

Probleemoplossing voor EVPN VLAN TRM op Catalyst 9000 Switches

Inhoud

[Inleiding](#)

[Voorwaarden](#)

[Vereisten](#)

[Gebruikte componenten](#)

[Achtergrondinformatie](#)

[Terminologie](#)

[Verifiëren](#)

[Verificatie Gemeenschappelijk voor alle scenario's](#)

[Controleer NVE-peeling](#)

[Controleer de RPF-interface in de huurder VRF](#)

[Controleer of het multicast-besturingsplane BGP gebruikt](#)

[Controleer MDT-groep](#)

[Scenario 1. AnyCast RP \(alleen SPT-bomen\) IPv4 en IPv6](#)

[Netwerkdigram](#)

[Controleer BGP EVPN- en MVPN-routers](#)

[Controleer de TRM-groep Leaf-01 \(FHR\).](#)

[Controleer TRM-groep Leaf-02 \(LHR\).](#)

[Scenario 2: PIM SSM in de fabric](#)

[Netwerkdigram](#)

[Controleer de volgorde van gebeurtenissen die voor dit scenario vereist zijn](#)

[Scenario 3: Single RP Inside the Fabric \(Regular Sparse-Mode\)](#)

[Netwerkdigram](#)

[Controleer de volgorde van gebeurtenissen die voor dit scenario vereist zijn](#)

[Scenario 4: RP Buiten de Stof \(RP geïmporteerd uit Border Leaf-02 vanuit IP-ruimte\)](#)

[Netwerkdigram](#)

[Controleer de invoer van border-switch van IP naar fabric](#)

[Scenario 5: Data MDT](#)

[Controleer MDT-gegevensgroep](#)

[Debug MDT-gegevensgroep](#)

[Problemen oplossen](#)

[Ongedetecteerde multicastbronnen](#)

[Andere nuttige debugs](#)

[Bronnen en ontvangers buiten de fabric](#)

[eBGP meervoudige AS \(spine-naar-spine\) topologie](#)

[Tunnel registreren met symmetrische L2VNI \(FHR vastgelegd in PIM-registratiestaat\)](#)

[Gerelateerde informatie](#)

Inleiding

Dit document beschrijft hoe u problemen met TRM (Tenant Routed Multicast) kunt oplossen via EVPN VxLAN.

Voorwaarden

- Aanbevolen wordt om bekend te zijn met de Unicast EVPN VxLAN-functie, BGP en MVPN (Multicast Virtual Private Network).
- Bovendien moet u begrijpen hoe multicast werkt en hoe multicast concepten werken

Vereisten

Deze gids veronderstelt BGP, zijn de peers NVE reeds correct. Als er problemen zijn met de basis-EVPN VxLAN (Unicast ping-fout, BGP, NVE-peers omlaag, enzovoort), raadpleeg dan zo nodig BGP, EVPN, route/switch-probleemoplossingsgidsen.

Beschikbaarheid van functies in elke softwarerelease

Release	Feature
17.1.1	TRMv4 met anycast RP
17.3.1	TRMv4 met externe RP of enkelvoudige RP
17.3.1	TRMv6 met anycast RP
17.3.1	TRMv6 met externe RP of enkelvoudige RP
17.3.1	TRMv4 met MVPN interworking (profile11) met Single RP aan de zijde van de stof
17.6.2 en 17.7.1	TRMv4 Data mdt met Anycast RP, Externe RP of Single RP

Gebruikte componenten

De informatie in dit document is gebaseerd op de volgende software- en hardware-versies:

- C9300
- C9400
- C9500
- C9600

De informatie in dit document is gebaseerd op de apparaten in een specifieke laboratoriumomgeving. Alle apparaten die in dit document worden beschreven, hadden een opgeschoonde (standaard)configuratie. Als uw netwerk live is, moet u zorgen dat u de potentiële impact van elke opdracht begrijpt.

Opmerking: raadpleeg de juiste configuratiehandleiding voor de opdrachten die worden gebruikt om deze functies op andere Cisco-platforms in te schakelen.

Achtergrondinformatie

Voor het configureren van EVPN TRM raadpleegt u: [BGP EVPN VXLAN Configuration Guide, Cisco IOS XE Amsterdam 17.3.x](#)

Tenant Routed Multicast (TRM) is een op BGP-EVPN gebaseerde oplossing die multicast routing tussen bronnen en ontvangers mogelijk maakt die zijn aangesloten op VTEPS in VxLAN-fabric [RFC7432]. TRM vertrouwt op routes die aanwezig zijn in de unicast EVPN om multicast bron en multicast RP te ontdekken. Net als bij NG-MVPN wordt informatie over multicastbronnen en -ontvangers door het BGP-protocol verspreid onder VTEP's die met de BGP MVPN-adresfamilie zijn geconfigureerd. Er worden geen PIM/IGMP-pakketten vanuit een TRM VTEP naar VxLAN-fabric verzonden.

Het belangrijkste probleem dat het TRM oplost is de capaciteit van multicast afzenders en ontvangers

die in verschillende VLANs maar in zelfde VRF worden gevestigd, met elkaar kunnen communiceren. Zonder TRM wordt multicast verkeer verzonden als deel van dezelfde BUM (Broadcast, Unicast en Multicast)-infrastructuur in de onderlaag, die een multicast-boom of een toegangsreplicatie kan zijn. Deze infrastructuur wordt per VLAN gebouwd en bijgevolg kunnen multicast-bronnen en ontvangers op hetzelfde VLAN communiceren, maar niet met die op verschillende VLANs. Met TRM, wordt Multicast verplaatst uit BUM en samengevoegd onder de oude VRF. Hierdoor is multicast communicatie volledig ingeschakeld, ongeacht de VLANs waarin de bron of de ontvanger zich bevindt.

TRM biedt multi-tenancy bewuste multicast doorsturen tussen afzenders en ontvangers binnen dezelfde of verschillende subnetten lokaal of over VTEP's. Zie gids [BGP EVPN VXLAN-configuratiehandleiding, Cisco IOS XE Amsterdam 17.3.x](#) voor meer informatie

Hoe te om jezelf te oriënteren in deze gids:

- De gids is onderverdeeld in 4 scenario's op basis van de RP locatie.
- Een scenario kan verwijzen naar CLI voorbeelden niet direct in de sectie waar u in bent. Bijvoorbeeld, SSM Scenario 2 verwijst u naar Scenario 1 om te begrijpen hoe te bepaalde CLIs te lezen.
- **Alleen scenario 1 heeft betrekking op IPv4 en IPv6**, aangezien de concepten fundamenteel hetzelfde zijn voor beide adresfamilies.
- **De vereisten die in deze scenario's worden vermeld gaan ervan uit dat de bron en ontvanger direct aan de VTEP's zijn gekoppeld** (zie de sectie 'Bronnen en ontvangers buiten de fabric' voor meer informatie over dit onderwerp).

Scenario's	IPv4/v6	Wat is in elk document opgenomen?
Gemeenschappelijke details voor alle andere scenario's	IPv4	Altijd vereist: MVPN onderlaag en NVE zijn gezond, RPF controle naar elke TRM bron is de L3VNI.
AnyCast RP (elke VTEP is een RP met een gemeenschappelijke RP IP)	IPv4/v6	BGP-, PIM-, IGMP-, MFIB- en FED-opdrachten volledig in detail voor zowel IPv4/v6 als voorbeelden van opnamen
Geen RP (SSM-overlay)	IPv4	GTM-specifieke informatie. (Raadpleeg scenario 1 voor algemene informatie)
RP In de stof (een gemeenschappelijke RP voor de stof)	IPv4	BGP-, PIM-, IGMP-, MFIB- en FED-opdrachten in alle details voor IPv4
RP Buiten de stof (de RP zit niet in de stof)	IPv4	Specifieke informatie over IP-naar-fabric-grens. (Zie scenario 3 voor algemene informatie)
RP Binnen de Stof (een gemeenschappelijke RP voor de stof) met symmetrische L2VNI	IPv4	Bezwaren voor het gebruik van een enkele RP in stof wanneer de VNI op zowel Afzender als Ontvanger VTEPs is. (Zie scenario 3 voor algemene informatie)

In dit document van het oplossen van problemen, zijn de commentaren toegevoegd aan het eind van bepaalde lijnen van de output van tonen bevelen. Dit is gedaan om een specifiek aspect van die productielijn te belichten of te verklaren. Als een commentaar in een nieuwe lijn begint, dan verwijst het naar de lijn van output die de commentaar voorafgaat. Deze notatie is in het gehele document gebruikt om de opmerkingen in de uitgangen van show opdrachten te markeren:

```
<#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.
```

