Verificar replicación de cabecera en fabric de acceso SD

Contenido

Introducción

Este documento describe cómo resolver problemas de replicación de HeadEnd en el fabric de

acceso SD (SDA).

Prerequisites

Requirements

Cisco recomienda que tenga conocimiento sobre estos temas:

- Reenvío de protocolo de Internet (IP)
- Protocolo de separación Localizador/ID (LISP)
- Modo disperso de multidifusión independiente de protocolo (PIM)

Componentes Utilizados

- C9000v en Cisco IOS® XE 17.10.1
- Cisco Catalyst Center Versión 2.3.5.3

La información que contiene este documento se creó a partir de los dispositivos en un ambiente de laboratorio específico. Todos los dispositivos que se utilizan en este documento se pusieron en funcionamiento con una configuración verificada (predeterminada). Si tiene una red en vivo, asegúrese de entender el posible impacto de cualquier comando.

Este documento también puede utilizarse con estas versiones de software y hardware:

- C9200
- C9300
- C9400
- C9500
- C9600
- Cisco IOS® XE 16.12 y versiones posteriores

Antecedentes

La replicación de cabecera SDA es una forma de multidifusión superpuesta, que se utiliza para transportar el tráfico multidifusión entre los dispositivos de fabric, encapsulando el tráfico multidifusión en un encabezado IP de unidifusión. La replicación de cabecera puede enrutar el tráfico de multidifusión entre los orígenes y los receptores en la misma VLAN o en una VLAN diferente (se puede enrutar la multidifusión de la misma VLAN).

El tráfico multidifusión entre orígenes y receptores en el mismo Fabric Edge no se reenvía mediante multidifusión superpuesta (encapsulación VXLAN), sino que el Fabric Edge lo enruta localmente.

Cualquier forma de multidifusión superpuesta (de cabecera o nativa) no puede enrutar el tráfico de multidifusión para grupos en el rango 224.0.0.0/24, o con TTL=1, esto se maneja a través de la Inundación de Capa 2

Nota: Significa que el lector toma nota. Las notas contienen sugerencias de gran ayuda o referencias a material que no se encuentra en el documento.

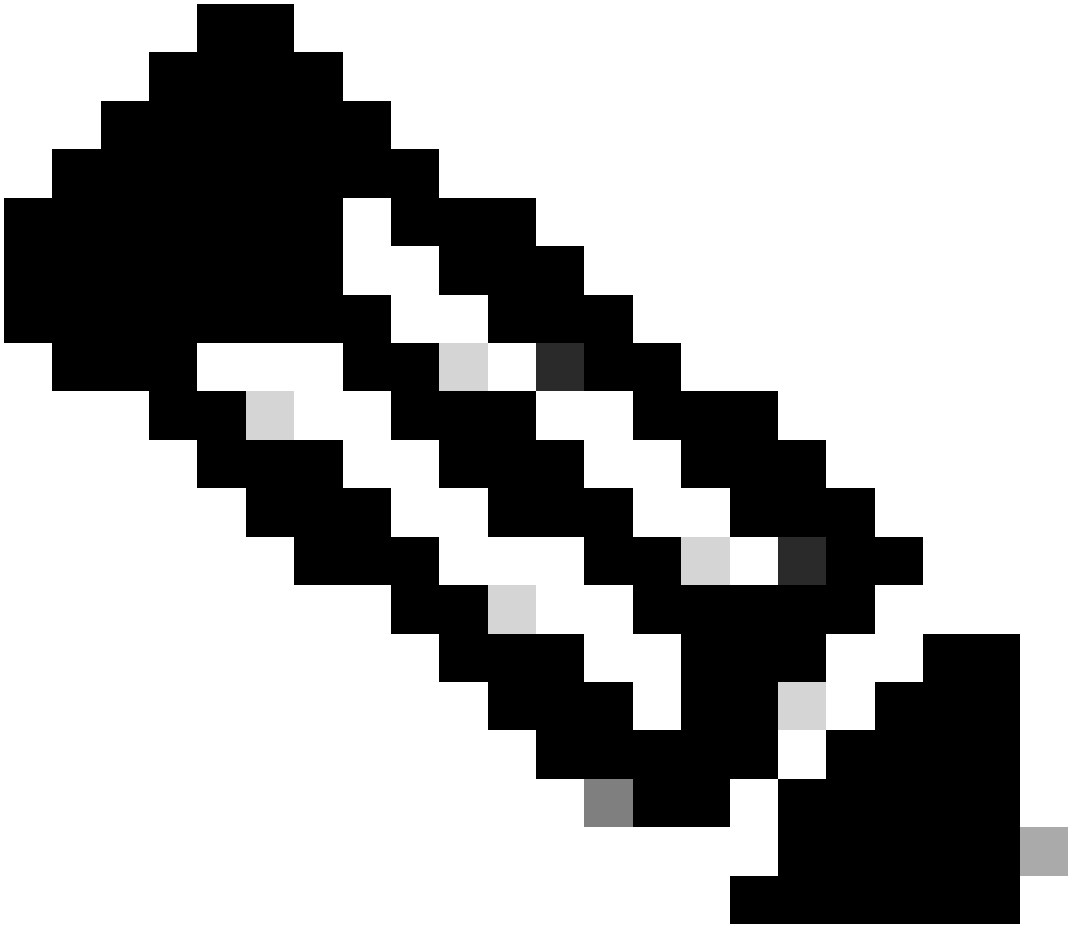

Nota:Los comandos de plataforma (fed) pueden variar. El comando puede ser "show platform fed <active|standby>" versus "show platform fed switch <active|standby>"". Si la sintaxis indicada en los ejemplos no se analiza, pruebe con la variante.

