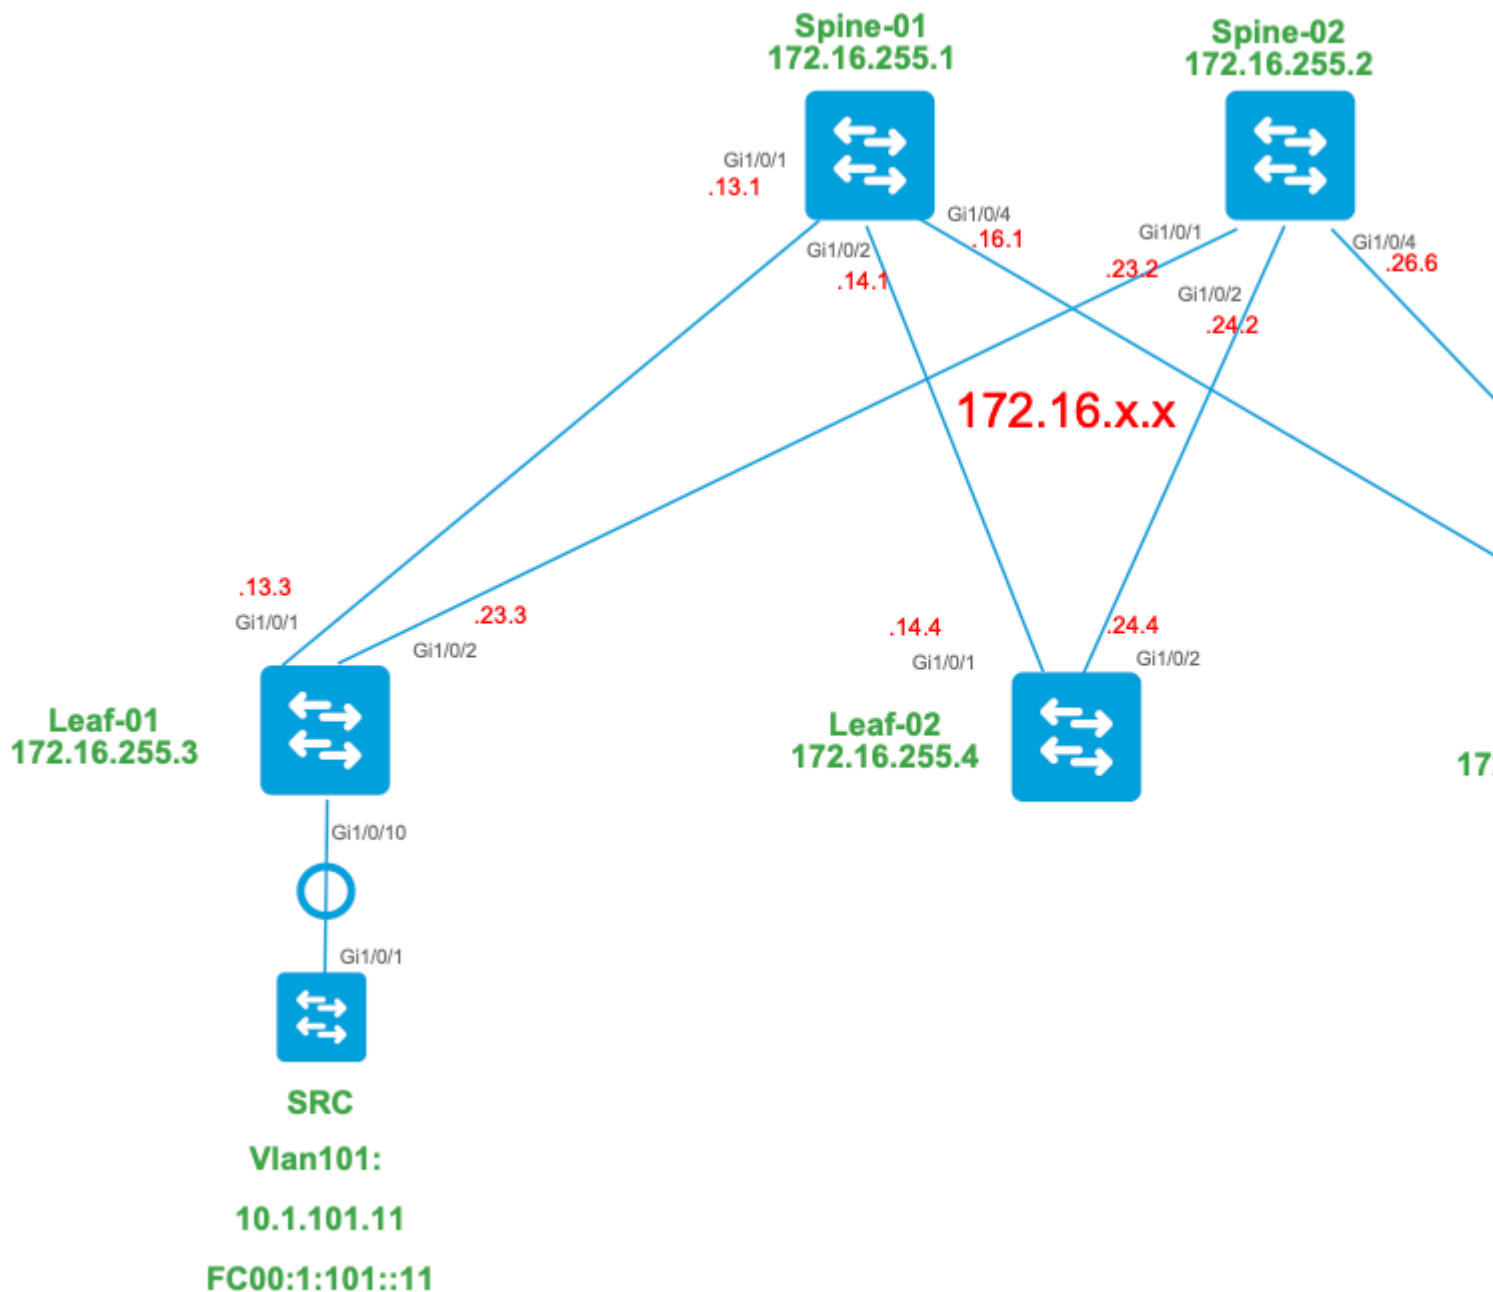
<...snip...>

Cenário 2: PIM SSM na malha

Neste modo, não há RP na Sobreposição e não são usados MVPN Tipo 5 ou Tipo 7 (a Sobreposição continua a operar como PIM ASM). **No SSM, o receptor envia e IGMPv3 S,G une** em direção ao VTEP LHR. Este VTEP executa pesquisa RPF para a Origem na RIB. Se L3VNI SVI for encontrado como a interface RPF, o LHR VTEP enviará o MVPN RT-7 para o FHR VTEP que recebe e instala essa rota. FHR VTEP informa PIM para adicionar L3VNI SVI como a interface de saída para S,G mroute.