Terminologie

EVPN	Ethernet Virtual Private Network	De uitbreiding die BGP in staat stelt Layer 2 MAC- en Layer 3 IP-informatie te transporteren, is EVPN en gebruikt Multi-Protocol Border Gateway Protocol (MP-BGP) als het protocol om bereikbaarheidsinformatie te distribueren die betrekking heeft op het VXLAN-overlay-netwerk.
VXLAN	Virtual Extensible LAN (Local Area Network)	VXLAN is ontworpen om de inherente beperkingen van VLANâ€™s en STP te overwinnen. Het is een voorgestelde IETF-standaard [RFC 7348] om dezelfde Ethernet Layer 2-netwerkservices te bieden als VLANâ€™s, maar met grotere flexibiliteit. Functioneel is het een MAC-in-UDP-inkapselingsprotocol dat als virtuele overlay op een Layer 3 underlay-netwerk wordt uitgevoerd.
VTEP	Virtueel tunnelendpoint	Dit is het apparaat dat de inkapseling en de-inkapseling doet
NVE	Virtuele interface voor netwerk	Logische interface waar de inkapseling en de-inkapseling voorkomen
VNI	VXLAN-netwerkidentificatie	Identificeert elke Layer 2-subnetprocessor of -segment op unieke wijze. Er zijn twee typen VNI: Symmetrisch (L2VNI): VTEP's hebben dezelfde VNI Asymmetrisch (L3VNI): VTEP's hebben niet hetzelfde VNI en worden gerouteerd via één gemeenschappelijk VNI.
MDT	Multicastdistributiestructuur	De multicast boom die tussen VTEPs voor inkapseling en het een tunnel graven van Huurder Multicast Verkeer wordt

		gebouwd.
STOMME	Uitzending, onbekende Unicast, Multicast	BUM-verkeer wordt verzonden via de Mcast-groep die is gekoppeld aan de VNI onder de NVE-configuratie.
RP	Rendezvous Point	Een rol die een apparaat uitvoert wanneer in PIM Sparse Mode. Het gemeenschappelijke vergaderpunt voor multicastbronnen en ontvangers.
AnyCast (RP)	AnyCast rendez-vous point	Twee of meer RP's zijn geconfigureerd met hetzelfde IP-adres op loopback interfaces. FHR registreert op de dichtstbijzijnde RP op basis van Unicast-routing.
RPT (structuur)	Root Path Tree	Wordt ook de gedeelde of *,G-boom genoemd. Dit pad is richting de RP
NBP (boom)	Structuur met kortste pad	Het kortste pad naar de bron zoals bepaald door de Unicast Routing Table
FHR	Eerste hoprouter	Het apparaat dat direct verbonden is (ARP naast) met de bron. De bron van de FHR-registers bij de RP.
LHR	Laatste hop router	Het apparaat waarop de ontvanger is aangesloten
RPF	Doorsturen van pad omkeren	Het Unicast-pad keert terug naar de bron. Inkomende multicast-pakketten worden niet aanvaard/doorgestuurd, tenzij ze hetzelfde pad ontvangen hebben als de unicast-routertabel. ('ip multicast multipath' niet van toepassing).
MIRIB	Multicast-routing informatiebasis	De software multicast routingstabel, ook wel route-tabel genoemd
MFIB	Multicast Forwarding Information Base	Het multicast-equivalent van CEF. Populated door updates van MRIB, en gebruikt om door te sturen door het vliegtuig van Gegevens.
FED	Forwarding Engine Driver	De component die de hardware van het apparaat programmeert.
IIF	Inkomende interface	PIM enabled interface die ook de unicast RPF upstream pad is terug naar de bron. (gezien in show ip mroute)

OIF	Uitgaande interface	PIM-enabled interface die stroomafwaarts naar de ontvanger is. (gezien in show ip mroute)
------------	---------------------	--

Verifiëren

Verificatie Gemeenschappelijk voor alle scenario's

In dit eerste deel worden de basisvereisten besproken die **voor een van de scenario's vereist zijn**.

- Zorg ervoor dat de vereiste NVE-peers zijn ingesteld
- Zorg ervoor dat de RPF-interface naar de bron in de huurder VRF de L3VNI SVI is. Als de RPF-interface niet de L3VNI SVI is, stuurt BGP geen type-7 route. In om het even welk scenario, moet de interface RPF aan deze Interface richten.
- Zorg ervoor dat het Underlay-pad (MDT-tunnel) tussen peers is voltooid.
- Zorg ervoor dat BGP wordt gebruikt voor multicast control-plane (gebruik MVPN versus PIM)

Opmerking: deze sectie is van toepassing op zowel IPv4 als IPv6 Tenant multicast verificatie.

Controleer NVE-peeling

Controleer of de NVE-peers tussen VTEP's zitten voor een van de scenario's in deze handleiding

- NVE-peers worden gevormd door adressen die van BGP zijn geleerd.

```
<#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 Leaf 02
```

```
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 Leaf 01
```

```
nve1      50901  L3CP 172.16.254.3   10b3.d56a.8fc8  50901    UP      A/M/6 17:56:19 <-- IPv6 peering with Leaf 01
```

Controleer de RPF-interface in de huurder VRF

Als deze interface een andere interface is dan de L3VNI SVI, komt BGP niet uit een MVPN Type-7-verbinding.

- Als u deze interface niet ziet, bevestig dan dat er geen probleem is met de configuratie dat de route terug naar de bron een interface zou maken die niet de L3VNI is.

```
<#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
```

Controleer of het multicast-besturingsplane BGP gebruikt

- **mdt overlay use-bgp:** informeert apparaten om BGP MVPN type 5/6/7 te gebruiken als het signaalprotocol (tegenover PIM-berichten)
- **spt-only:** extra trefwoord informeert apparaat om alleen SPT-bomen te gebruiken in AnyCast RP Scenario. Aangezien elke VTEP een RP is, **wordt er geen MVPN Type-6 route gebruikt.**

```
<#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
```

Controleer MDT-groep

De MDT-groep is gemeenschappelijk voor alle scenario's aangezien dit de buitentunnelgroep is waarin de TRM-groep is ingekapseld.

Controleer of de MDT-groep aan de bronkant correct is geprogrammeerd

- Inkomende interface van MDT groep is de bronkant Loopback
- Uitgaande interface van MDT-groep is de Underlay Interface

Controleer Leaf-01: de MDT-route is correct in 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

Controleer Leaf-01: FED-vermeldingen voor de MDT-groep

<#root>

Leaf-01#

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

MROUTE ENTRY

vrf 0

(

172.16.254.3, 239.1.1.1/32

)

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

HW Handle: 139738317079128 Flags:

RPF interface: Null0

(1):

<-- Leaf-01 the Source (Null0)

HW Handle:139738317079128 Flags:A

Number of OIF: 2

Flags: 0x4

Pkts : 71

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

OIF Details:

Null0 A

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

GigabitEthernet1/0/2

F

NS

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

Htm: 0x7f175cc0beb8 Si: 0x7f175cc0a6b8

Di: 0x7f175cc09df8

Rep_ri: 0x7f175cc0a1d8

<-- The DI (dest index) handle

DI details

Handle:0x7f175cc09df8 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICAST
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

=====

Controleer of de MDT-groep aan ontvangerzijde correct is geprogrammeerd

- Inkomende interface van MDT-groep is de RPF-interface terug naar de bronzijde Loopback
- Uitgaande interface van MDT-groep is Encap/Decap Tunnel-interface

Controleer Leaf-02: de MDT-route is correct in 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 Lol 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 hardware
```

```
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
```

Controleer Leaf-02: FED-vermeldingen voor de MDT-groep

<#root>

Leaf-02#

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

MROUTE ENTRY

vrf 0

(

172.16.254.3, 239.1.1.1/32

)

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

HW Handle: 140397391831832 Flags:

RPF interface: GigabitEthernet1/0/2

(57)):

<-- RPF interface to 172.16.254.3

HW Handle:140397391831832 Flags:A

Number of OIF: 2

Flags: 0x4

Pkts : 1585

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

OIF Details:

Tunnel0 F NS

<-- Send to decap tunnel to remove VxLAN header

(Adj: 0x73)

<-- Tunnel0 Adjacency

GigabitEthernet1/0/2 A

<-- Accept MDT packets from this interface

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

RI details

<-- Rewrite Index is used for VxLAN decapsulation

Handle:0x7fb0d0ed07a8 Res-Type:ASIC_RSC_RI_REP Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_MULTICAST
priv_ri/priv_si Handle:(nil) Hardware Indices/Handles: index0:0x38 mtu_index/l3u_ri_index0:0x0 index1:0x0

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

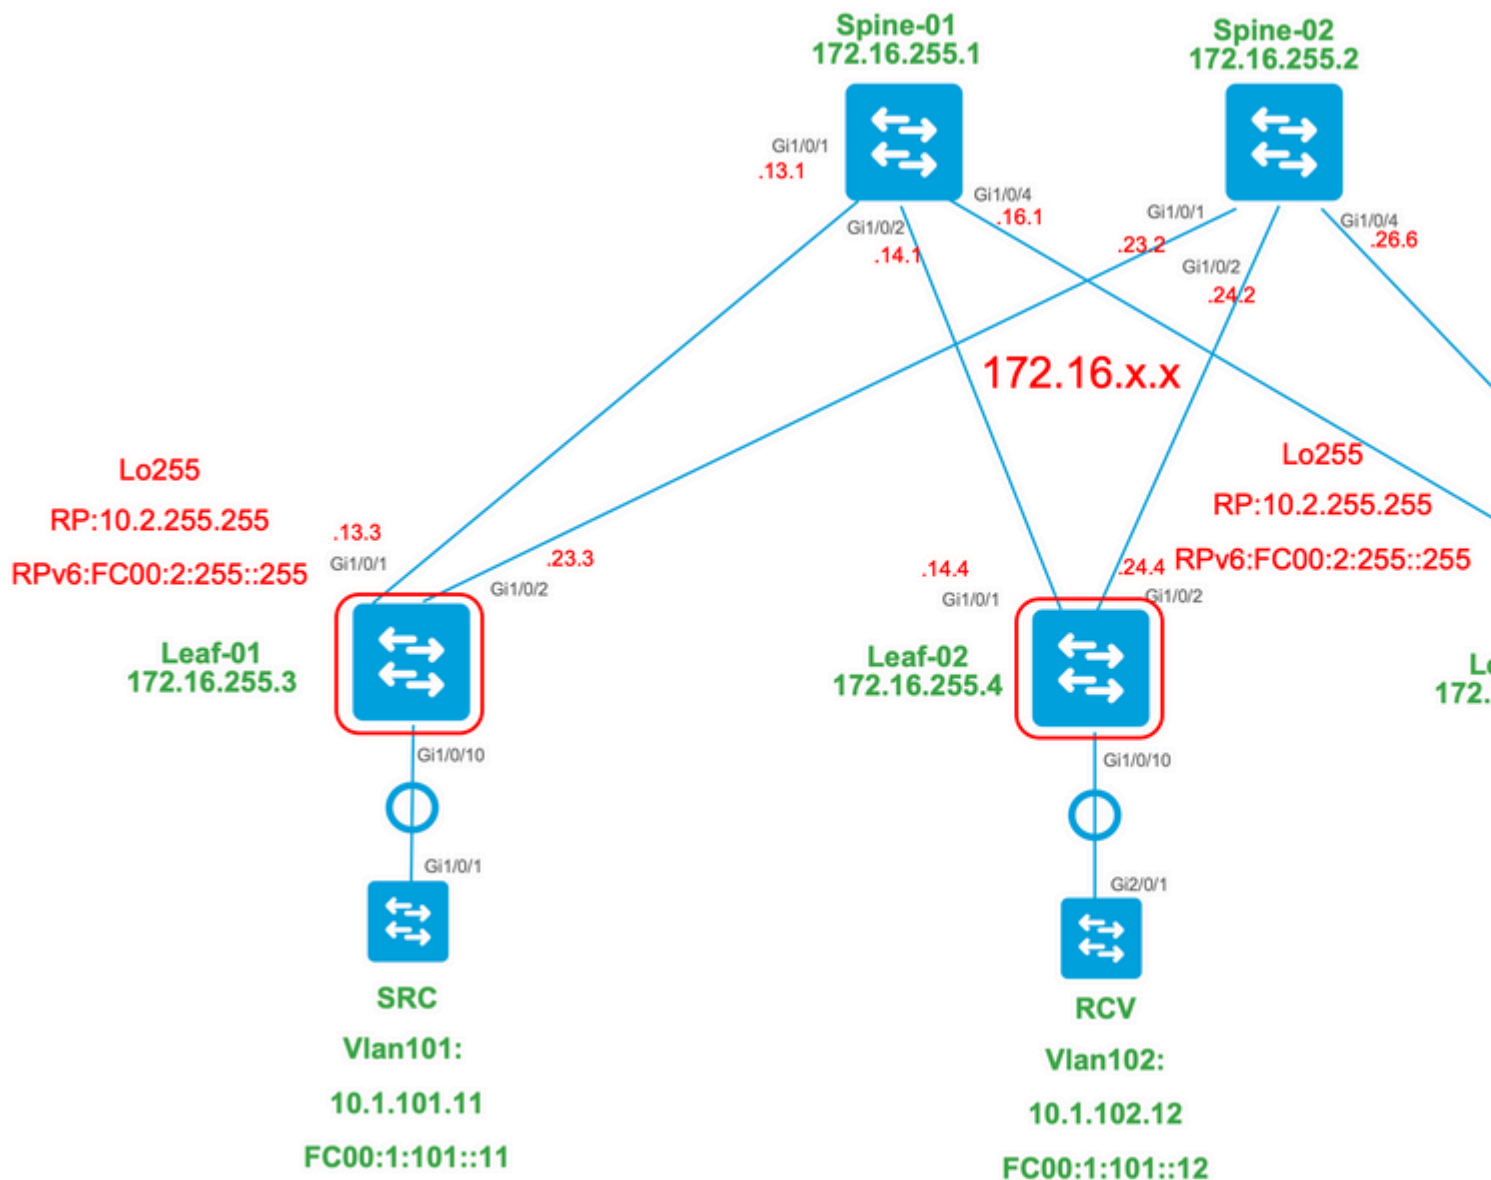
```
0)
ri[0]=0xE803
Dynamic port=88ri_ref_count:1 dirty=0
Leaf-02#
sh platform hardware fed sw active fwd-asic resource asic all rewrite-index range 0xE803 0xE803