Topología

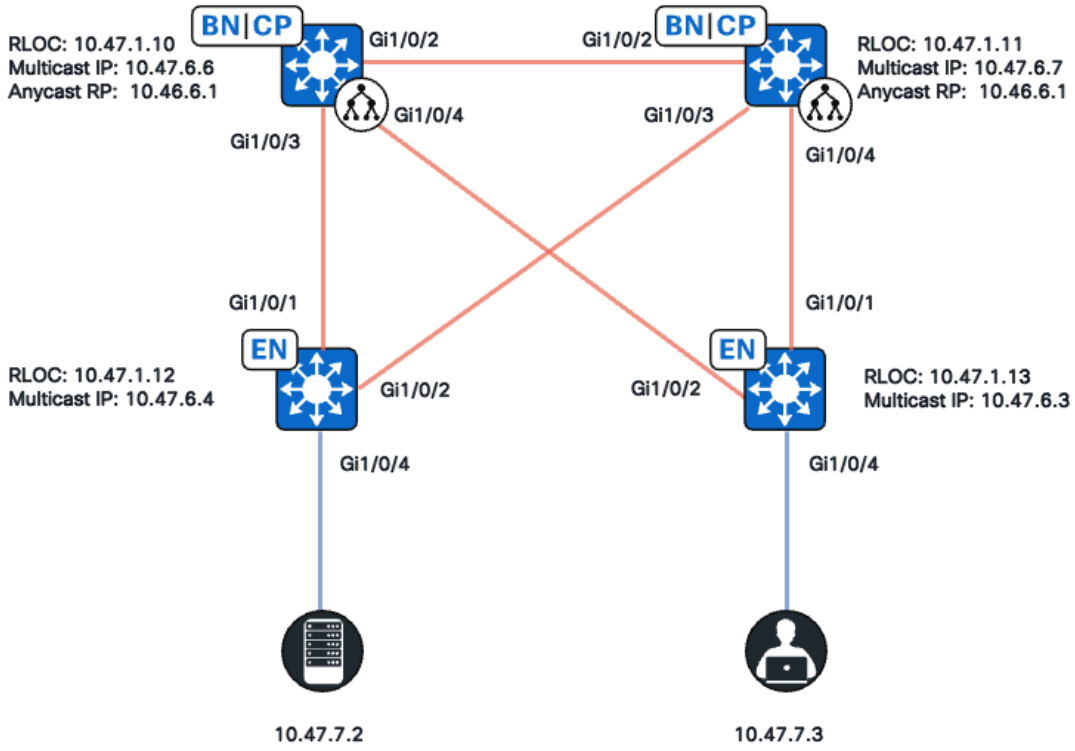

Topología de red

En esta topología:

- 10.47.10 y 10.47.1.11 son fronteras entre ubicaciones que también funcionan como punto de encuentro de difusión ilimitada (RP) con protocolo de transmisión de fuente multidifusión (MSDP) entre los dos en la red virtual (VN)/VRF.
- 10.47.1.12 y 10.47.1.13 son nodos de extremo de fabric
- 10.47.7.2 es el origen de multidifusión
- 10.47.7.3 es el receptor de multidifusión
- 239.1.1.1 es la dirección de destino de grupo (GDA) de multidifusión

Configuración

Se supone que Cisco Catalyst Center se utiliza para aprovisionar el fabric SDA con la configuración predeterminada:

- La implementación de replicación es replicación de cabecera
- RP de difusión ilimitada con MSDP para una multidifusión de multidifusión de cualquier origen (ASM) en los bordes colocados

Después de la correcta configuración desde Catalyst Center, la configuración relevante por dispositivo contiene varias secciones:

Configuración de Fabric Edge (10.47.1.12)

```
ip multicast-routing vrf blue_vn
ip multicast vrf blue_vn multipath
!
interface LISP0.4100
vrf forwarding blue_vn
ip pim sparse-mode
end
!
interface Loopback4100
vrf forwarding blue_vn
ip address 10.47.6.4 255.255.255.255
ip pim sparse-mode
end
!
ip pim vrf blue_vn register-source Loopback4100
ip pim vrf blue_vn rp-address 10.47.6.1 ASM_ACL_IPV4_blue_vn_10.47.6.1
!
interface Vlan1025
description Configured from Cisco DNA-Center
mac-address 0000.0c9f.fb87
vrf forwarding blue_vn
ip address 10.47.7.1 255.255.255.0
ip helper-address 10.47.9.9
no ip redirects
ip pim passive
ip route-cache same-interface
ip igmp version 3
ip igmp explicit-tracking
no lisp mobility liveness test
lisp mobility blue-IPV4
end
!
ip access-list standard ASM_ACL_IPV4_blue_vn_10.47.6.1
10 permit 239.0.0.0 0.255.255.255
```

Configuración de Fabric Edge (10.47.1.13)

```
ip multicast-routing vrf blue_vn
ip multicast vrf blue_vn multipath
!
interface LISP0.4100
vrf forwarding blue_vn
ip pim sparse-mode
end
!
interface Loopback4100
vrf forwarding blue_vn
ip address 10.47.6.4 255.255.255.255
ip pim sparse-mode
end
!
ip pim vrf blue_vn register-source Loopback4100
ip pim vrf blue_vn rp-address 10.47.6.1 ASM_ACL_IPV4_blue_vn_10.47.6.1
!
interface Vlan1025
```

```
description Configured from Cisco DNA-Center
mac-address 0000.0c9f.fb87
vrf forwarding blue_vn
ip address 10.47.7.1 255.255.255.0
ip helper-address 10.47.9.9
no ip redirects
ip pim passive
ip route-cache same-interface
ip igmp version 3
ip igmp explicit-tracking
no lisp mobility liveness test
lisp mobility blue-IPV4
end
!
ip access-list standard ASM_ACL_IPV4_blue_vn_10.47.6.1
10 permit 239.0.0.0 0.255.255.255
```

Configuración de RP de difusión ilimitada/borde compartido (10.47.1.10)

```
router bgp 69420
address-family ipv4 vrf blue_vn
aggregate-address 10.47.6.0 255.255.255.0 summary-only
!
router lisp
site site_uci
eid-record instance-id 4100 10.47.6.0/24 accept-more-specifics
!
ip multicast-routing vrf blue_vn
ip multicast vrf blue_vn multipath
!
interface LISP0.4100
vrf forwarding blue_vn
ip pim sparse-mode
end
!
interface Loopback4100
vrf forwarding blue_vn
ip address 10.47.6.1 255.255.255.255
ip pim sparse-mode
end
!
interface Loopback4600
vrf forwarding blue_vn
ip address 10.47.6.6 255.255.255.255
ip pim sparse-mode
end
!
ip pim vrf blue_vn rp-address 10.47.6.1 ASM_ACL_IPV4_blue_vn_10.47.6.1
ip pim vrf blue_vn register-source Loopback4100
!
ip access-list standard ASM_ACL_IPV4_blue_vn_10.47.6.1
10 permit 239.0.0.0 0.255.255.255
!
ip msdp vrf blue_vn peer 10.47.6.7 connect-source Loopback4600
ip msdp vrf blue_vn cache-sa-state
ip msdp vrf blue_vn originator-id Loopback4600
```