Esta seção mostra as diferenças do Cenário 1. **As etapas e os métodos idênticos são observados apenas no Cenário 1.**

- Consulte as etapas de verificação e depuração para BGP e PIM do Cenário 1, pois as operações BGP e PIM são as mesmas

Diagrama de Rede

Spine-01
172.16.255.1

Spine-02
172.16.255.2

Gi1/0/1
.13.1

Gi1/0/4
.16.1

Gi1/0/2
.14.1

Gi1/0/1
.23.2

Gi1/0/4
.26.6

Gi1/0/2
.24.2

172.16.x.x

.13.3
Gi1/0/1

.23.3
Gi1/0/2

.14.4
Gi1/0/1

.24.4
Gi1/0/2

Leaf-01
172.16.255.3

Leaf-02
172.16.255.4

17

Gi1/0/10

Gi1/0/1

SRC

Vlan101:

10.1.101.11

FC00:1:101::11

Para este modo, considere estes tipos de rota BGP e suas origens

Criado por: Source VTEP

- EVPN Route-type 2. Usado para obter informações de Unicast e VRI para a Origem e adicionado à rota C-Multicast (MVPN tipo 7) quando o VTEP se une à árvore STP.

Criado por: Receiver VTEP

- MVPN Route-type 7. As informações da camada IGMP ou MLD e do EVPN Tipo 2 são usadas para criar essa junção do tipo BGP. O Type-7 orienta a criação do MRIB OIF no lado da origem.

Requisitos de EVPN tipo 2:

1. O FHR (VTEP de origem) verifica o ARP (ou ND) e a adjacência de CEF (confirma que a origem está diretamente conectada).
2. O FHR origina a atualização do BGP EVPN Tipo 2

Requisitos de MVPN Tipo-7:

1. A entrada EVPN Tipo-2 está presente (necessária para construir a rota C-Multicast tipo-7 com VRI correto e enviada do VTEP de origem)
2. VTEP do Receptor: o relatório de associação específico da origem IGMPv3 foi recebido e processado pelo VTEP LHR
3. A interface RPF LHR VTEP é a interface L3VNI da estrutura

Para esse modo, há uma configuração adicional necessária no LHR VTEP para ativar o intervalo SSM e processar relatórios de associação IGMPv3

Configure Leaf-03: defina o consultante IGMP para a Versão 3 na SVI do Locatário

```
<#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
```

Verificar Leaf-03: o consultante IGMP está definido como **versão 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
```

Habilitar Leaf-03: o intervalo SSM necessário para o VRF de Locatário

```
<#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
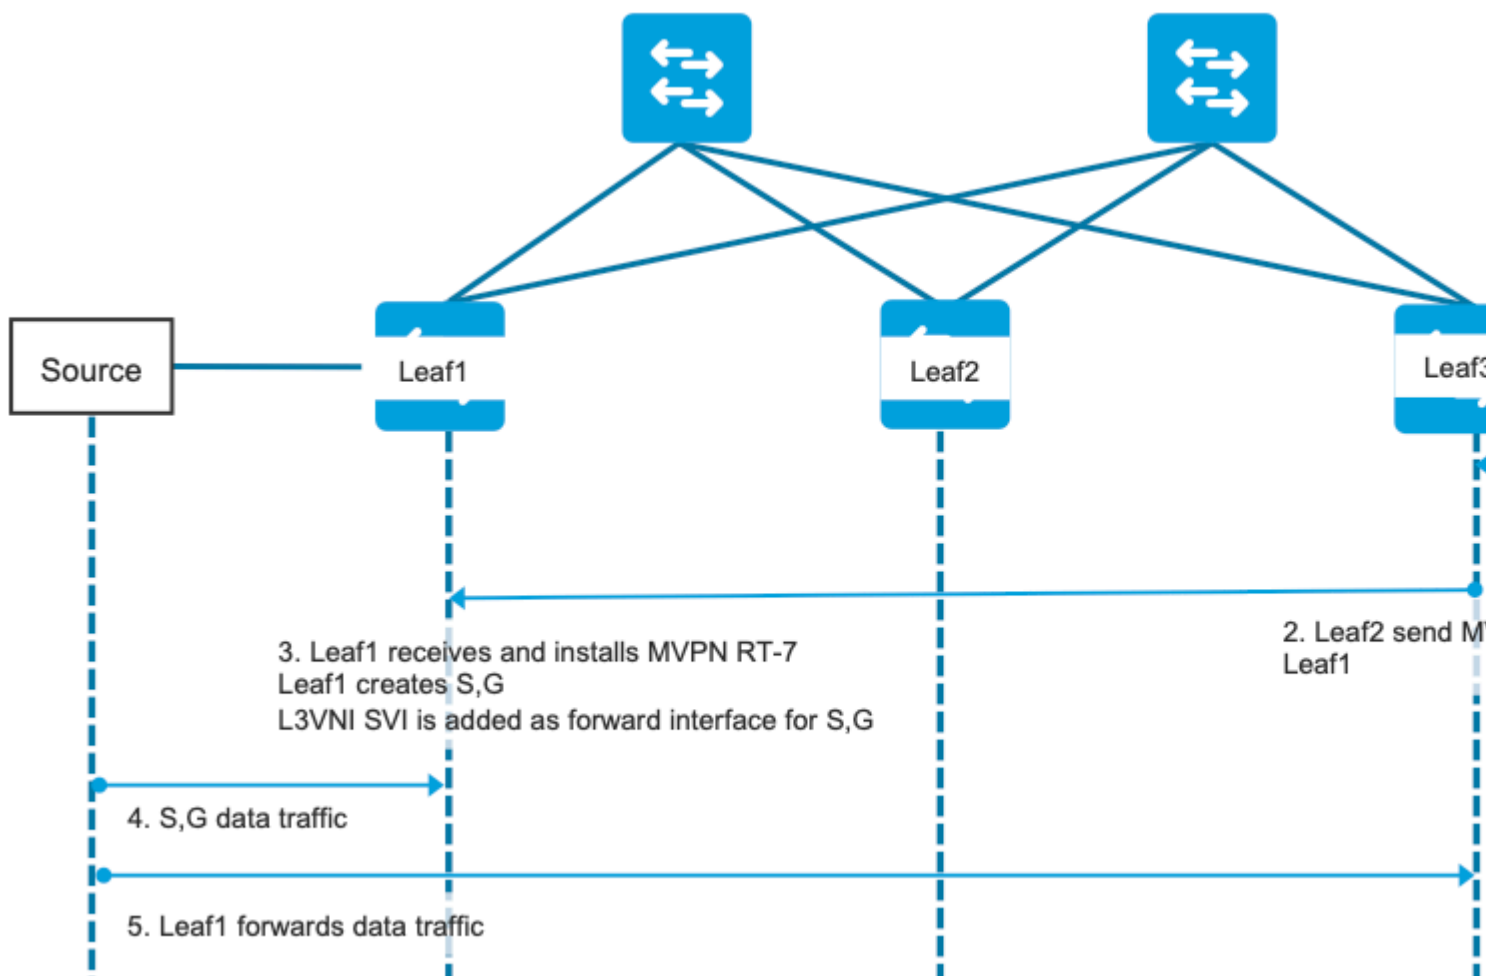