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

Scenario 1. AnyCast RP (alleen SPT-bomen) IPv4 en IPv6

In deze modus bevindt zich op elke VTEP een RP. Deze VTEP's synchroniseren geen geleerde bronnen via MSDP en er is geen gedeelde boom. In plaats daarvan maakt de MDT-modus gebruik van BGP-informatie om alleen SPT-multicastbomen te maken. Deze modus wordt onderling verwisselbaar aangeduid als SPT-only-modus of gedistribueerde Anycast-RP-modus. In deze modus is elke VTEP de PIM RP. Zo wordt de (*,G) boom op elke plaats ingekort bij de lokale VTEP zelf. Het is niet nodig om (*,G) toe te voegen of MVPN RT-6 over de stof te sturen.

Netwerkdigram



Voor deze modus kunt u rekening houden met 3 BGP-routetypen:

1. EVPN-routetype 2. Dit staat andere PE's toe die een route C-Multicast (MVPN type6/7) terug naar de oorsprong moeten construeren PE, om juiste RT van de Invoer van C-multicast vast te maken zodat de initiator PE de route C-Multicast (RFC 6514 11.1.3) kan invoeren [RFC6514]. Het gebruik van deze VRI is afhankelijk van de opdracht "**mdt overlay use-bgp**" VRF-opdracht.
2. MVPN router-type 5. Dit is hetzelfde als in MVPN, en is de advertentie van een multicast bron/groep beschikbaar
3. MVPN router-type 7. Informatie van de IGMP of MLD-laag en van EVPN Type-2 worden gebruikt om dit BGP-type aan te maken. De Type-7 drijft de creatie van MRIB OIF aan de bronkant aan.

Eisen van EVPN Type 2:

1. Direct Connected multicast-bron wordt online geleverd.
2. FHR (bron-VTEP) verifieert ARP (of ND) en CEF-nabijheid (bevestigt dat de bron direct verbonden is).
3. FHR levert de EVPN Type-2 BGP update

MVPN Type-5 vereisten:

1. Het vereiste voor de Direct Connect-bron is opgelost
2. RP is lokaal dus FHR registreert bij zichzelf
3. FHR genereert MVPN Type-5 BGP-update

MVPN Type-7 vereisten:

1. Er is een EVPN Type-2 vermelding aanwezig (vereist voor het construeren van de C-Multicast-route type-7 met de juiste VRI en verzonden vanuit Source VTEP)
2. De MVPN Type-5-ingang is aanwezig (vereist om op te lossen welk Source/Group-paar beschikbaar is voor SPT-verbinding)
3. IGMP- of MLD-lidmaatschapsrapport is ontvangen en verwerkt door de LHR VTEP
4. LHR VPN RPF-interface is de Fabric L3VNI-interface

Tip: Bij de uitgang LHR VTEP PIM controleert het pad naar de bron. PIM moet een route in de RIB vinden die de L3VNI als RPF-interface is. Als L3VNI niet goed is geconfigureerd, is down enzovoort. VTEP probeert niet om de type-7 BGP-verbinding te maken.

Controleer BGP EVPN- en MVPN-routers

Controleer Leaf-01: de EVPN Type-2 is gemaakt

```
<#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, 13-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, 13-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

Controleer Leaf-01: ARP/IPv6 ND en EVPN Debugs tonen ARP/ND wordt geleerd, dan Route-type 2 gemaakt en verzonden

<#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][10.1.1.1]
*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][FC00:1:101::11]/42
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

```

Controleer Leaf-02: Source side **Route-type 2** wordt geleerd in BGP aan ontvangerzijde

```

<#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 Lol
```

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

Controleer blad-02: Source Route-type 5 wordt geleerd in BGP op ontvanger 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

Controleer Leaf-02: heeft BGP info van Leaf-01 nodig om Type-7 te creëren. Uiteindelijk is IGMP of MLD heeft een lidmaatschapsrapport verwerkt dat de VTEP informeert dat er een geïnteresseerde Ontvanger is.

<#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 V1102
```

```
<-- 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
```

Controleer Leaf-02: MVPN Debugs tonen Route-type 7 wordt gecreëerd wanneer IGMP/MLD lidmaatschapsrapport aankomt en de vereiste EVPN Type-2 en Type-5 zijn al geïnstalleerd.

```
<#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, 1:1:10.1.101.11/32

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,
```

Controleer Leaf-01: De MVPN Type-7 ontvangen van 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 MVPN
```

```
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
```

Controleer Leaf-01: MVPN Debugs tonen Route-type 7 ontvangen met MVPN VRI Route-Target

```
<#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)

Controleer Leaf-02: Complete BGP-tabel bevat de Leaf-01 EVPN Type-2 en MVPN Type-5, en de Type-7 gegenereerd door 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 Lol
-----	--------------	---	-----	-----	----------------------

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 L01

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-01

*>

[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

Controleer de TRM-groep Leaf-01 (FHR).

Controleer of de MDT- en TRM-groepen op de juiste wijze aan de bronkant zijn gevormd.

- Inkomende interface van TRM-groep is de SVI die aan de client-VRF is gekoppeld
- Uitgaande interface van de TRM-groep is de L3VNI SVI

Controleer Leaf-01: de TRM-groep 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 hardware

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 hardware

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

Controleer Leaf-01: de TRM-groep in de 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 : {Sv1}

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 local

OIF :

Vlan101 A {Remote}

Vlan901 F {Remote}

<-- Forward via L3VNI interface

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

Controleer blad-01: de nabijheid is juist

<#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)

Controleer TRM-groep Leaf-02 (LHR).

Controleer of de MDT- en TRM-groepen aan de ontvangerzijde correct zijn gevormd.

- Inkomende interface van de TRM-groep is de SVI gekoppeld aan de L3VNI
- De uitgaande interface van de TRM-groep is de client SVI waar de IGMP-verbinding is verwerkt.

Controleer Leaf-02: de TRM (Tenant multicast route) route in 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
```

Controleer Leaf-02: de TRM-groep in de 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: {Sv1}
```

```
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_MULTICAST
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: {Sv1}

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_MULTICAST
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

Controleer Leaf-02: pakketopname toont de buitenste MDT-tunnelgroep met binnenste clientverkeer

```
<#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...>
```


Scenario 2: PIM SSM in de fabric

In deze modus is er geen RP in de Overlay en worden er geen MVPN Type-5 of Type-7 gebruikt (de Underlay blijft functioneren als PIM ASM). **In SSM, verzendt de ontvanger en IGMPv3 S,G toetreden** naar LHR VTEP. Dit VTEP voert RPF-raadpleging uit voor de bron in de RIB. Als L3VNI SVI wordt gevonden als de RPF-interface, stuurt de LHR VTEP MVPN RT-7 naar de FHR VTEP die deze route ontvangt en installeert. FHR VTEP informeert PIM vervolgens om L3VNI SVI toe te voegen als de Uitgaande interface voor de S,G route.