Configuración de RP de difusión ilimitada/borde compartido (10.47.1.11)

```
router bgp 69420
address-family ipv4 vrf blue_vn
aggregate-address 10.47.6.0 255.255.255.0 summary-only
!
router lisp
site site_uci
eid-record instance-id 4100 10.47.6.0/24 accept-more-specifics
!
ip multicast-routing vrf blue_vn
ip multicast vrf blue_vn multipath
!
interface LISP0.4100
vrf forwarding blue_vn
ip pim sparse-mode
end
!
interface Loopback4100
vrf forwarding blue_vn
ip address 10.47.6.1 255.255.255.255
ip pim sparse-mode
end
!
interface Loopback4600
vrf forwarding blue_vn
ip address 10.47.6.7 255.255.255.255
ip pim sparse-mode
end
!
ip pim vrf blue_vn rp-address 10.47.6.1 ASM_ACL_IPV4__blue_vn_10.47.6.1
ip pim vrf blue_vn register-source Loopback4100
!
ip access-list standard ASM_ACL_IPV4_blue_vn_10.47.6.1
10 permit 239.0.0.0 0.255.255.255
!
ip msdp vrf blue_vn peer 10.47.6.6 connect-source Loopback4600
ip msdp vrf blue_vn cache-sa-state
ip msdp vrf blue_vn originator-id Loopback4600
```

Verificación del plano de control

A continuación, verifique el protocolo de pertenencia a grupos de Internet (IGMP) y PIM.

El receptor de multidifusión envía el informe de afiliación IGMP

El receptor de multidifusión (10.47.7.3) envía un informe de pertenencia IGMP (MR) o una unión IGMP para indicar interés en recibir tráfico de multidifusión. Puede configurar una captura de paquetes integrada (EPC) para confirmar que se ha recibido una MR de IGMP:

<#root>

Edge-2#

monitor capture 1 interface GigabitEthernet1/0/5 IN

Edge-2#

monitor capture 1 match any

Edge-2#

monitor capture 1 buffer size 10

Edge-2#

monitor capture 1 start

Edge-2#

monitor capture 1 stop

Edge-1#

show monitor capture 1 buff display-filter igmp brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exi
145 63.730527 10.47.7.4 -> 239.1.1.1 IGMPv2 60 Membership Report group 239.1.1.1
```

A continuación, asegúrese de que el borde del entramado sea el router designado (DR) PIM para la VLAN en la que se encuentra el receptor de multidifusión. Esto también se conoce como el router de último salto (LHR). Puede utilizar el comando "show ip pim vrf <VN Name> interface vlan <vlan> detail" | incluir PIM DR"

<#root>

Edge-2#

show ip pim vrf blue_vn interface vlan 1025 detail | i PIM DR

```
PIM DR: 10.47.7.1 (this system)
```

Valide que IGMP Snooping haya detectado IGMP MR con el comando "show ip igmp vrf <VN Name> snooping group"

<#root>

```
Edge-1#
```

show ip igmp vrf blue_vn snooping groups

```
Vlan Group           Type Version Port List
------------------------------------------------------------------------
1025 239.255.255.254 igmp v2      Gi1/0/5
```

Creación de árbol compartido en modo disperso de PIM

El Edge-2, que es el DR en ese segmento envía una unión PIM (*,G) hacia el RP de difusión ilimitada. Si la dirección RP de difusión ilimitada no se resuelve en la memoria caché de mapas de LISP, el proceso LISP EID Watch es responsable de activar las solicitudes de mapas de LISP. Puede utilizar el comando "show lisp instance-id <LISP L3 ID> ipv4/ipv6 eid-watch | iniciar RLOC"

<#root>

```
Edge-2#
```

show lisp instance-id 4100 ipv4 eid-watch | begin RLOC

```
LISP IPv4 EID Watches for Table (RLOC mapping in vrf blue_vn IPv4) IID (4100), 1 watch entries
Watch entries for prefix 10.47.6.1/32
```

10.47.6.1

,

multicast

```
Edge-2#
```

show lisp instance-id 4100 ipv4 map-cache 10.47.6.1

```
LISP IPv4 Mapping Cache for LISP 0 EID-table vrf blue_vn (IID 4100), 1 entries
10.47.6.1/32, uptime: 9w1d, expires: 20:19:57, via map-reply, complete
Sources: map-reply
State: complete, last modified: 9w1d, map-source: 10.47.1.10
Active, Packets out: 577721(21849998 bytes), counters are not accurate (~ 00:00:12 ago)
Locator Uptime State Pri/Wgt Encap-IID
```

10.47.1.10

```
 9w1d up 10/10 -
Last up-down state change: 1w1d, state change count: 3
Last route reachability change: 9w1d, state change count: 1
Last priority / weight change: never/never
RLOC-probing loc-status algorithm:
Last RLOC-probe sent: 1w1d (rtt 272ms)
```

10.47.1.11

```
 9w1d up 10/10 -
Last up-down state change: 9w1d, state change count: 1
Last route reachability change: 9w1d, state change count: 1
Last priority / weight change: never/never
RLOC-probing loc-status algorithm:
Last RLOC-probe sent: 1w1d (rtt 602ms)

Edge-2#
```

show ip rpf vrf blue_vn 10.47.6.1

```
RPF information for (10.47.6.1)
RPF interface: LISP0.4100
RPF neighbor: ? (10.47.1.10)
RPF route/mask: 10.47.6.1/32
RPF type: unicast ()
Doing distance-preferred lookups across tables
Multicast Multipath enabled.
RPF topology: ipv4 multicast base
```

Valide la entrada (*,G) en el Edge-2 con el comando "show ip mroute vrf <VN Name> <multicast group>"

<#root>

```
Edge-2#
```

show ip mroute vrf blue_vn 239.1.1.1

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 4d05h/00:02:12, RP
```

10.47.6.1

```
, flags: SC
```

<-- Anycast RP IP address 10.47.6.1

```
Incoming interface: LISP0.4100, RPF nbr
10.47.1.10 <-- Reverse Path Forwarding (RPF) neighbor to get to the Anycast RP IP

Outgoing interface list:
Vlan1025
, Forward/Sparse-Dense, 4d05h/00:02:12, flags:
<-- Outgoing interface list (OIL) is populated via PIM Join or IGMP Membership Report
```

Vecinos PIM en la superposición

Una vez que el vecino RPF que está representado por su localizador de ruteo (RLOC) y es accesible a través de la interfaz LISP, se agrega como vecino PIM en el VRF/VN.