> **Dica**: os grupos SSM não criam um *,G mroute. Se você vir *,G para o grupo, verifique se sua configuração está correta para SSM.

Verificar a Sequência de Eventos Necessários para este Cenário

Etapa 0 EVPN (Leaf-03): Verifique se há um prefixo EVPN que o BGP pode encontrar o VRI para usar no MVPN tipo 7.

<#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
```

Etapa 1 (Folha-03): Relatório de associação IGMPv3 recebido e contém uma fonte

<#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

Etapa 2 (Leaf-03): o BGP é informado dessa junção, cria e envia a junção MVPN Tipo-7.

<#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
```

Etapa 3 (Leaf-01): a Source Leaf recebe e instala a rota de ingresso MVPN Type-7 e informa ao PIM para instalar o 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 Typ
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 Ac

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
```

Etapa 4 e 5 (Leaf-01 e Leaf-03): o multicast chega à folha do FHR e é enviado sobre a tela para a folha do LHR. Resumo dos comandos de validação fornecidos aqui. Você pode verificar a validação detalhada desses comandos no Cenário 1.

```
<#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
```

Cenário 3: RP único dentro da estrutura (modo disperso regular)

Esse modo é alternadamente chamado de modo RP não Anycast ou RP externo. Neste modo, há apenas um RP na sobreposição. Assim, a árvore (*,G) na sobreposição pode abranger vários sites. O BGP usa um MVPN RT-6 para anunciar a associação (*,G) através da estrutura. Se o RP e o FHR estiverem em locais diferentes, os registros PIM serão enviados pela tela. Este é o modo operacional padrão para o PIM SM na sobreposição.