Dit deel laat de verschillen zien met scenario 1. De stappen en methoden die hetzelfde zijn worden alleen in scenario 1 genoteerd.

- Zie verificatie en debug stappen voor BGP en PIM van scenario 1, omdat de BGP- en PIM-bewerkingen hetzelfde zijn

Netwerkdigram



Voor deze modus kunt u rekening houden met deze BGP-routetypen en hun oorsprong

Gemaakt door: Source VTEP

- EVPN-routetype 2. Wordt gebruikt om Unicast- en VRI-informatie voor de bron te verkrijgen en toegevoegd aan de C-Multicast-route (MVPN type-7) wanneer de VTEP lid wordt van de STP-structuur.

Gemaakt door: ontvanger VTEP

- MVPN router-type 7. Informatie van de IGMP of MLD-laag en van EVPN Type-2 worden gebruikt om dit BGP-type aan te maken. De Type-7 drijft de creatie van MRIB OIF aan de bronkant aan.

Eisen van EVPN Type 2:

1. FHR (bron-VTEP) verifieert ARP (of ND) en CEF-nabijheid (bevestigt dat de bron direct verbonden is).

2. FHR levert de EVPN Type-2 BGP update

MVPN Type-7 vereisten:

1. Er is een EVPN Type-2 vermelding aanwezig (vereist voor het construeren van de C-Multicast-route type-7 met de juiste VRI en verzonden vanuit Source VTEP)
2. Ontvanger VTEP: IGMPv3 Het bronspecifieke lidmaatschapsrapport is ontvangen en verwerkt door LHR VTEP
3. LHR VPN RPF-interface is de Fabric L3VNI-interface

Voor deze modus is er extra configuratie nodig op de LHR VTEP om SSM-bereik mogelijk te maken en IGMPv3-lidmaatschapsrapporten te verwerken

Configureer Leaf-03: stel de IGMP querier in op Versie 3 onder de huurder SVI

```
<#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
```

Controleer Leaf-03: de IGMP-query is ingesteld op versie 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
```

Schakel Leaf-03 in: het SSM-bereik dat voor de huurder VRF vereist is

```
<#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
```

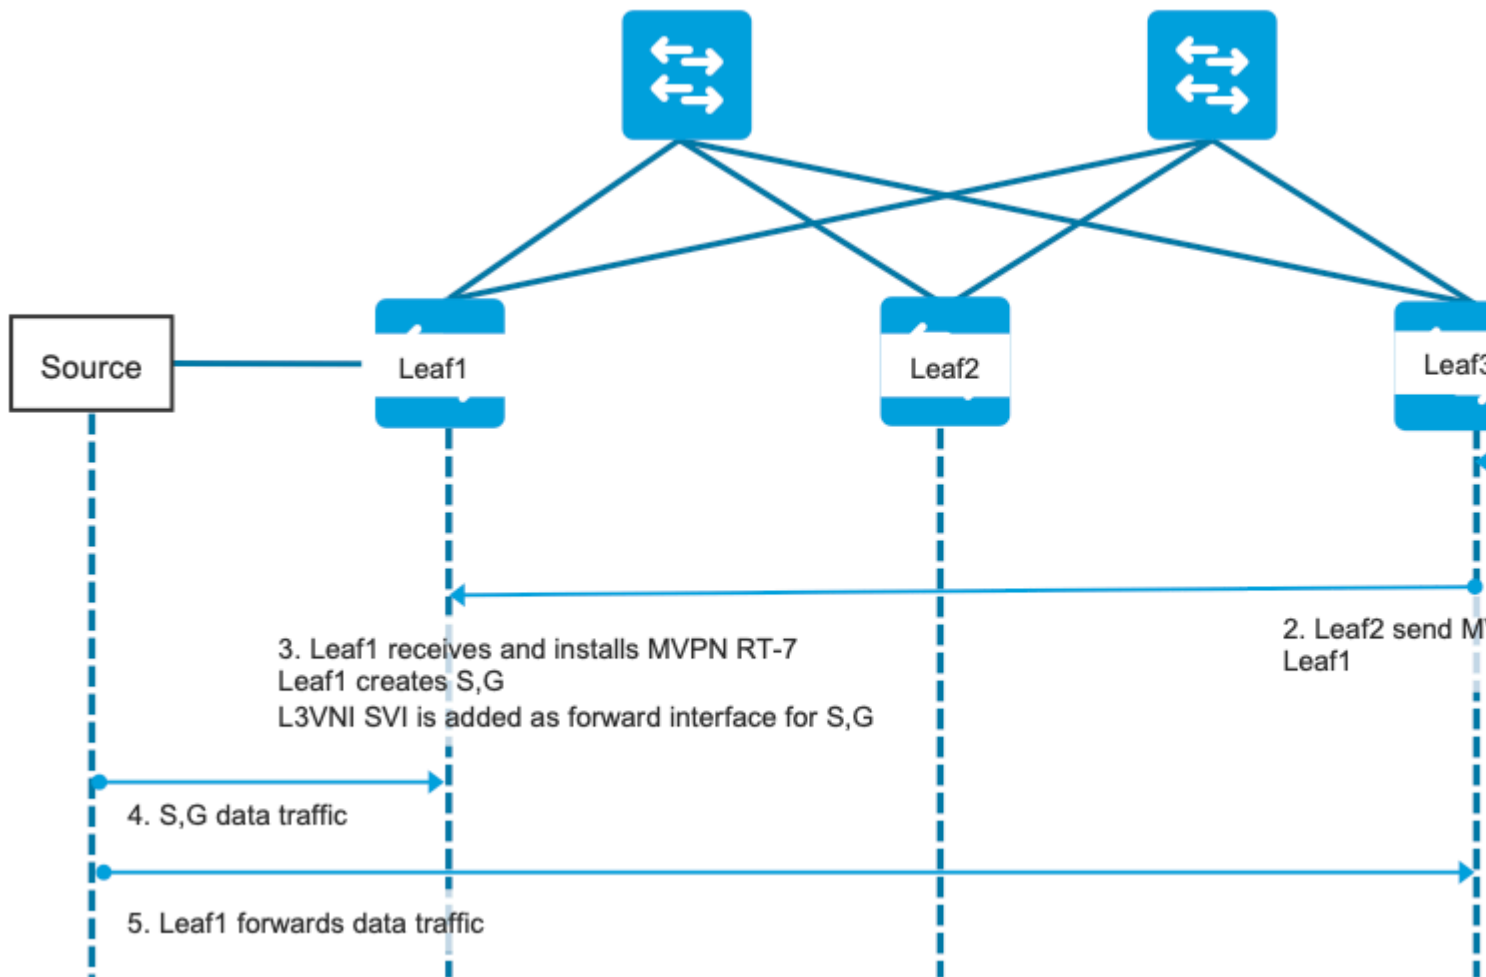
```
ACL for group range
```

```
to be used for SSM
```

```
<-- use an ACL to define a non-default SSM range
```

Tip: SSM-groepen maken geen *,G-route. Als u * ziet, verifieert G voor de groep, uw configuratie correct voor SSM is.

Controleer de volgorde van gebeurtenissen die voor dit scenario vereist zijn



Step 0 EVPN (Leaf-03): Controleer of er een EVPN prefix is dat BGP kan vinden de VRI te gebruiken in de MVPN type-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

```

Step 1 (Leaf-03): IGMPv3 Membership-rapport ontvangen en bevat een bron

<#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

```

Stap 2 (Leaf-03): BGP is op de hoogte van deze koppeling, maakt en verstuurt de Type-7 MVPN-koppeling.

```
<#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 ty
```

```

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

Step 3 (Leaf-01): Source Leaf ontvangt en installeert MVPN Type-7 toetreden route, en informeert PIM om L3VNI OIF te installeren

<#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 Acc

Leaf-01#

sh bgp ipv4 mvpn all

<...snip...>

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

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

172.16.255.6

0 100 0 ?

<-- Recieved from Reciever Leaf-03

* i	172.16.255.6	0	100	0 ?
-----	--------------	---	-----	-----

Leaf-01#

```
sh ip mroute vrf green 226.1.1.1
```

```
<...snip...>
```

```
(10.1.101.11, 226.1.1.1), 00:42:41/stopped, flags: sTGx
```

```
<-- s = SSM Group, G = Received BGP
```

```
Incoming interface: Vlan101, RPF nbr 10.1.101.11
```

```
Outgoing interface list:
```

```
Vlan901, Forward/Sparse, 00:42:41/stopped
```

```
<-- L3VNI installed as OIF interface
```

Stap 4 & 5 (Leaf-01 & Leaf-03): Multicast komt naar het FHR-blad en wordt via stof naar LHR-blad gestuurd. Samenvatting van de hier gegeven valideringsopdrachten. U kunt de gedetailleerde validatie van deze opdrachten controleren in Scenario 1.

```
<#root>
```

```
show ip mroute vrf green 226.1.1.1 count
```

```
<-- software m
```

```
show ip mfib vrf green 226.1.1.1
```