Hay varios aspectos que debe tener en cuenta:

- La verificación RPF, que se utiliza para enviar el PIM (*,G) Join desencadena la creación de vecinos PIM con un temporizador de vencimiento de dos minutos. Si no se envían mensajes de incorporación de PIM durante 2 minutos, el vecino agota el tiempo de espera.
- PIM tiene que crear explícitamente una estructura de vecino para el RLOC correspondiente, ya que los mensajes de saludo PIM no se envían en la superposición SDA

<#root>

```
Edge-2#
show ip pim vrf blue_vn neighbor

PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
L - DR Load-balancing Capable
Neighbor    Interface Uptime/Expires Ver DR
Address                                  Prio/Mode
10.47.1.10 LISP0.4100 4d23h/00:01:37 v2  0 /
```

Creación de RP de difusión ilimitada (*,G)

Basándose en la unión PIM (*,G) recibida desde el borde 2, el borde 1 crea (*,G) con el aceite hacia el RLOC del borde 2

<#root>

Border-1#

show ip mroute vrf blue_vn 239.1.1.1

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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 4d23h/00:02:48, RP 10.47.6.1, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:

LISP0.4100

,

10.47.1.13

, Forward/Sparse, 4d23h/00:02:33, flags:

<-- RLOC of Edge-2

Registro de origen de multidifusión

El origen de multidifusión 10.47.7.2 envía tráfico de multidifusión, que ingresa al Edge-1. El Edge-1 dirige el paquete a la CPU para crear el estado (S,G), y el Edge-1 registra el origen en el RP de difusión ilimitada.

<#root>

Edge-1#

monitor capture 1 interface GigabitEthernet1/0/4 IN

Edge-1#

monitor capture 1 match any

Edge-1#

```
monitor capture 1 buffer size 10


Edge-1#
monitor capture 1 start


Edge-1#
monitor capture 1 stop



Edge-1#
show monitor capture 1 buffer brief


Starting the packet display ........ Press Ctrl + Shift + 6 to exit
1 0.000000 10.47.7.2 -> 239.1.1.1 ICMP 98 Echo (ping) request id=0x0007, seq=107/27392, ttl=5
2 0.355071 10.47.7.3 -> 239.1.1.1 ICMP 98 Echo (ping) request id=0x0007, seq=107/27392, ttl=5
3 1.096757 10.47.7.3 -> 239.1.1.1 ICMP 98 Echo (ping) request id=0x0007, seq=108/27648, ttl=5
4 1.102425 10.47.7.3 -> 239.1.1.1 ICMP 98 Echo (ping) request id=0x0007, seq=108/27648, ttl=5
```

Una vez que Border-1 recibe el paquete de multidifusión a través del registro PIM, Border-1 tiene (S,G) y lo anuncia al Border-2 a través de MSDP

<#root>

```
Border-1#
show ip mroute vrf blue_vn 239.1.1.1 10.47.7.2


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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(10.47.7.2, 239.1.1.1), 00:02:26/00:00:34, flags: T
A <-- A flag indicates that this is a candidate for MSDP advertisement


Incoming interface: LISP0.4100, RPF nbr 10.47.1.12
```

```
Outgoing interface list:
LISP0.4100, 10.47.1.13, Forward/Sparse, 00:02:26/00:02:36, flags:
```

Anuncio de origen MSDP

Utilice el comando "show ip msdp vrf <VN name> sa-cache" para ver la caché activa de origen. Puede utilizar el comando "show ip msdp vrf <nombre de VPN> summary" para ver el peer MSDP

<#root>

```
Border-1#
```

show ip msdp vrf blue_vn sa-cache

```
MSDP Source-Active Cache - 1 entries
(10.47.7.2, 239.1.1.1), RP 10.47.6.7, BGP/AS 23456, 00:00:34/00:05:25, Peer 10.47.6.7
```

```
Border-1#
```

show ip msdp vrf blue_vn summary

```
MSDP Peer Status Summary
Peer Address AS    State Uptime/  Reset SA   Peer Name
                       Downtime Count Count
```

10.47.6.7

```
    23456
```

Up

```
    1w1d     0     1
```

Border-2 recibe información (S,G) de Border-1 a través del anuncio MSDP. Si Border-2 ha recibido una PIM (*,G) Join desde Edge-2, Border-2 crea una entrada (S,G) y hereda el LISP OIL del (*,G) que apunta al RLOC del Edge-2. La regla general es que las entradas de SA MSDP sólo se instalan en la base de información de ruteo multicast (MRIB) si existe un (*,G).

<#root>

```
Border-2#
```

show ip msdp vrf blue_vn sa-cache

```
MSDP Source-Active Cache - 1 entries
(10.47.7.2, 239.1.1.1), RP 10.47.6.6, BGP/AS 23456, 00:13:59/00:03:28, Peer 10.47.6.6
```

```
Border-2#
```

show ip mroute vrf blue_vn 239.1.1.1

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 00:21:04/00:00:06, RP 10.47.6.1, flags: SP
Incoming interface: Null, RPF nbr 0.0.0.0
```

Outgoing interface list: Null <-- Indicates no PIM (*,G) Join received, if there was an OIL, ther

Border-1 envía un PIM (S,G) Join hacia la fuente 10.47.7.2 para atraer el tráfico multicast de forma nativa, lo que permite que Edge-1 actualice el OIL (S,G)

<#root>

```
Edge-1#
```

show ip mroute vrf blue_vn 239.1.1.1 10.47.7.3

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
```

```
Interface state: Interface, Next-Hop or VCD, State/Mode
(10.47.7.2, 239.1.1.1), 01:19:57/00:02:29, flags: FT
Incoming interface:
```

Vlan1025

```
, RPF nbr 0.0.0.0
```

<-- Multicast source 10.47.7.2 is in VLAN 1025

```
Outgoing interface list:
LISP0.4100,
```

10.47.1.10

```
, Forward/Sparse, 01:19:55/00:02:30, flags:
```

<-- RLOC of Border-1

El tráfico multidifusión de 10.47.7.2 a 239.1.1.1 se reenvía fuera de 10.47.6.6 (frontera-1) a través de la encapsulación VXLAN de unidifusión. Border-1 desencapsula el tráfico VXLAN y lo vuelve a encapsular en Edge-2 (10.47.1.13)

<#root>

```
Border-1#
```

show ip mroute vrf blue_vn 239.1.1.1

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 5d01h/00:03:14, RP 10.47.6.1, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
LISP0.4100, 10.47.1.13, Forward/Sparse, 5d01h/00:02:54, flags:
```

(

10.47.7.2