Diagrama de Rede
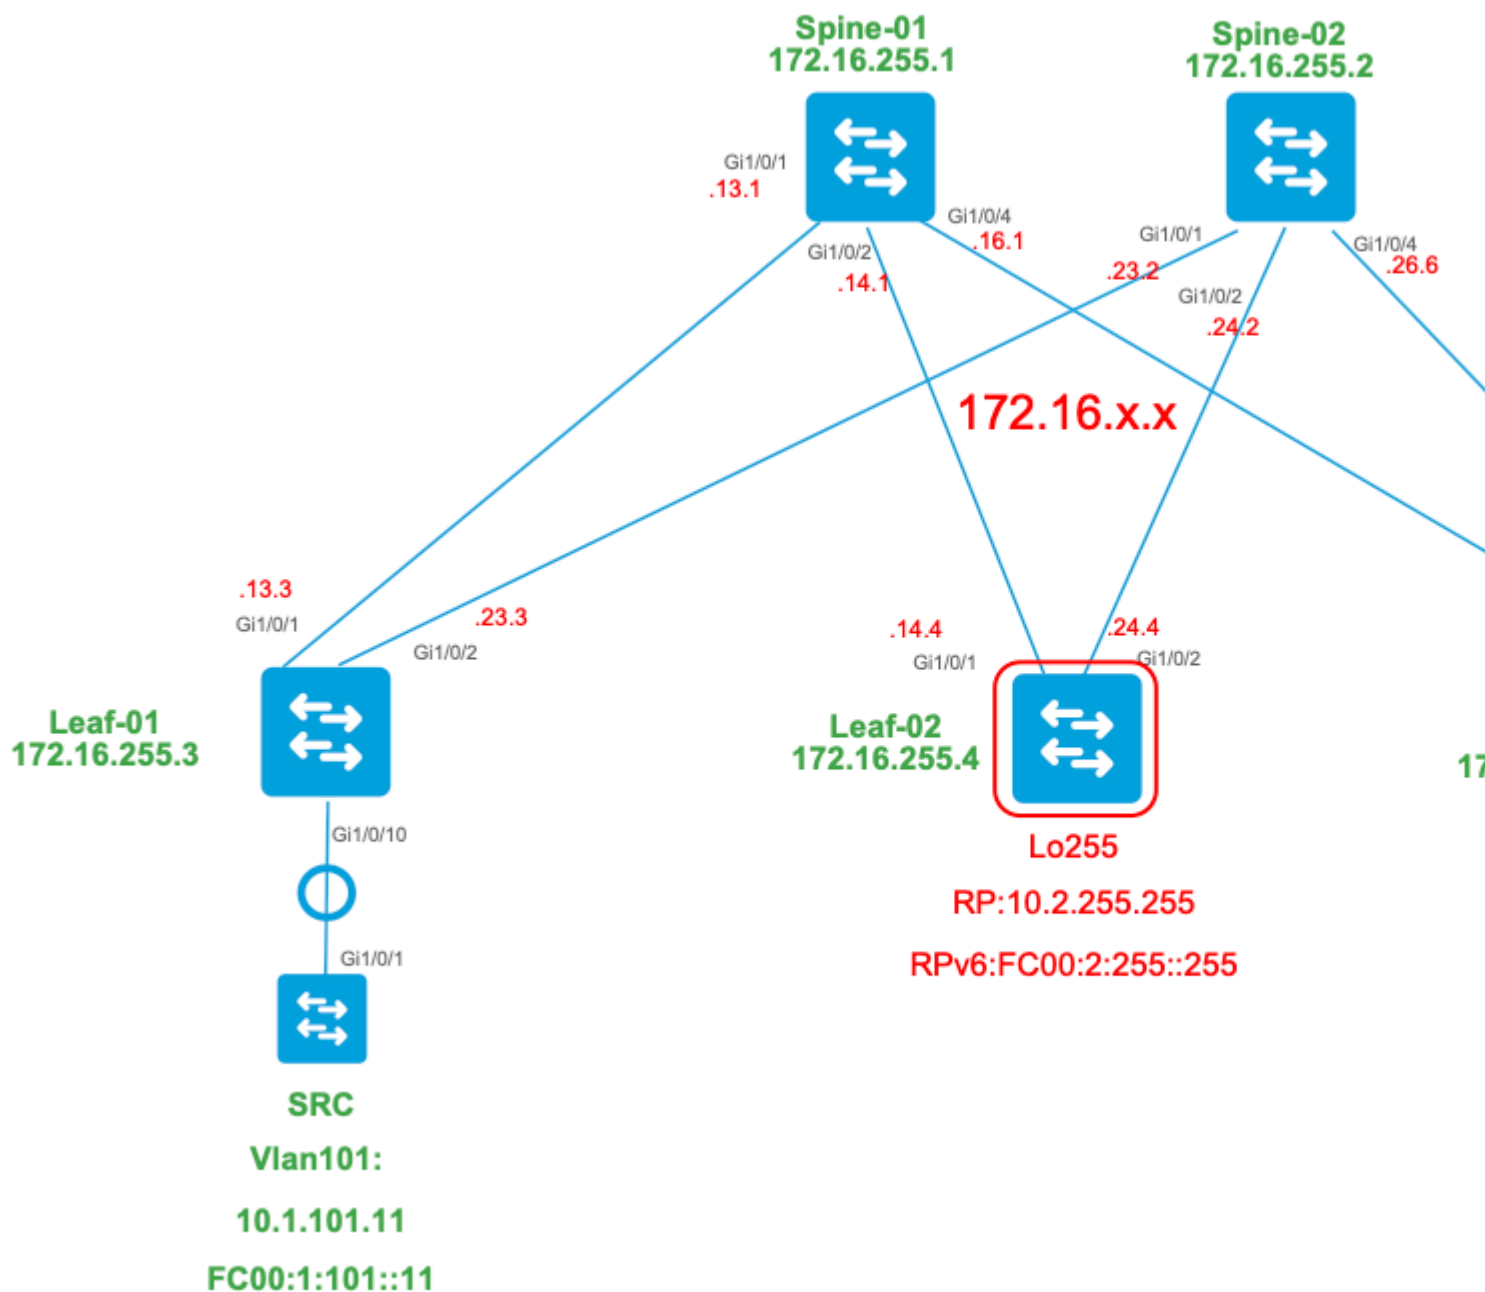
Para este modo, considere estes tipos de rota BGP e suas origens

Criado por: Source VTEP

- EVPN Route-type 2. Usado para obter informações de Unicast e VRI para a Origem e adicionado à rota C-Multicast (MVPN tipo 7) quando o VTEP se une à árvore STP.
- MVPN Route-type 5. Rota A-D de origem enviada a VTEPs para S,G

Criado por: RP VTEP

- EVPN Route-type 5. Usado para obter informações de Unicast e VRI para loopback RP. O loopback não cria o tipo de rota 2, portanto, o tipo 5 é usado.
- MVPN Route-type 7. Esta é a junção IGMP + detalhes de VRI RT retirados do EVPN Tipo 2 e enviados para o VTEP de origem, e conduz a criação do MRIB OIF.

Criado por: Receiver VTEP

- MVPN Route-type 6**.** Tipo de rota criado pelo VTEP do receptor para unir a Árvore compartilhada *,G (árvore RPT) em direção ao RP.
- MVPN Route-type 7. As informações da camada IGMP ou MLD e do EVPN Tipo 2 são usadas para criar essa junção do tipo BGP. O Type-7 orienta a criação do MRIB OIF no lado da origem.

Requisitos de EVPN tipo 2:

1. O FHR (VTEP de origem) verifica o ARP (ou ND) e a adjacência de CEF (confirma que a origem está diretamente conectada).
2. O FHR origina a atualização do BGP EVPN Tipo 2

Requisitos de EVPN tipo 5:

1. O loopback RP é configurado e anunciado no BGP

Requisitos de MVPN Tipo-5:

Neste modo, Leaf no site de origem anuncia mensagens A-D ativas de origem para um (S,G) somente se essas duas condições forem atendidas.

1. Recebe tráfego na interface RPF em direção à origem. (a origem envia mcast para o FHR)