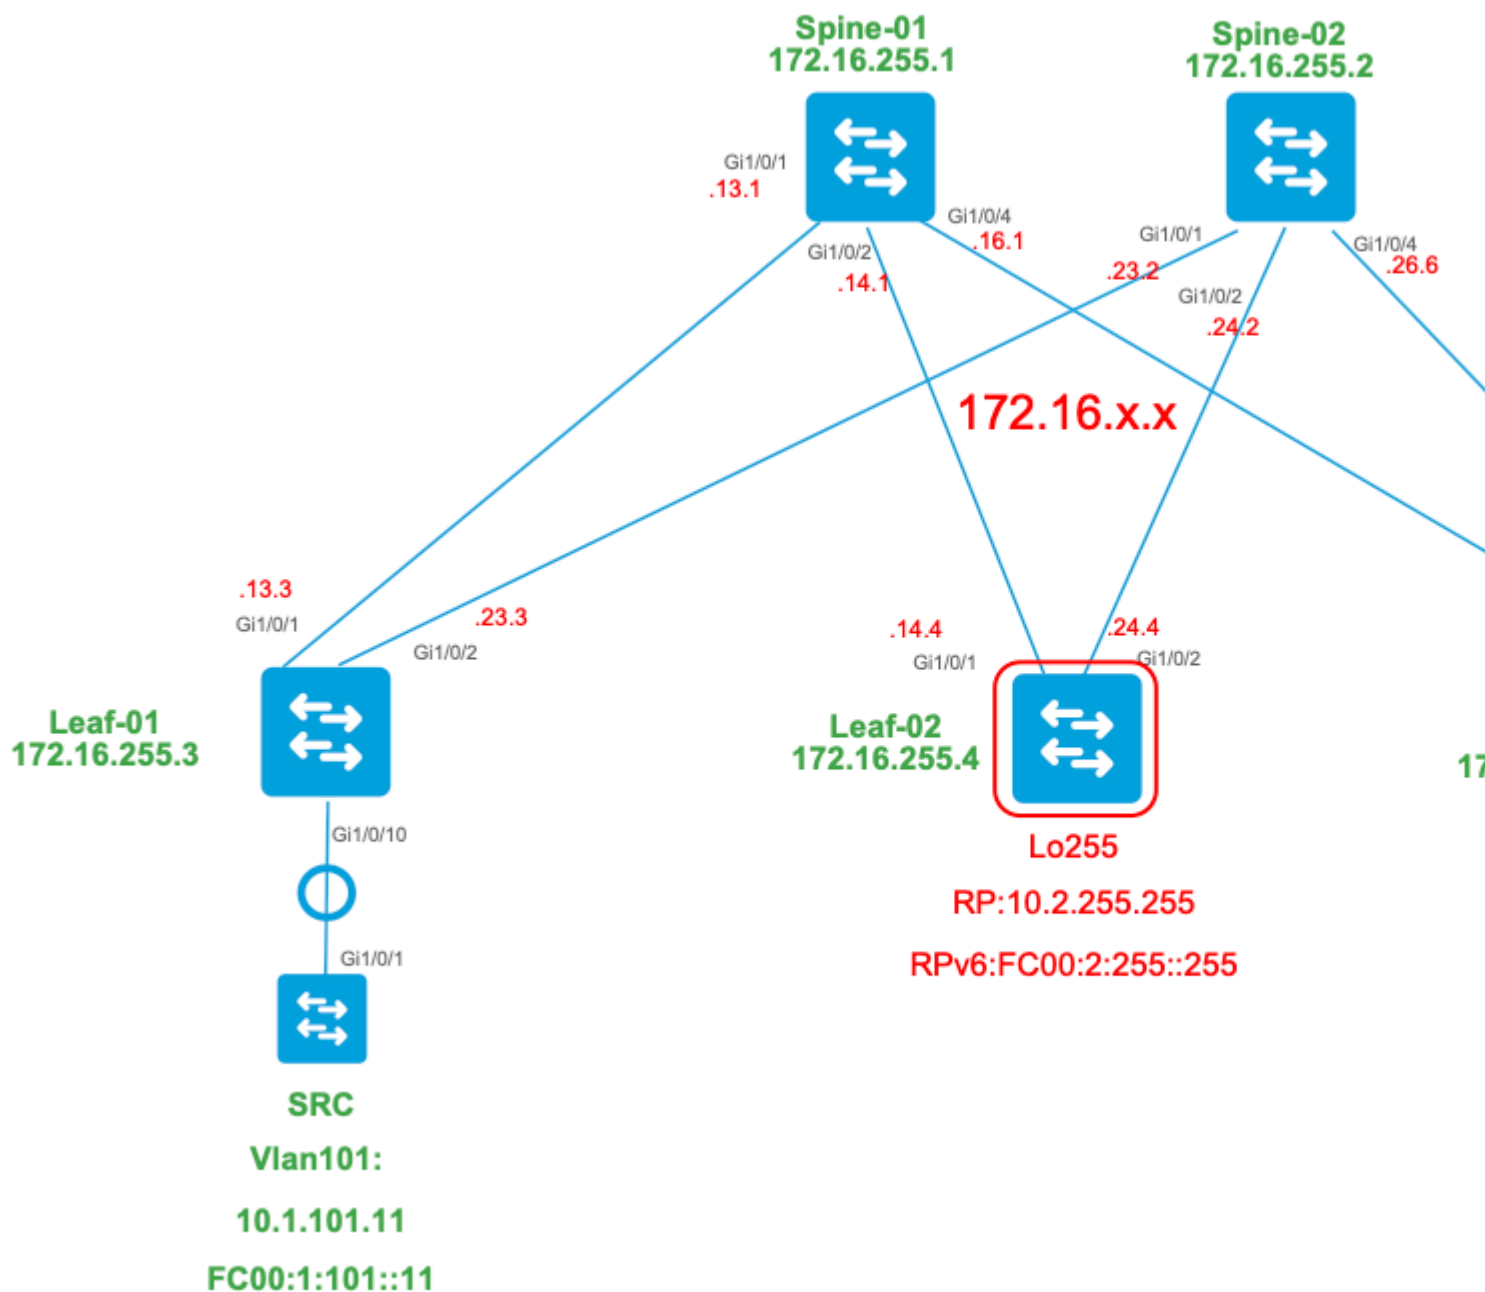
```
<-- hardware mroute details & counters
```

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

Scenario 3: Single RP Inside the Fabric (Regular Sparse-Mode)

Deze modus wordt onderling verwisselbaar aangeduid als niet-Anycast RP of externe RP-modus. In deze modus is er slechts één RP in de overlay. Zodoende kan een (*,G)-boom in de overlay over meerdere locaties lopen. BGP maakt gebruik van een MVPN RT-6 om (*,G)-lidmaatschap te adverteren over de stof. Als RP en FHR zich op verschillende locaties bevinden, worden PIM-registers over de stof verzonden. Dit is de standaard operationele modus voor PIM SM in de overlay.

Netwerkdigram



Voor deze modus kunt u rekening houden met deze BGP-routetypen en hun oorsprong

Gemaakt door: Source VTEP

- EVPN-routetype 2. Wordt gebruikt om Unicast- en VRI-informatie voor de bron te verkrijgen en toegevoegd aan de C-Multicast-route (MVPN type-7) wanneer de VTEP lid wordt van de STP-structuur.
- MVPN router-type 5. Bron-A-D route naar VTEP's verzonden voor S,G

Gemaakt door: RP VTEP

- EVPN Route-type 5. Gebruikt om Unicast en VRI info voor RP loopback te verkrijgen. Loopback maakt geen Route-type 2, dus type 5 wordt gebruikt.
- MVPN router-type 7. Dit is de IGMP-koppeling + RT VRI-gegevens van de EVPN Type-2 en verzonden naar de bron VTEP, en het stimuleert de creatie van de MRIB OIF.

Gemaakt door: ontvanger VTEP

- MVPN router-type 6. Route-type gemaakt door ontvanger VTEP om de gedeelde boom *,G (RPT-boom) naar de RP toe te voegen.
- MVPN router-type 7. Informatie van de IGMP of MLD-laag en van EVPN Type-2 worden gebruikt om dit BGP-type aan te maken. De Type-7 drijft de creatie van MRIB OIF aan de bronkant aan.

Eisen van EVPN Type 2:

1. FHR (bron-VTEP) verifieert ARP (of ND) en CEF-nabijheid (bevestigt dat de bron direct verbonden is).
2. FHR levert de EVPN Type-2 BGP update

Eisen van EVPN Type 5:

1. RP loopback wordt geconfigureerd en geadverteerd in BGP

MVPN Type-5 vereisten:

In deze modus adverteert Leaf op de bron-site alleen actieve A-D-berichten voor een (S, G) als aan deze twee voorwaarden is voldaan.

1. Het ontvangt verkeer op de RPF-interface naar de Bron. (de bron wordt verzonden mcast naar de FHR)
2. L3VNI SVI-interface wordt toegevoegd als een voorwaartse interface voor (S,G)-ingang, als resultaat van een S,G-verbinding van de RP als deel van het PIM-registratieproces. (Het L3VNI SVI wordt in de OIF-lijst geïnstalleerd)

MVPN Type-6 vereisten:

1. RP adverteert zijn EVPN Type-5 route die zijn VRI en Unicast bereikbaarheidsdetails bevat.
2. IGMP wordt ontvangen via LHR, waardoor een BGP-update naar RP wordt geactiveerd