```
, 239.1.1.1), 00:02:28/00:00:30, flags: MT
```

<-- Unicast Source

```
Incoming interface: LISP0.4100, RPF nbr
```

10.47.1.12

<-- RPF neighbor to get to the source (Edge-1)

```
Outgoing interface list:
LISP0.4100,
```

10.47.1.13

```
, Forward/Sparse, 00:02:28/00:03:14, flags:
```

<-- RLOC of Edge-2

Transición al árbol de ruta más corta (SPT)

Una vez que el router de último salto (LHR) Edge-2, recibe el paquete multicast a lo largo del árbol (*,G), intenta hacer la transición SPT y enviar una unión PIM (S,G) al Edge-1.

<#root>

```
Edge-2#
```

show ip mroute vrf blue_vn 239.1.1.1

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 4d23h/stopped, RP 10.47.6.1, flags: SJC
Incoming interface: LISP0.4100, RPF nbr 10.47.1.10
Outgoing interface list:
```

Vlan1025

, Forward/Sparse-Dense, 4d23h/00:02:40, flags:

<-- LHR creates the OIL because of receipt of an IGMP MR

(

10.47.7.2

, 239.1.1.1), 00:00:02/00:02:57, flags: JT

<-- Unicast Source

Incoming interface: LISP0.4100, RPF nbr

10.47.1.12

<-- RPF neighbor to get to 10.47.7.2, which is Edge-1 RLOC

Outgoing interface list:

Vlan1025

, Forward/Sparse-Dense, 00:00:02/00:02:57, flags:

<-- Multicast traffic is forwarded into VLAN 1025, where 10.47.7.3 is

El FHR (Edge-1) tiene (S,G) apuntando directamente hacia el RLOC del Edge-2

<#root>

Edge-1#

show ip mroute vrf blue_vn 239.1.1.1

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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 5d01h/stopped, RP 10.47.6.1, flags: SCF
Incoming interface: LISP0.4100, RPF nbr 10.47.1.10

```
Outgoing interface list:
Vlan1025, Forward/Sparse-Dense, 5d01h/00:01:40, flags:

(

10.47.7.2

, 239.1.1.1), 01:53:06/00:02:42, flags: FT

<-- Unicast Source


Incoming interface: Vlan1025, RPF nbr 0.0.0.0
Outgoing interface list:
LISP0.4100,

10.47.1.13

, Forward/Sparse, 00:14:22/00:03:07, flags:

<-- Edge-2's RLOC
```

Verificación del plano de datos (independiente de la plataforma)

Puede haber varios problemas que pueden impedir que el origen o el receptor de multidifusión envíen o reciban el tráfico. Esta sección se centra en la validación de los problemas que pueden afectar tanto al origen de multidifusión como al receptor de multidifusión, haciendo hincapié en los problemas que no están relacionados con la programación de hardware.

Verificación del lado de origen

Para validar el origen de multidifusión y la capacidad de creación de FHR (S,G), valide la función de seguridad integrada en el switch (SISF), LISP, Cisco Express Forwarding (CEF) y, a continuación, RPF.

El origen de multidifusión debe estar en el seguimiento de dispositivos (IPDT) de ISF/IP, que controla el resto de LISP, CEF y, en última instancia, RPF.

Puede utilizar el comando "show device-tracking database address <IP address>" para asegurarse de que el origen de multidifusión tenga una entrada IPDT válida.

<#root>

```
Edge-1#

show device-tracking database address 10.47.7.2


Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol, DH4 - I
Preflevel flags (prlvl):
0001:MAC and LLA match 0002:Orig trunk 0004:Orig access
```

```
0008:Orig trusted trunk 0010:Orig trusted access 0020:DHCP assigned
0040:Cga authenticated 0080:Cert authenticated 0100:Statically assigned
     Network Layer Address Link Layer Address Interface vlan prlvl age  state     Time left
DH4  10.47.7.2             5254.0012.521d     Gi1/0/4   1025 0024  163s REACHABLE 81 s try 0(8428
```

A continuación, asegúrese de que la base de datos LISP del FHR tenga una entrada para el origen de multidifusión. Utilice el comando "show lisp instance-id <LISP L3 ID> ipv4 database ip address/32"

<#root>

```
Edge-1#
```

show lisp instance-id 4100 ipv4 database 10.47.7.2/32

```
LISP ETR IPv4 Mapping Database for LISP 0 EID-table vrf blue_vn (IID 4100), LSBs: 0x1
Entries total 1, no-route 0, inactive 0, do-not-register 1
10.47.7.3/32, dynamic-eid blue-IPV4, inherited from default locator-set rloc_691b1fe4-5264-44c2-b
Uptime: 1w2d, Last-change: 1w2d
Domain-ID: local
Service-Insertion: N/A
Locator Pri/Wgt Source State
10.47.1.13 10/10 cfg-intf site-self, reachable
Map-server Uptime ACK Domain-ID
10.47.1.10 1w2d Yes 0
10.47.1.11 1w2d Yes 0
```

```
Edge-1#
```

show ip lisp instance-id 4100 forwarding eid local 10.47.7.2

```
Prefix
10.47.7.2/32
```

CEF crea una entrada basada en LISP, CEF apunta a una entrada de host /32, no a LISP.

<#root>

```
Edge-1#
```

show ip cef vrf blue_vn 10.47.7.2

```
10.47.7.2/32
nexthop 10.47.7.2 Vlan1025
```

A continuación, RPF se deriva de CEF

<#root>

```
Edge-1#
```

show ip rpf vrf blue_vn 10.47.7.2

```
RPF information for (10.47.7.2)
RPF interface: Vlan1025
RPF neighbor: ? (
```

10.47.7.2

```
) - directly connected
RPF route/mask: 10.47.7.2/32
RPF type:
```

unicast (lisp)

```
Doing distance-preferred lookups across tables
Multicast Multipath enabled.
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

Si no hay una entrada válida en SISF/IPDT, lo que resulta en ninguna asignación de base de datos LISP en el FHR, lo que resulta en CEF y RPF apuntando a las fronteras. Si el origen multicast envía puntos RPF de tráfico a la interfaz incorrecta, lo que resulta en una falla RPF, (S,G) no se forma.

<#root>