2. A interface L3VNI SVI é adicionada como uma interface de encaminhamento para entrada (S,G), como resultado de uma junção S,G do RP como parte do processo de registro PIM. (A SVI L3VNI é instalada na lista OIF)

Requisitos de MVPN Tipo-6:

1. O RP anunciou sua rota EVPN Tipo-5 que contém seus detalhes de alcançabilidade de VRI e unicast.
2. Junção de IGMP recebida em LHR que dispara uma atualização de BGP para RP

Requisitos de MVPN Tipo-7:

1. A entrada EVPN Tipo-2 está presente (necessária para construir a rota C-Multicast tipo-7 com VRI correto e enviada do VTEP de origem)
2. A entrada MVPN Tipo-5 está presente (necessária para resolver qual par de Origem/Grupo está disponível para união STP)
3. VTEP do Receptor: o relatório de associação IGMP foi recebido e processado pelo VTEP LHR
4. RP VTEP: o RP recebeu pacotes de registro multicast, tem rotas EVPN e tem um receptor para S,G (aprendido através do tipo 6)
5. A interface RPF LHR VTEP é a interface L3VNI da estrutura

---

Dica: na saída, o PIM de VTEP de LHR verifica o caminho em direção à origem. O PIM deve encontrar uma rota no RIB que seja o L3VNI como a interface RPF. Se o L3VNI não estiver configurado corretamente, estiver inoperante e assim por diante. o VTEP não cria a junção BGP tipo 7.

Verificar a Sequência de Eventos Necessários para este Cenário

Valide as etapas necessárias para que o VTEP do receptor se una inicialmente à Árvore compartilhada e depois passe para a Árvore de caminho mais curto. Isso envolve verificações das tabelas BGP, IGMP e estados de criação de MRIB.
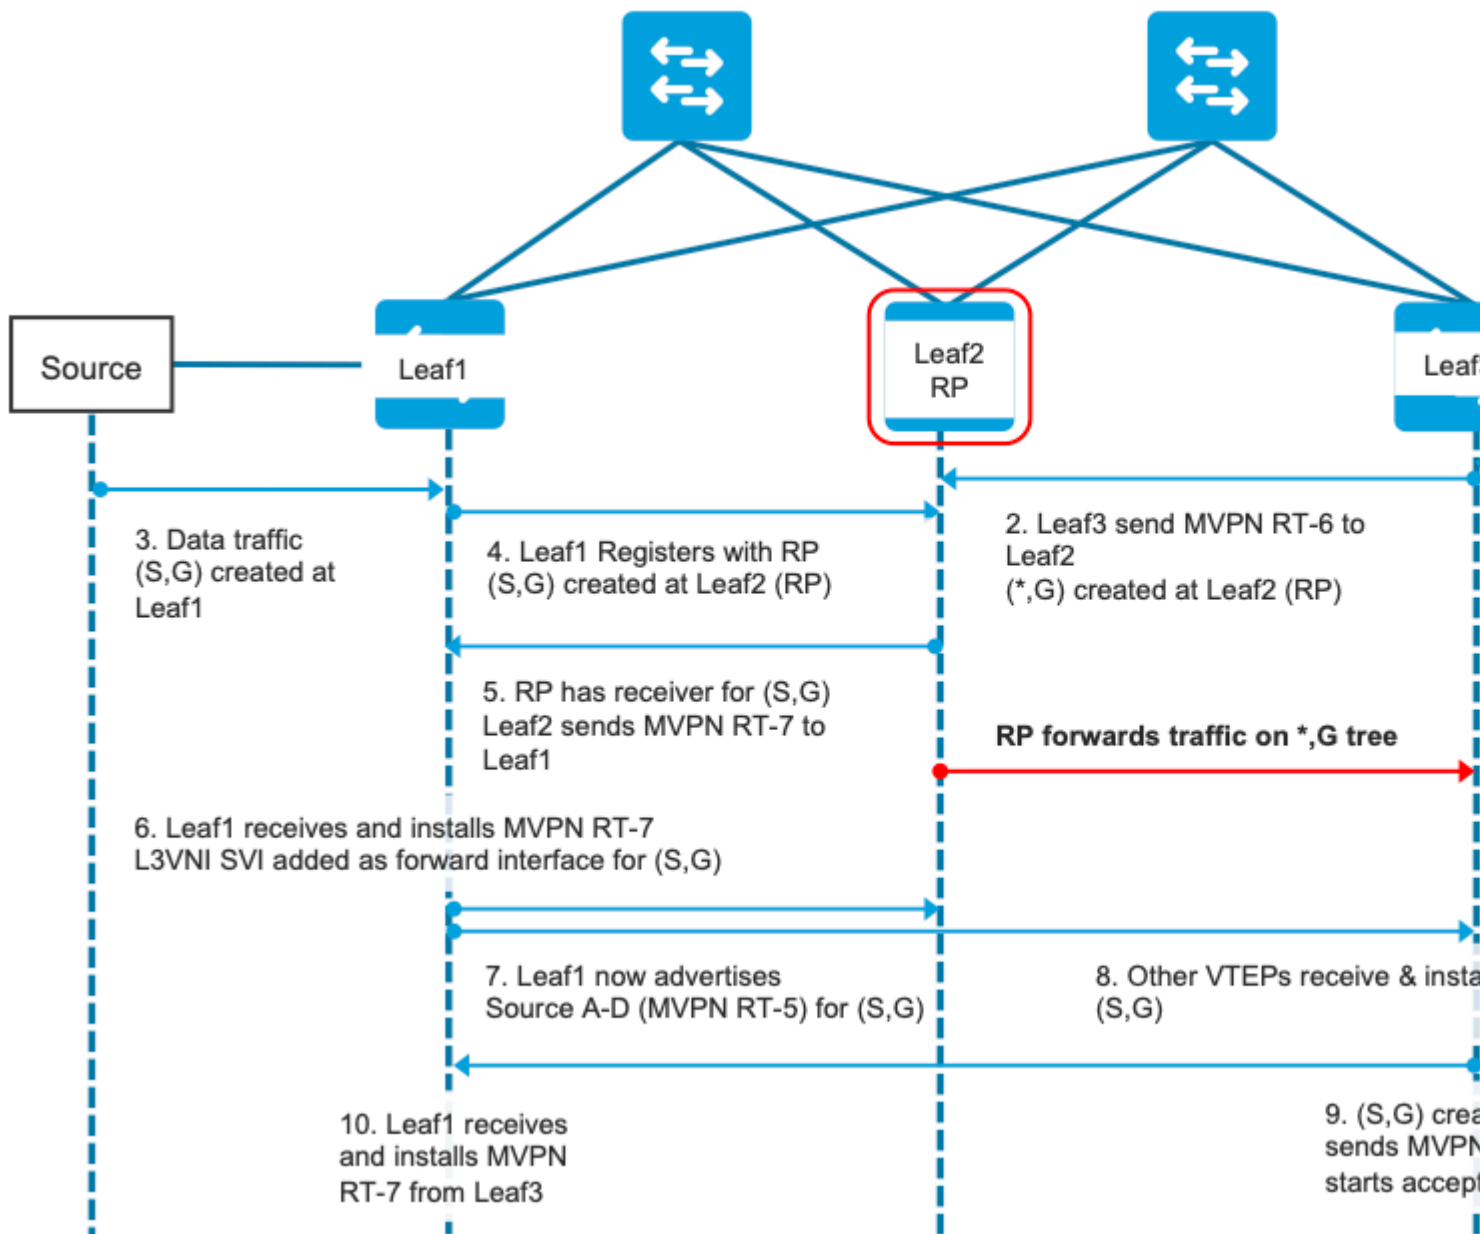
Etapa EVPN (Leaf-03): EVPN Type-5 do RP é aprendido no LHR. Isso é necessário para que o VTEP do receptor crie uma rota MVPN Tipo-6