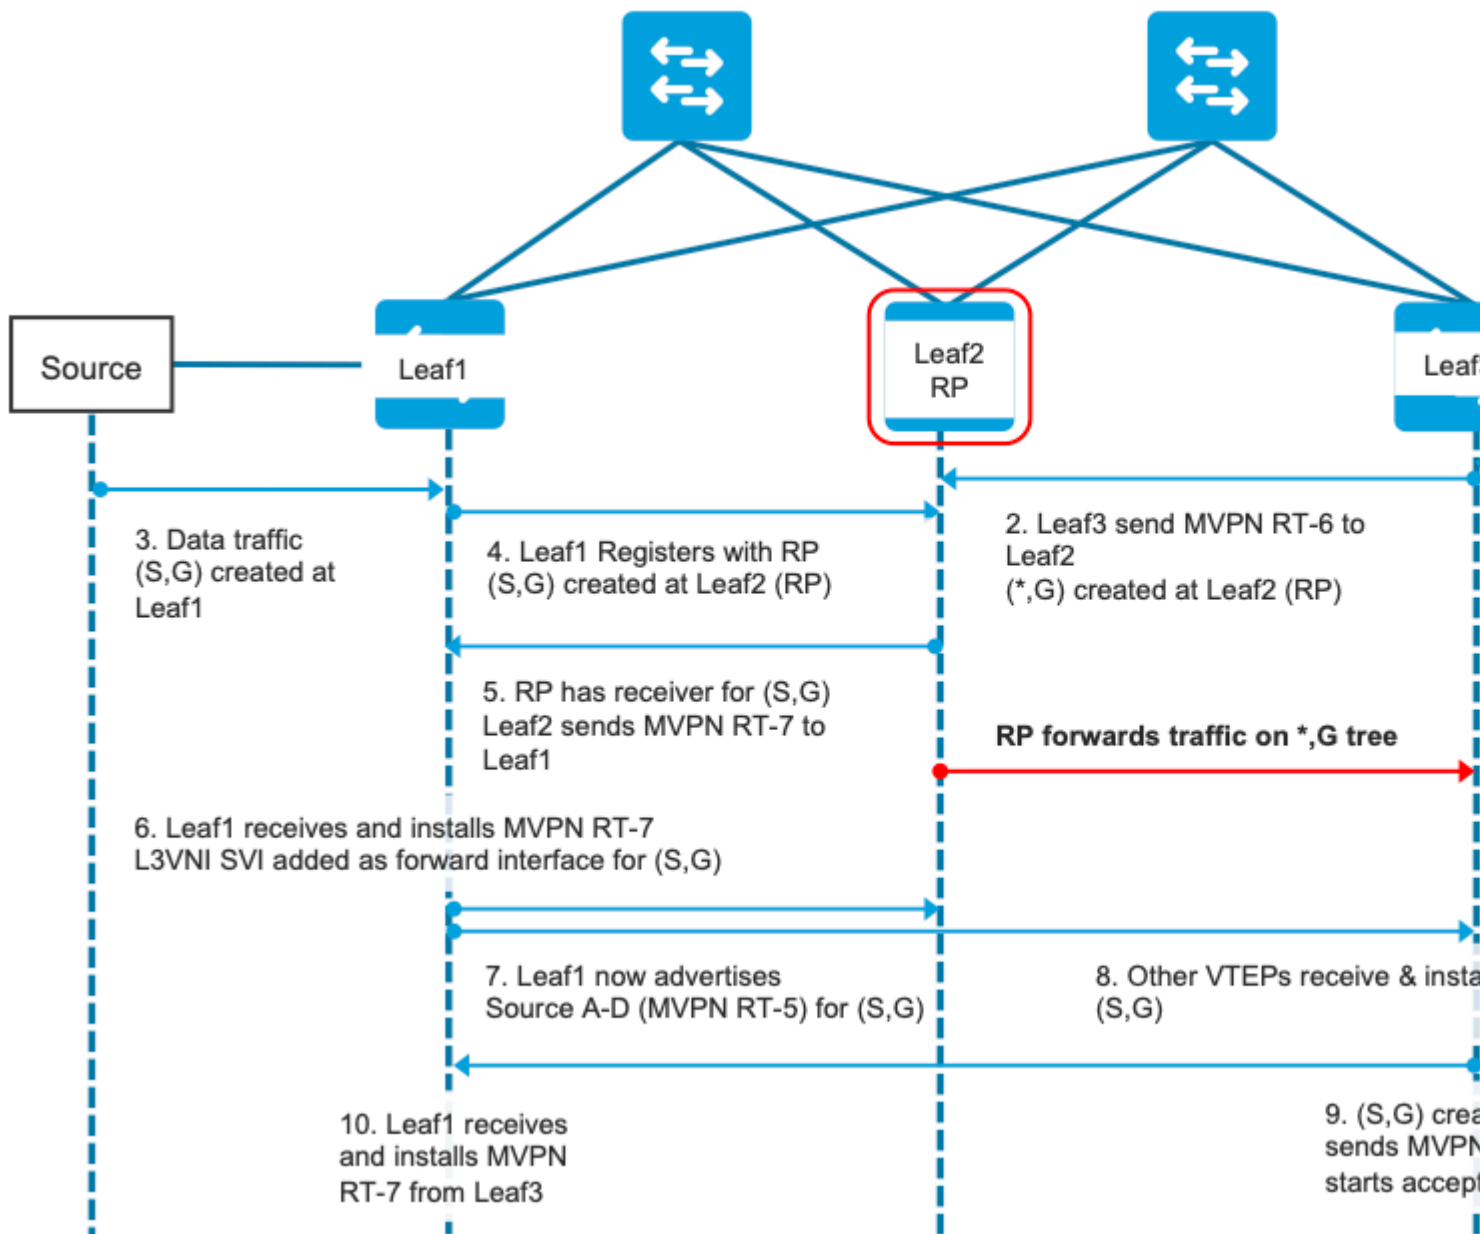
MVPN Type-7 vereisten:

1. Er is een EVPN Type-2 vermelding aanwezig (vereist voor het construeren van de C-Multicast-route type-7 met de juiste VRI en verzonden vanuit Source VTEP)
2. De ingang van MVPN type-5 is aanwezig (vereist om te besluiten welk bron/groepspar voor STP beschikbaar is toetreden)
3. Ontvanger VTEP: IGMP-lidmaatschapsrapport is ontvangen en verwerkt door de LHR VTEP
4. RP VTEP: RP heeft multicast Register-pakketten ontvangen, heeft EVPN-routes, en heeft een ontvanger voor S,G (geleerd via de type-6)
5. LHR VPN RPF-interface is de Fabric L3VNI-interface

Tip: Bij de uitgang LHR VTEP PIM controleert het pad naar de bron. PIM moet een route in de RIB vinden die de L3VNI als RPF-interface is. Als L3VNI niet goed is geconfigureerd, is down enzovoort. de VTEP maakt geen type-7 BGP-verbinding.

Controleer de volgorde van gebeurtenissen die voor dit scenario vereist zijn

Valideer de stappen die nodig zijn om de VTEP-ontvanger in eerste instantie te laten toetreden tot de gedeelde boom en vervolgens over te schakelen naar de kortste padboom. Dit omvat controles van de BGP-tabellen, IGMP en MRIB-ontwerpstaten.



Step EVPN (Leaf-03): EVPN Type-5 van RP wordt geleerd op LHR. Dit is vereist om een MVPN Type-6 route te maken voor de ontvanger VTEP

```
<#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

Stap 1 (Leaf-03): IGMP-lidmaatschapsrapport ontvangen

<#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

Stap 2 (Leaf-03): MVPN Type-6 gemaakt, verzonden naar RP en ontvangen door RP (Leaf-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

Step 1 & 2 Debugs (Leaf-01): IGMP-rapport, EVPN-bronraadpleging en MVPN Type-6 maken

<#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][226.1.1.1/32]/22
*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][65001][10.2.255.255/32]
*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)

```

Stap 3 & 4 (Leaf-01): Vanuit FHR-perspectief valideren van de S, G creëren en registreren gebeurtenissen (S, G creëren & registreren gebeurt bijna op hetzelfde moment)

3. Het dataverkeer begint en S,G wordt gecreëerd bij FHR VTEP. De vereisten die in de sectie "Ongedetecteerde multicastbronnen" zijn opgenomen, zijn hier van toepassing.

4. Leaf-01 voert bronregistratie aan RP via zijn PIM-tunnel uit

```
<#root>
```

```
Leaf-01#
```

```
debug ip pim vrf green 226.1.1.1
```

```
PIM debugging is on
```

```
Leaf-01#
```

```
debug ip mrouting vrf green 226.1.1.1
```

```
IP multicast routing debugging is on
```

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

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

```
*Jan 29 18:18:58.426:
```

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

```
*Jan 29 18:18:58.427:
```

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

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

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

```
*Jan 29 18:18:58.428:
```

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

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

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

```
Leaf-01#
```

```
sh int tunnel4
```

```
Tunnel4 is up, line protocol is up
```

```
  Hardware is Tunnel
```

```
  Description:
```

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

```
  Interface is unnumbered.
```

```
  Using address of Loopback901 (10.1.255.1)
```

```
  <-- Local Loopback
```

S,G is created when Source sends data traffic

Leaf-01#

sh ip mroute vrf green 226.1.1.1

IP Multicast Routing Table

<...snip...>

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join

Timers: Uptime/Expires

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

(* , 226.1.1.1), 00:00:16/stopped, RP 10.2.255.255, flags: SPF

Incoming interface: Vlan901, RPF nbr 172.16.254.4

Outgoing interface list: Null

(10.1.101.11, 226.1.1.1)

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

Incoming interface: Vlan101

,

RPF nbr 10.1.101.11

,

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

Outgoing interface list:

Vlan901

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

<-- OIF is the L3VNI SVI

Checking S,G in Hardware

Leaf-01#

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

MROUTE ENTRY

vrf 2

(10.1.101.11, 226.1.1.1/32)

<-- VRF 2 is the ID for vrf green

HW Handle: 140213987784872 Flags: {Sv1}

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 I

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

Stap 4 (Leaf-02): Vanuit RP-perspectief, bevestig bronregistratie bereikt de RP en S,G wordt gecreëerd.

```
<#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-02
```

```
*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 (10.1.101.11)
```

```
*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

Step 5 (Leaf-02): RP heeft een ontvanger, dus direct gemaakt Type-7 MVPN-bronboom Join route

<#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

Stap 6 (Leaf-01): Source Leaf-01 ontvangt en installeert MVPN Route-Type 7. (L3 VNI SVI wordt geïnstalleerd als een voorwaartse interface voor 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][10.1.101.11/32]

*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

Stap 7 (Leaf-01): Leaf-01 adverteert met MVPN-bron A-D Type-5 voor 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, exten

Stap 8 (Leaf-03): Ontvanger VTEP krijgt de Type-5 en installeert de Source A-D route voor S,G

<#root>

Leaf-03#