```
Edge-1#
```

show device-tracking database address 10.47.7.2

```
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol, DH4 - I
Preflevel flags (prlvl):
0001:MAC and LLA match 0002:Orig trunk 0004:Orig access
0008:Orig trusted trunk 0010:Orig trusted access 0020:DHCP assigned
0040:Cga authenticated 0080:Cert authenticated 0100:Statically assigned
Network Layer Address Link Layer Address Interface vlan prlvl age state Time left
```

```
Edge-1#
```

show lisp instance-id 4100 ipv4 database 10.47.7.2/32

```
% No database-mapping entry for 10.47.7.2/32.
```

```
Edge-1#
```

show ip cef vrf blue_vn 10.47.7.2

```
10.47.7.0/24
nexthop 10.47.1.10
LISP0.4100 <-- Result of a LISP Negative Map-Reply, so the LISP interface is now the RPF interfac

nexthop 10.47.1.11
LISP0.4100 <-- Result of a LISP Negative Map-Reply, so the LISP interface is now the RPF interfac


Edge-1#
show ip rpf vrf blue_vn 10.47.7.2


RPF information for (10.47.7.2)
RPF interface:
LISP0.4100


RPF neighbor: ? (
10.47.1.11
)
RPF route/mask: 10.47.7.2/32
RPF type: unicast ()
Doing distance-preferred lookups across tables
Multicast Multipath enabled.
RPF topology: ipv4 multicast base
```

Para evitar esto, trate el origen de multidifusión como un host silencioso, donde la difusión dirigida IP, la inundación, los enlaces estáticos SISF/IPDT pueden superar este problema.

Registro de origen

El registro PIM es un flujo de paquetes de unidifusión, que utiliza LISP/VXLAN como cualquier otro paquete de unidifusión. Hay varias comprobaciones necesarias para validar que el FHR puede registrar correctamente el origen de multidifusión en el RP de difusión ilimitada.

Primero, asegúrese de que el RP de difusión ilimitada esté configurado correctamente para el GDA.

```
<#root>
Edge-1#
show ip pim vrf blue_vn rp 239.1.1.1
```

```
Group: 239.1.1.1, RP: 10.47.6.1, uptime 5d22h, expires never
```

Asegúrese de que se ha formado el túnel de registro PIM.

```
<#root>

Edge-1#

show ip pim vrf blue_vn tunnel


Tunnel1
Type : PIM Encap
RP :

10.47.6.1 <-- This is from "ip pim vrf blue_vn rp-address 10.47.6.1 ASM_ACL_IPV4_blue_vn_10.47.6.

Source :

10.47.6.4 <-- This is from "ip pim vrf blue_vn register-source Loopback4100"

State : UP
Last event : Created (1w2d)
```

Garantizar la disponibilidad de IP para el RP de difusión ilimitada

```
<#root>

Edge-1#

show ip cef vrf blue_vn 10.47.6.1


10.47.6.1/32
nexthop

10.47.1.10

 LISP0.4100

<-- RLOC of Border-1


nexthop

10.47.1.11

 LISP0.4100

<-- RLOC of Border-2


Edge-1#
```

```
ping vrf blue_vn 10.47.6.1 source lo4100


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.47.6.1, timeout is 2 seconds:
Packet sent with a source address of 10.47.6.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 257/275/294 ms
```

Verificación del lado receptor

- Asegúrese de que el receptor multicast esté enviando un IGMP MR.
- Asegúrese de que la función IGMP Snooping esté habilitada. Las VPN solo de L2 son el único tipo de VPN que no tiene activada la función IGMP Snooping
- Asegúrese de que no haya ninguna ACL de puerto, ACL de VLAN, ACL de puerto ruteado configurada que descarte IGMP MR.
- Valide la versión de IGMP MR; de forma predeterminada, es IGMPv2, si el receptor multicast es IGMPv3, que requiere "ip igmp version 3"
- Asegúrese de que "ip option drop" no esté configurado

Verificación de LHR PIM (*,G)

- Asegúrese de que el LHR sea el DR PIM para la subred/segmento del receptor
- Asegúrese de que no esté configurado "ip multicast group-range"
- Asegúrese de que no haya ninguna ACL de puerto, ACL de VLAN, ACL de puerto ruteado configurada que descarte IGMP MR.
- Asegúrese de que no haya una CPU alta ni una política de plano de control (CoPP) que descarte IGMP MR.

Verificación de árbol compartido PIM de LHR

Asegúrese de que el RP esté configurado para el grupo multicast

<#root>

```
Edge-2#

show ip mroute vrf blue_vn 239.1.1.1


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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 239.1.1.1), 6d01h/stopped,
```

RP 10.47.6.1

```
, flags: SCF
Incoming interface: LISP0.4100, RPF nbr 10.47.1.10
Outgoing interface list:
Vlan1025, Forward/Sparse-Dense, 6d01h/00:01:34, flags:
```

Asegúrese de que RPF al RP de difusión ilimitada sea correcto

<#root>

```
Edge-2#
```

show ip cef vrf blue_vn 10.47.6.1

```
10.47.6.1/32
nexthop 10.47.1.10 LISP0.4100
nexthop 10.47.1.11 LISP0.4100

Edge-2#
```

show ip rpf vrf blue_vn 10.47.6.1

```
RPF information for (10.47.6.1)
RPF interface: LISP0.4100
RPF neighbor: ? (10.47.1.10)
RPF route/mask: 10.47.6.1/32
RPF type: unicast ()
Doing distance-preferred lookups across tables
Multicast Multipath enabled.
RPF topology: ipv4 multicast base
```

Reenvío de MFIB: verificación del lado de origen

Puede utilizar el comando "show ip mfib vrf <VN Name> <multicast group> <unicast source> verbose" para obtener información adicional sobre el reenvío de paquetes

<#root>

Edge-1#

show ip mfib vrf blue_vn 239.1.1.1 10.47.7.2 verbose

Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
ET - Data Rate Exceeds Threshold, K - Keepalive
DDE - Data Driven Event, HW - Hardware Installed
ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client,
e - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
NS - Negate Signalling, SP - Signal Present,
A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
MA - MFIB Accept, A2 - Accept backup,
RA2 - MRIB Accept backup, MA2 - MFIB Accept backup
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
I/O Item Counts: HW Pkt Count/FS Pkt Count/PS Pkt Count Egress Rate in pps
VRF blue_vn
(10.47.7.2,239.1.1.1) Flags: K HW DDE
0x42 OIF-IC count: 0, OIF-A count: 1
SW Forwarding: 0/0/0/0, Other: 272/272/0
HW Forwarding: 7431223059161284608/0/0/0, Other: 0/0/0
Vlan1025 Flags: RA A MA NS
LISP0.4100,

10.47.1.13

Flags: RF F NS

<-- RLOC of Edge-2

CEF: Adjacency with MAC: 450000000000400000011647 70A2F010D0A2F010C000012B500000000084000000010040 0

Pkts: 0/0/0 Rate: 0 pps

Edge-1#

show adjacency lisp0.4100

Protocol Interface Address
IP LISP0.4100 10.47.1.10(23)
IP LISP0.4100 10.47.1.11(27)
IP LISP0.4100

10.47.1.13

(8)

Edge-2#

show adjacency lisp0.4100 10.47.1.13 detail

Protocol Interface Address
IP LISP0.4100

10.47.1.13