```
<#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

Etapa 1 (Leaf-03): Relatório de associação IGMP recebido

<#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

Etapa 2 (Folha-03): MVPN Tipo-6 criado, enviado para RP e recebido por RP **(Folha-02)**

<#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
```

Etapa 1 e 2 - Depurações (Leaf-01): Relatório IGMP, pesquisa de origem EVPN e criação de MVPN Type-6

<#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 s
*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][6
*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, ex

<-- Advertise to RR

 (

172.16.255.1)

Etapas 3 e 4 (Leaf-01):Da perspectiva do FHR, valide os eventos S,G create & Register (S,G create & Register ocorrem quase ao mesmo tempo)

3. O tráfego de dados começa e S,G é criado em FHR VTEP. Os requisitos indicados na seção "Fontes de multicast não detectadas" aplicam-se aqui.

4. Leaf-01 executa o registro de Origem para RP através de seu túnel PIM

<#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 (MF

*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**

```
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

Etapa 4 (Folha-02): da perspectiva do RP, confirme se o registro de Origem alcança o RP e S,G é criado.

<#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
```

Etapa 5 (Leaf-02): o RP tem um receptor, por isso criou imediatamente a rota Type-7 MVPN Source Tree Join

```
<#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, 1:1:1

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, 1:1:1

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
```

Etapa 6 (Leaf-01): a Source Leaf-01 recebe e instala o MVPN Route-Type 7. (O VNI SVI de L3 é instalado como uma interface de encaminhamento para S,G)

<#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][10.1.101.11
*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
```

Etapa 7 (Leaf-01): Leaf-01 anuncia MVPN Source A-D Type-5 para S,G

```
<#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
```

Etapa 8 (Folha-03): o VTEP do receptor obtém o Tipo-5 e instala a rota de Origem A-D para 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 ?, localpr

*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
```

Etapa 9 (Leaf-03): S,G é criado, Leaf-03 envia MVPN Type-7 para entrar na árvore SPT e começa a aceitar tráfego

<#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
```