```
debug bgp ipv4 mvpn updates
```

```
BGP updates debugging is on for address family: MVPNV4 Unicast  
Leaf-03#
```

```
debug bgp ipv4 mvpn updates events
```

```
BGP update events debugging is on for address family: MVPNV4 Unicast
```

```
*Jan 29 19:18:53.318: BGP(15): 172.16.255.1 rcvd UPDATE w/ attr: nexthop 172.16.255.3, origin ?, localp
```

```
*Jan 29 19:18:53.319: BGP(15): 172.16.255.1 rcvd [5][1:1][10.1.101.11][226.1.1.1]/18 <-- Type-5 Receiv
```

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

```
Leaf-03#
```

```
sh bgp ipv4 mvpn all route-type 5 10.1.101.11 226.1.1.1  
...or you can also use:
```

```
Leaf-03#
```

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

```
BGP routing table entry for
```

```
[5][1:1][10.1.101.11][226.1.1.1]/18
```

```
, version 41
```

```
<-- Type-5 A-D route from Leaf-01
```

```
Paths: (2 available, best #2, table MVPNV4-BGP-Table, not advertised to EBGp peer)
```

```
Flag: 0x100
```

```
Not advertised to any peer
```

```
Refresh Epoch 1
```

```
Local
```

```
172.16.255.3
```

```
(metric 3) from 172.16.255.1 (172.16.255.1)
```

```
<-- Leaf-01 IP
```

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

```
Community: no-export
```

```
Extended Community: RT:1:1
```

```
Originator: 172.16.255.3
```

```
, Cluster list: 172.16.255.1
```

```
rx pathid: 0, tx pathid: 0x0
```

```
Updated on Jan 29 2021 19:18:53 UTC
```

Stap 9 (Leaf-03): S,G is gemaakt, Leaf-03 stuurt MVPN Type-7 om zich aan te sluiten bij SPT-boom en begint verkeer te accepteren

<#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

Stap 10 (Leaf-01): Leaf-01 ontvangt en installeert MVPN Type-7 van 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][10.1.101.11/32]
```

```
*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
```

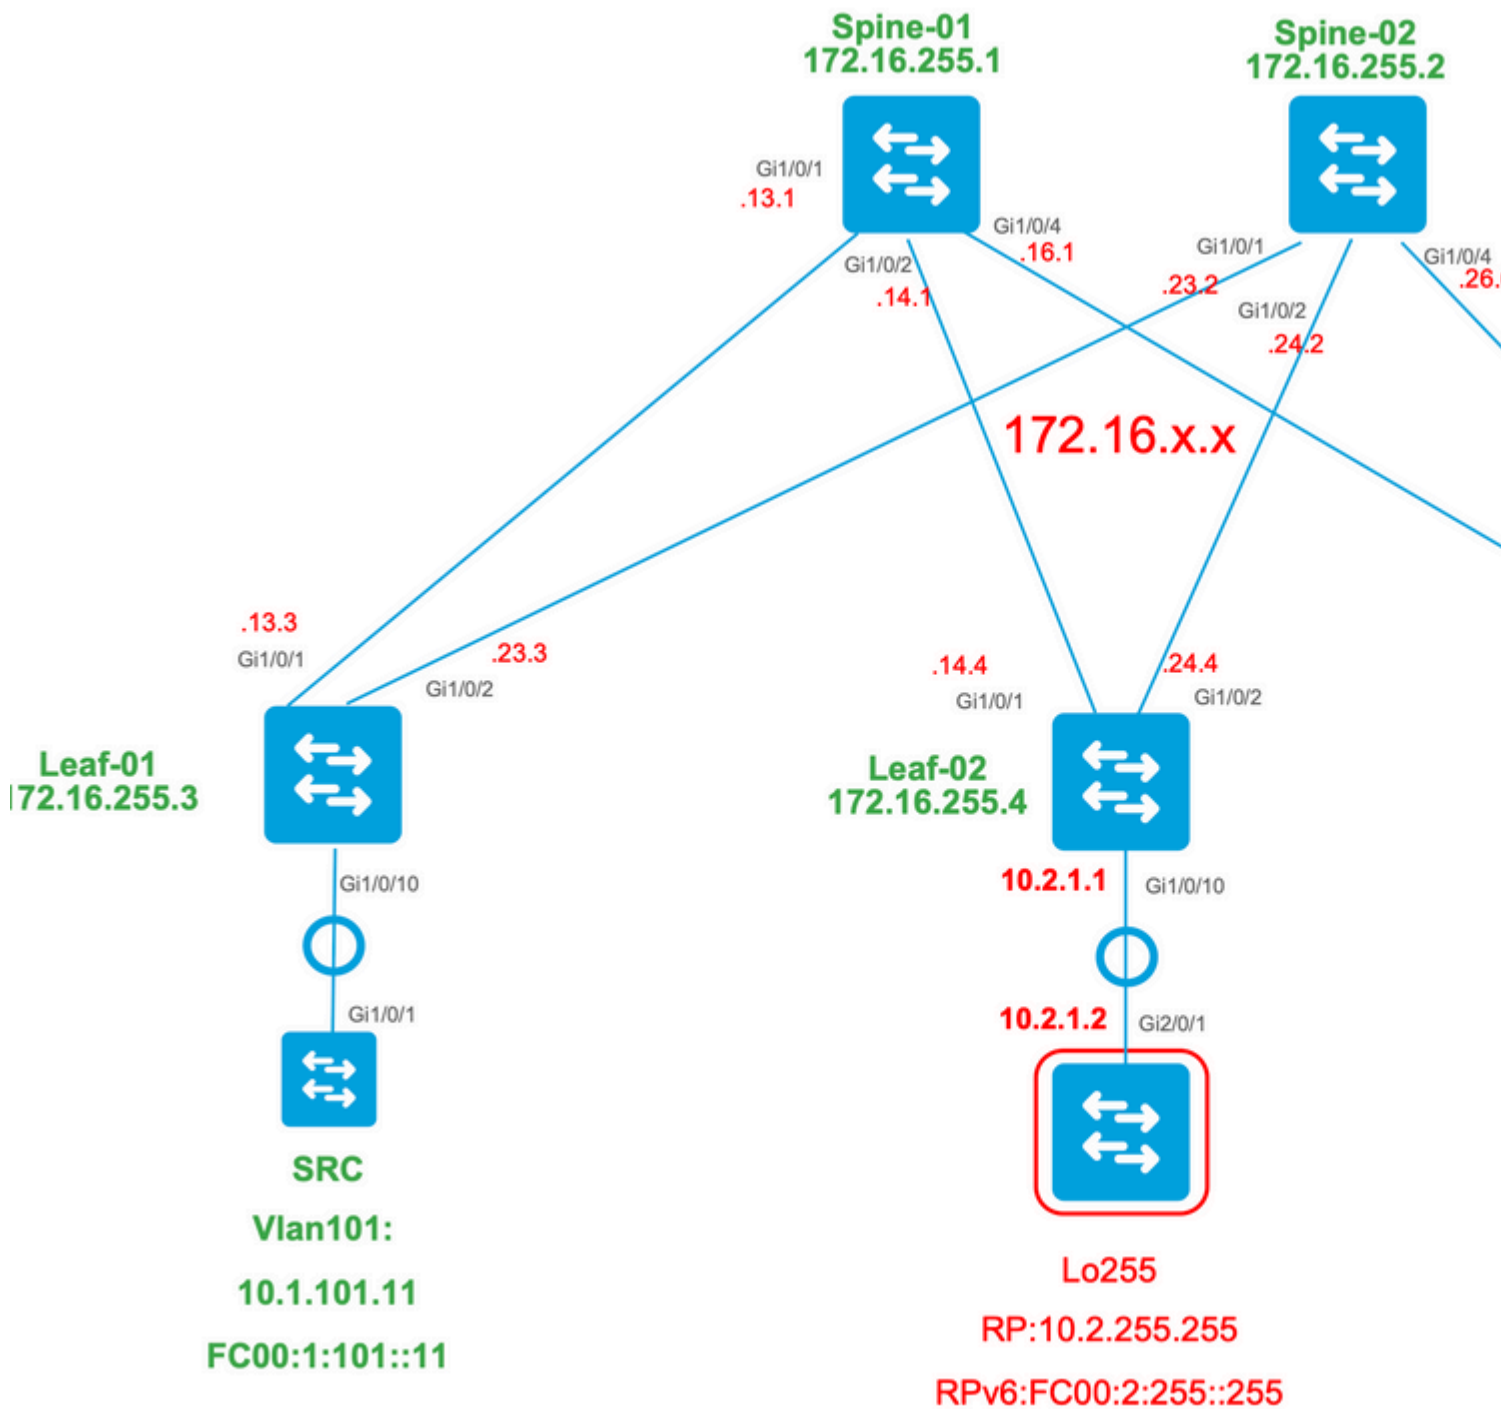
Scenario 4: RP Buiten de Stof (RP geïmporteerd uit Border Leaf-02 vanuit IP-ruimte)

Dit scenario is in wezen hetzelfde als scenario 2. Er is een enkele RP gebruikt door de Fabric in zijn geheel. Het verschil is dat de RP IP moet worden geïmporteerd van een niet-fabric IP-ruimte in Fabric en geadverteerd in BGP.

Dit deel laat de verschillen zien met scenario 3. De stappen en methoden die hetzelfde zijn, worden alleen in scenario 3 vermeld

- Zie **Controleer de volgorde van gebeurtenissen die voor dit scenario van scenario 3 vereist zijn** omdat de BGP- en PIM-bewerkingen hetzelfde zijn

Netwerkdigram



Controleer de invoer van border-switch van IP naar fabric

Het belangrijkste verschil met dit ontwerp ten opzichte van scenario 3 is de noodzaak om eerst de RP IP van de IP ruimte naar EVPN te importeren.