```
(8)
0 packets, 0 bytes
epoch 0
sourced in sev-epoch 14
Encap length 50
450000000000400000 1164770A2F010D
0A2F010C000012B50000000008400000
00100400BA25CDF4AD3852540017FE73
0000
L2 destination address byte offset 0
L2 destination address byte length 0
Link-type after encap: ip
LISP
Next chain element:
```

IP adj out of GigabitEthernet1/0/1

```
, addr 10.47.1.6
```

Se puede utilizar un EPC para validar la encapsulación VXLAN del paquete de multidifusión

<#root>

```
Edge-1#monitor capture 1 interface GigabitEthernet1/0/4 IN
Edge-1#monitor capture 1 interface GigabitEthernet1/0/1 OUT
Edge-1#monitor capture 1 match any
Edge-1#monitor capture 1 buffer size 10
Edge-1#monitor capture 1 limit pps 1000
Edge-1#monitor capture 1 start
Edge-1#monitor capture 1 stop

Edge-1#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
1 0.000000 10.47.7.2 -> 239.1.1.1 ICMP 98 Echo (ping) request id=0x0008, seq=28213/13678,
```

ttl=5 <-- Packet as it ingresses the FHR, TTL is 5

```
2 0.014254 10.47.7.2 -> 239.1.1.1 ICMP 148 Echo (ping) request id=0x0008, seq=28213/13678,
```

ttl=4 <-- Packet as it leaves the FHR, TTL is 4 as is it decremented

Reenvío de MFIB: verificación del lado del receptor

La red subyacente enruta este paquete desde el Edge-1 al Edge-2 mediante el ruteo unicast.

<#root>

```
Edge-2#
```

show ip mroute vrf blue_vn 239.1.1.1 10.47.7.2

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(
```

10.47.7.2

```
,
```

239.1.1.1

```
), 00:01:39/00:01:20, flags: JT
Incoming interface: LISP0.4100, RPF nbr
```

10.47.1.12

```
Outgoing interface list:
```

Vlan1025

```
, Forward/Sparse-Dense, 00:01:39/00:02:45, flags:
```

Con el comando "show ip mfib vrf <VN Name> <group address> <unicast source> counters" puede asegurarse de que los contadores de reenvío de hardware aumenten

<#root>

```
Edge-2#
```

show ip mfib vrf blue_vn 239.1.1.1 counters

```
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)
VRF blue_vn
```

```
12 routes, 7 (*,G)s, 4 (*,G/m)s
Group: 239.1.1.1
RP-tree,
SW Forwarding: 0/0/0/0, Other: 0/0/0
HW Forwarding: 0/0/2/0, Other: 0/0/0
Source: 10.47.7.2,
SW Forwarding: 0/0/0/0, Other: 2/1/1
HW Forwarding:
```

6118996613340856320

```
/0/0/0, Other: 0/0/0
Totals - Source count: 1, Packet count:
```

6118996613340856320

```
Edge-2#
```

show ip igmp snooping groups vlan 1025 239.1.1.1

```
Vlan Group     Type Version Port List
------------------------------------------------------------------------
1025 239.1.1.1 igmp v2      Gi1/0/4
```

Puede utilizar los contadores de multidifusión de salida para validar si el tráfico de multidifusión ha salido o no del LHR, hacia el receptor de multidifusión. Utilice el comando "show controllers ethernet-controller <interface> | incluir multidifusión|transmitir"

<#root>

```
Edge-2#
```

show controllers ethernet-controller g1/0/4 | include Multicast|Transmit
Transmit

```
                  GigabitEthernet1/0/5          Receive

426729240 Total bytes                                    100803109 Total bytes

5732 Unicast frames                                      949355 Unicast frames

5732 Unicast bytes                                       93563018 Unicast bytes
```

4388433

```
 Multicast frames                                 32346 Multicast frames
```

4388433

```
 Multicast bytes                                    7236178 Multicast bytes
<snip>
Edge-2#
```

show controllers ethernet-controller g1/0/5 | include |Multicast|Transmit

Transmit

```
                  GigabitEthernet1/0/5          Receive
426742895 Total bytes                                   100813570 Total bytes
5733 Unicast frames                                     949456 Unicast frames
5733 Unicast bytes                                      93573016 Unicast bytes
```

4388569

```
 Multicast frames                                    32348 Multicast frames
```

4388569

```
 Multicast bytes                                    7236641 Multicast bytes
```

Otra manera de validar el tráfico multicast que sale del LHR es hacer un EPC hacia el receptor multicast.

<#root>

```
Edge-2#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
1 0.168401 10.47.7.2 -> 239.1.1.1 ICMP 106 Echo (ping) request id=0x0008, seq=35903/16268, ttl=3
2 0.969138 10.47.7.2 -> 239.1.1.1 ICMP 106 Echo (ping) request id=0x0008, seq=35904/16524, ttl=3
```

Verificación del plano de datos (según la plataforma)

Creación (S,G): Ruta de punteo de CPU

Para que FHR cree el estado (S,G), un par de los paquetes multicast enviados desde el origen multicast se envían a la CPU para ser procesados por la MFIB. Los paquetes multicast se envían a la cola FED "CPU_Q_MCAST_DATA"