Etapa 10 (Leaf-01): Leaf-01 recebe e instala o MVPN Type-7 da Leaf-03

<#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][10.1.101.11
*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
```

Cenário 4: RP fora da estrutura (RP importado da folha de borda-02 do espaço IP)

Esse cenário é basicamente o mesmo que o cenário 2. Há um único RP usado pela estrutura como um todo.
A diferença é que o IP do RP deve ser importado de um espaço IP sem estrutura para a estrutura e anunciado
no BGP.

**Esta seção mostra as diferenças do Cenário 3. As etapas e os métodos idênticos são observados apenas
no Cenário 3**

- Consulte **Verificação da Sequência de Eventos Necessários para este Cenário** do Cenário 3, já que
  as operações BGP e PIM são as mesmas

Diagrama de Rede
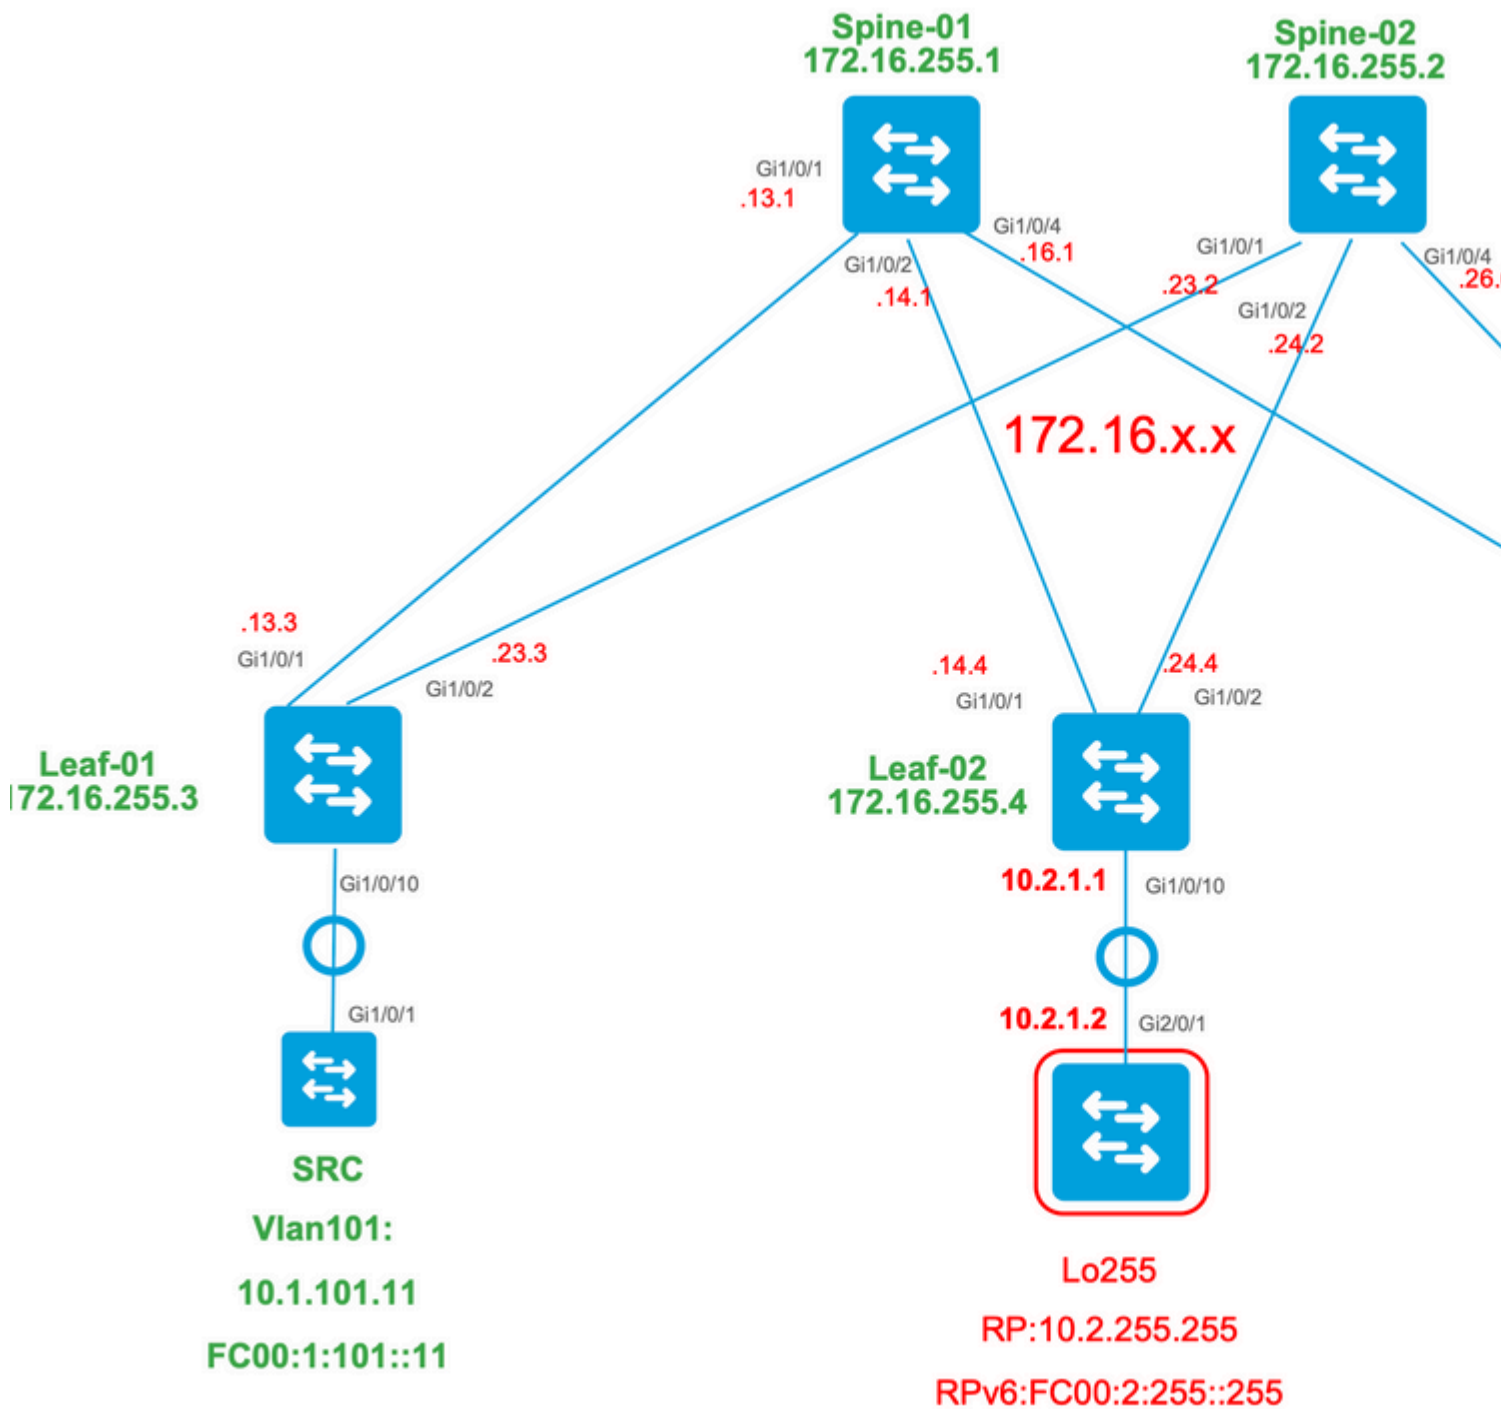
Verifique as importações de switch de borda de IP para malha

A principal diferença desse design em relação ao cenário 3 é a necessidade de primeiro importar o IP do RP do espaço IP para o EVPN.

A borda precisa conter determinados comandos para importar/exportar de e para espaços de estrutura e IP:

- Comandos **de costura route-target <value>** na seção de configuração do VRF
- **anuncie l2vpn evpn** na família de endereços BGP vrf

Verificar (Folha-02): Configuração

<#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 VF

exit-address-family
```

Verificar (Folha-02): Importação e Anúncio de Prefixo

<#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 10.2.1.2
*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
```

Verificar (Folha-02): Caminho da Borda para RP

<#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

Cenário 5: MDT de dados

Verificar grupo de dados MDT

O grupo de Dados MDT é semelhante a outro grupo Padrão MDT onde o grupo de túnel externo para o TRM ser encapsulado. No entanto, ao contrário do MDT Default, este grupo só terá VTEP's para se juntar a esta árvore se eles tiverem receptores interessados para o grupo TRM.

Configuração **necessária**

<#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
!

Verifique se o grupo MDT está programado corretamente no lado da origem

- A interface de entrada do grupo MDT é o Loopback do lado da origem
- A interface de saída do grupo MDT é a Interface Subjacente

Verifique a Leaf-01: a mroute MDT está correta em MRIB/MFIB

<#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

Verifique Leaf-01: entradas FED para o grupo MDT

<#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_MULTICAS
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_in
```

```
<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


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


Verifique se o grupo MDT está programado corretamente no lado do receptor

- A interface de entrada do grupo MDT é a interface RPF de volta ao Loopback do lado da origem
- A interface de saída do grupo MDT é a interface Encap/Decap Tunnel

Verifique a Leaf-02: a mroute MDT está correta em MRIB/MFIB