De rand moet bepaalde opdrachten bevatten om te importeren/exporteren naar en van weefsel- en IP-ruimtes:

- **route-doel <waarde>-stikkende** opdrachten onder het VRF-configuratiedeelte
- **adverteren l2vpn evpn** onder de BGP vrf-adresfamilie

Controleer (blad-02): configuratie

<#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' R
```

```
redistribute connected
redistribute static
```

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

```
exit-address-family
```

Verifiëren (Leaf-02): Prefix import en advertentie

<#root>

```
debug bgp vpv4 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. (M
```

```
*>
```

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

Controleer (Leaf-02): Rand pad naar 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 - encaps-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

Scenario 5: Data MDT

Controleer MDT-gegevensgroep

De MDT Data-groep is vergelijkbaar met andere MDT Default-groep, waar de buitentunnelgroep voor TRM wordt ingekapseld in. Echter, in tegenstelling tot de MDT Default, zal deze groep alleen VTEP's hebben die zich bij deze boom aansluiten als ze geïnteresseerde ontvangers hebben voor de TRM groep.

Vereiste configuratie

<#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
!
```

Controleer of de MDT-groep aan de bronkant correct is geprogrammeerd

- Inkomende interface van MDT groep is de bronkant Loopback
- Uitgaande interface van MDT-groep is de Underlay Interface

Controleer Leaf-01: de MDT-route is correct in 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
```

Controleer Leaf-01: FED-vermeldingen voor de MDT-groep

<#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_MULTICAST
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

=====

Controleer of de MDT-groep aan ontvangerzijde correct is geprogrammeerd

- Inkomende interface van MDT-groep is de RPF-interface terug naar de bronzijde Loopback
- Uitgaande interface van MDT-groep is Encap/Decap Tunnel-interface

Controleer Leaf-02: de MDT-route is correct in 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
```

Controleer Leaf-02: FED-vermeldingen voor de MDT-groep

```
<#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 platform software fed switch active fwd-asic resource asic all rewrite-index range 0xE803 0XE803
```

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

```
Rewrite_type
```

```
:AL_RRM_REWRITE_L2_PAYLOAD_
```

```
IPV4_EVPN_DECAP
```

```
(118) Mapped_rii:LVX_EVPN_DECAP(143)
```

```
<snip>
```

Debug MDT-gegevensgroep

Gebruik de debug MVPN om de Data MDT cutover gebeurtenis te controleren

Bron-VTEP

```
<#root>
```

```
Leaf#
```

```
debug mvpn
```

```
<snip>
```

```
*Mar 27 12:12:11.115: MVPN: Received local withdraw for (10.1.101.11, 239.1.1.1) with RD: 1:1, Route Type
```

```
*Mar 27 12:12:11.115: MVPN: Sending BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nh 0.0.0.0, W
```

```
*Mar 27 12:12:11.115: MVPN: Route Type 5 deleted [(10.1.101.11, 239.1.1.1), nh 0.0.0.0] rd:1:1 send:1
```

```
*Mar 27 12:12:11.115: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: 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, O
```

```
*Mar 27 12:13:00.431: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: 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 fordatabmt 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 ontvangerzijde

<#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: 172.16.255.3

*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 Type: 5

*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: 5

*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:0

*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: 5

*Mar 27 12:28:44.235: MVPN: Received BGP prefix=[5: 1:1 : (10.1.101.11,239.1.1.1)] len=19, nexthop: 172.16.255.3

*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 Type: 5

*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, nexthop: 172.16.255.3

*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, nexthop: 172.16.255.3

*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:172.16.255.3

*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:172.16.255.3

*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:172.16.255.3

Leaf-03#

Problemen oplossen

Ongedetecteerde multicastbronnen

Alvorens u bekijkt waarom een multicast stroom niet werkt, is het belangrijk om de relatie van ARP en multicast het door:sturen te begrijpen

Gewoonlijk wanneer een gastheer actief wordt en verkeer verzendt, worden ARP ingangen voltooid door de regelmatige bronopsporingsprocedures. Maar in het geval van multicast bronnen, is het mogelijk dat de bron begint om verkeer bij FHR verwerkt dit multicast verkeer zonder resolutie van ARP voor de bron.

ARP-voltooing speelt om twee redenen een belangrijke rol in de TRM-functie.

1. "**Direct connected**" controle bij de eerste hop router roept een FIB API die op zijn beurt afhankelijk is van ARP succesvolle controle. Als ARP naar de multicast bron niet is voltooid, blijft de CEF-nabijheid naar de bron onverbonden controle FALSE terug.
2. Source detectie activeert **reclame voor EVPN RT-2 in de EVPN stof**. Deze EVPN-route die in L3RIB bij een bron geïnstalleerd, wordt gebruikt als de RPF-route naar de bron. Dus, als de bron niet is gedetecteerd, kan RPF niet worden gevonden. In dit geval blijft ofwel de RPF NULL of wordt een minder specifieke route (indien aanwezig) gevonden.

Zorg ervoor dat ARP is opgelost en dat de Source bereikbaar is binnen de EVPN-stof.

Andere nuttige debugs

In deze sectie zijn andere debugs die kunnen helpen bij het isoleren van TRM problemen

- **debug mvpn** (alle MVPN-gebeurtenissen, zie bijvoorbeeld scenario 2)
- **debug van ip|ipv6 PIM <vrf>** (PIM-protocolactiviteit)
- **debug ip mrib <vrf> trans** (MRIB, klassieke PIM-vertaling)
- **debug ip mfib <vrf> pak|ps|fs** (Packet Forwarding| Process switching| Fast switching)

Bronnen en ontvangers buiten de fabric

In sommige gevallen kan de Bron en/of ontvanger een of meer L3 hop leven weg van de stof VTEP(s).

Dit is een geldig ontwerp, maar verandert welk EVPN routetype vervoerders de VRI, en welk proces verantwoordelijk is voor de verbindingen bij de Ontvanger VTEP.

- Als de bron buiten de stof is ziet de toegang VTEP de bron via een buur PIM, niet direct verbonden en verzendt het verkeer naar de ontvanger VTEP. De VRI is opgenomen in type 5.
- Als de ontvanger buiten de stof is, komt de koppeling via een PIM-koppeling bij IGMP. Informatie in de PIM wordt gebruikt om de MVPN Type-7 te maken.

eBGP meervoudige AS (spine-naar-spine) topologie

In sommige gevallen kan de topologie van BGP vereisen om updateinformatie naar een andere AS/Fabric te

Het is mogelijk dat de informatie over het BGP-besturingsplane gedurende maximaal 30 seconden samenkomt en d

- **Dit komt door het standaard eBGP-advertentieinterval van 30 seconden.**
- Als er een probleem is met lange convergentietijden als gevolg van vertraging in BGP-updates, kan het worden verkort om updates vaker te verzenden.
- Raadpleeg de BGP-configuratiegids in het referentiegedeelte van dit artikel voor meer informatie over

eBGP inter-as **vereist** een extra opdracht

Gebruik het inter-as sleutelwoord voor de MVPN adresfamilieroutes om de BGP autonome systeem (AS) grenzen t

```
<#root>
```

```
Border-Leaf(config-vrf-af)#
```

```
mdt auto-discovery vxlan inter-as
```

Tunnel registreren met symmetrische L2VNI (FHR vastgelegd in PIM-registratiestaat)

In gevallen waarin het VNI bestaat op de FHR en op andere VTEP's, is het mogelijk om de FHR vast te zette

Dit komt doordat de PIM Register Tunnel bron IP de AnyCast gateway is. Wanneer de RP een PIM-register ontvan
juiste VTEP is om de register stop te verzenden, omdat het IP gebruikelijk is voor meerdere apparaten.

Probleem met tunnelroute in PIM-register

(Leaf-01) Dit is de eigenlijke FHR: Verzendt Register-berichten naar RP

```
<#root>
```

```
Leaf-01#sh ip pim vrf green tunnel
```

```
Tunnel15*
```

```
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): Deze VTEP (en mogelijk andere) bevat hetzelfde SVI- en IP-adres als de FHR

```
<#root>
```

```
Leaf-03#sh ip pim vrf green tunnel
```

```
Tunnel14
```

```
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): De FHR blijft vastzitten in Register (Het ontvangt geen register-stop van de 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) Dit is de RP: In dit geval bezit het ook dezelfde AnyCast IP als de FHR, en stuurt het daardoor de register-stop naar de verkeerde VTEP kunnen worden gestuurd.

Als RP niet de l2vni maar 2 of 3 andere vteps heeft, zou register-stop naar de verkeerde VTEP kunnen worden gestuurd. De juiste kan selecteren.

```
<#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 op RP toont het probleem waar RP deze route heeft als 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.
```

Probleemoplossing voor PIM-register tunnelroute

De oplossing is om een unieke Loopback IP op alle VTEPs te gebruiken en de configuratie te gebruiken die in deze

```
<#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
```

Gerelateerde informatie

[Configuratiehandleiding EVPN VLAN TRM](#)

[EVN VXLAN Unicast-probleemoplossing](#)

[MVPN Configuration Guide 17.3.x \(Catalyst 9300 Switches\)](#)

[MVPN Configuration Guide 17.3.x \(Catalyst 9500 Switches\)](#)

[BGP-configuratiehandleiding](#)

Over deze vertaling

Cisco heeft dit document vertaald via een combinatie van machine- en menselijke technologie om onze gebruikers wereldwijd ondersteuningscontent te bieden in hun eigen taal. Houd er rekening mee dat zelfs de beste machinevertaling niet net zo nauwkeurig is als die van een professionele vertaler. Cisco Systems, Inc. is niet aansprakelijk voor de nauwkeurigheid van deze vertalingen en raadt aan altijd het oorspronkelijke Engelstalige document ([link](#)) te raadplegen.