```
<#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

Verifique Leaf-02: entradas FED para o grupo MDT

<#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_MULT
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x1a mtu_index/l3u_ri_index0:0x0 index1:0x


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>
```

Depurar grupo de dados MDT

Use a depuração de MVPN para verificar o evento de cutover de MDT de dados

VTEP **do** lado **da origem**

```
<#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 Typ
*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, Wi
*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: UNKN
*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, Route
*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, Or
*Mar 27 12:13:00.431: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNKN
*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#
```

VTEP do **lado** do **receptor**

<#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 T
*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: UNKN
*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 send
*Mar 27 12:28:27.648: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: UNKN
*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 T
*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:0
*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, r
*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, r
*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:172.1
Leaf-03#
```

Troubleshooting

Fontes de multicast não detectadas

Geralmente, quando um host se torna ativo e envia tráfego, as entradas ARP são concluídas pelos procedimentos no
Mas, no caso de origens multicast, é possível que a origem comece a enviar tráfego e o plano L2 no FHR processe e
resolução do ARP para a origem.

A conclusão ARP desempenha um papel importante na funcionalidade do TRM por dois motivos.

1. **A verificação "diretamente conectada"** no primeiro salto do roteador chama uma API FIB que, por sua vez
   para a verificação bem-sucedida. Se o ARP em direção à origem de multicast não for concluído, a adjacência
   permanecerá incompleta e a verificação diretamente conectada retornará FALSE.
2. A detecção de origem aciona o **anúncio de EVPN RT-2 na estrutura de EVPN**. Essa rota EVPN instalada e
   receptor é usada como a rota RPF em direção à origem. Portanto, se a origem não for detectada, o RPF para a
   encontrado. Nesse caso, o RPF permanece NULO ou uma rota menos específica (se houver) é instalada no RI

Certifique-se de que o ARP seja resolvido e que a Origem esteja acessível dentro da estrutura EVPN.

Outras depurações úteis

Nesta seção há outras depurações que podem ser úteis no isolamento de problemas de TRM

- **debug mvpn** (todos os eventos de MVPN, consulte o Cenário 2, por exemplo)
- **debug ip|ipv6 pim <vrf>** (atividade do protocolo PIM)
- **debug ip mrib <vrf> trans** (MRIB, tradução PIM clássica)
- **debug ip mfib <vrf> pak|ps|fs** (Encaminhamento de pacotes| Comutação de processos| Comutação rápida)

Fontes e receptores fora da malha

Em alguns casos, a origem e/ou o receptor podem residir a um ou mais saltos L3 de distância dos VTEPs da e

Esse é um design válido, mas altera o tipo de rota EVPN que transporta o VRI e o processo responsável pela criação
receptor.

- Se a origem estiver fora da estrutura, o VTEP de entrada verá a origem por meio de um vizinho PIM, e não di
  enviará um EVPN tipo 5 para o VTEP do receptor. O VRI está contido neste Tipo-5.
- Se o receptor estiver fora da estrutura, a junção virá por meio de um PIM join IGMP. As informações no join
  MVPN Tipo 7.

Topologia eBGP Multiple AS (Spine to Spine)

Em alguns casos, a topologia pode exigir que o BGP envie informações de atualização para outro AS/Fabric.

É possível que até 30 segundos se passem para que as informações do plano de controle do BGP converjam e o mul

- **Isso ocorre devido ao intervalo de anúncio padrão do eBGP de 30 segundos**.
- Se houver um problema com tempos de convergência longos devido ao atraso nas atualizações de BGP
  eBGP pode ser reduzido para enviar atualizações com mais frequência.
- Consulte o guia de Configuração do BGP na seção Referência deste artigo para obter mais informações

eBGP inter-as **requer** um comando adicional

Use a palavra-chave inter-as para as rotas da família de endereços MVPN para cruzar os limites do sistema autônom

<#root>

Border-Leaf(config-vrf-af)#

mdt auto-discovery vxlan inter-as

Túnel de Registro com L2VNI Simétrico (FHR Preso no Estado de Registro PIM)

Nos casos em que o VNI existe no FHR e noutros VTEP, é possível que o FHR fique bloqueado no estado de r

Isso se deve ao fato de que o IP de origem do túnel de registro PIM é o gateway AnyCast. Quando o RP recebe um
qual é o VTEP correto para enviar a parada de registro, já que o IP é comum para vários dispositivos.

Problema de Rota de Túnel de Registro PIM

(Leaf-01) Este é o FHR real: Envia mensagens de registro ao RP

<#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): Este VTEP (e possivelmente outros) contém o mesmo SVI e endereço IP que o FHR

<#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): O FHR permanece preso no Registro (ele não recebe uma parada de registro do RP)

```
<#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) Este é o RP: Neste caso, ele também possui o mesmo IP do AnyCast que o FHR e, portanto, envia o regi

Se o RP não tiver o l2vni, mas 2 ou 3 outros vteps tiverem, register-stop poderá ser enviado para o VTEP errado, já
selecionar o correto.

```
<#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): Debug on RP mostra o problema em que o RP tem essa rota como Connected Local

```
<#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.
```

Solução de Problema de Rota de Túnel de Registro PIM

A **solução** é usar um IP de loopback exclusivo em todos os VTEPs e usar a configuração observada nesta seção.

```
<#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
```

Informações Relacionadas

Guia de Configuração do EVPN VxLAN TRM

Troubleshooting de EVPN VxLAN Unicast

Guia de configuração de MVPN 17.3.x (Switches Catalyst 9300)

Guia de configuração de MVPN 17.3.x (Switches Catalyst 9500)

[Manual de configuração de BGP](#)