<#root>

```
Edge-1#
```

show platform software fed switch active punt cpuq 30

```
Punt CPU Q Statistics
==========================================
CPU Q Id : 30
CPU Q Name : CPU_Q_MCAST_DATA

Packets received from ASIC : 27124


Send to IOSd total attempts : 27124


Send to IOSd failed count : 0
RX suspend count : 0
RX unsuspend count : 0
RX unsuspend send count : 0
RX unsuspend send failed count : 0
RX consumed count : 0
RX dropped count : 0
RX non-active dropped count : 0
RX conversion failure dropped : 0
RX INTACK count : 0
RX packets dq'd after intack : 0
Active RxQ event : 0
RX spurious interrupt : 0
RX phy_idb fetch failed: 0
RX table_id fetch failed: 0
RX invalid punt cause: 0

Replenish Stats for all rxq:
-----------------------------------------
Number of replenish : 0
Number of replenish suspend : 0
Number of replenish un-suspend : 0
-----------------------------------------
```

Además, la cola CoPP para datos MCAST no debe tener ninguna pérdida. Utilice el comando "show platform hardware fed active qos queue stats internal cpu policer | incluir datos MCAST|QId"

<#root>

```
Edge-1#

show platform hardware fed active qos queue stats internal cpu policer | include MCAST Data|QId


QId PlcIdx                                    Queue  Name               Enabled   Rate     Rate
30   9                                        MCAST  Data                 No      500      400
```

Si el tráfico proviene de una fuente conectada directamente, es procesado por la cola de Linux Shared Memory Punt Interface (LSMPI) para "fuente conectada directamente" si es de una unión (S,G), es "Señalización Mcast PIM"

Utilice el comando "show platform software infrastructure lsmpi punt | incluir causa|multidifusión"

<#root>

```
Edge-1#
show platform software infrastructure lsmpi punt | include Cause|Mcast

Cause                           Total        Total      Length        Dot1q encap        Other
Mcast Directly Connected Source
 0
27038
      0             0                   0           0
Mcast IPv4 Options data packet  0           0          0             0                  0
Mcast Internal Copy             0           0          0             0                  0
Mcast IGMP Unroutable           0           0          0             0                  0
Mcast PIM signaling
 0            0          0            0                  0           0
Mcast punt to RP                0           0          0             0                  0
Mcast UDLR                      0           0          0             0                  0
```

A continuación, se puede realizar una captura de paquetes FED Punject para ver los paquetes multicast del origen y el grupo en la CPU, lo que confirma la interfaz entrante y la cola de la CPU.

<#root>

```
Edge-1#
debug platform software fed switch active punt packet-capture set-filter "ip.addr==239.1.1.1"
Edge-1#
debug platform software fed switch active punt packet-capture start

Edge-1#
debug platform software fed switch active punt packet-capture stop

Punt packet capturing stopped. Captured 2 packet(s)
Edge-1#
show platform software fed switch active punt packet-capture brief

Punt packet capturing: disabled. Buffer wrapping: disabled
```

```
Total captured so far: 2 packets. Capture capacity : 4096 packets
Capture filter : "ip.addr==239.255.255.254"
------ Punt Packet Number: 1, Timestamp: 2024/08/26 15:38:27.341 ------
interface : physical:
```

GigabitEthernet1/0/4

```
[if-id: 0x0000000c], pal:
```

Vlan1025

```
 [if-id: 0x0000001d]
metadata : cause: 12 [
```

Mcast Directly Connected Source

```
], sub-cause: 0, q-no: 30, linktype: MCP_LINK_TYPE_IP [1]
ether hdr : dest mac: 0100.5e7f.fffe, src mac: 5254.0012.521d
ether hdr : ethertype: 0x0800 (IPv4)
ipv4 hdr : dest ip:
```

239.1.1.1,

```
 src ip: 10.47.7.2
ipv4 hdr : packet len: 84, ttl: 5, protocol: 1 (ICMP)
icmp hdr : icmp type: 8, code: 0
```

Programación de Hardware de Mroute - IOS Mroute

La programación de hardware de (S,G) utiliza la misma estructura que cualquier otra ruta de programación: IOS a FMAN RP a FMAN FP, a FED.

<#root>

```
Edge-1#
```

show ip mroute vrf blue_vn 239.1.1.1

```
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, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
t - LISP transit group
Timers: Uptime/Expires
```

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

(*, 239.255.255.254), 00:08:29/stopped, RP
10.47.6.1
, flags: SCF
<-- Anycast RP address

Incoming interface: LISP0.4100, RPF nbr
10.47.1.10 <-- RLOC of Border-1

Outgoing interface list:
Vlan1025, Forward/Sparse-Dense, 00:08:29/00:00:30, flags:

(
10.47.7.2
,
239.1.1.1
), 00:08:28/00:02:54, flags: FT
<-- Unicast source

Incoming interface:
Vlan1025
, RPF nbr 0.0.0.0
<-- Multicast source is in VLAN 1025

Outgoing interface list:
LISP0.4100
,
10.47.1.13
, Forward/Sparse, 00:08:23/00:03:07, flags:
<-- Forwarding to Edge-2
```

Programación de hardware Mroute - IOS MFIB

A continuación, las rutas de multidifusión se agregan a la Base de información de reenvío de multidifusión (MFIB), que es similar a cómo se agrega la Base de información de routing (RIB) a Cisco Express Forwarding (CEF); el equivalente de multidifusión es la MFIB.

```
Edge-1#
```

show ip mfib vrf blue_vn 239.1.1.1 10.47.7.2 verbose

```
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
ET - Data Rate Exceeds Threshold, K - Keepalive
DDE - Data Driven Event, HW - Hardware Installed
ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client,
e - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
NS - Negate Signalling, SP - Signal Present,
A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
MA - MFIB Accept, A2 - Accept backup,
RA2 - MRIB Accept backup, MA2 - MFIB Accept backup
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
I/O Item Counts: HW Pkt Count/FS Pkt Count/PS Pkt Count Egress Rate in pps
VRF blue_vn
(
```

10.47.7.2,239.1.1.1

```
) Flags: K HW DDE
```

<-- Multicast source and GDA

```
0x21 OIF-IC count: 0, OIF-A count: 1
SW Forwarding: 0/0/0/0, Other: 2/2/0
HW Forwarding: 0/0/0/0, Other: 0/0/0
Vlan1025 Flags: RA A MA NS
```

LISP0.4100, 10.47.1.13

```
 Flags: RF F NS
```

<-- RLOC of Edge-2 and the RPF interface to reach 10.47.1.13

```
CEF: Adjacency with MAC: 4500000000004000001164770A2F010D0A2F010C000012B5000000000840000000100400
Pkts: 0/0/0 Rate: 0 pps
```

Programación de hardware Mroute - RP MFIB

Utilice el comando "show platform software ip switch active r0 mfib vrf index <VRF index> group <GDA/32>"

<#root>

```
Edge-1#
```

show vrf detail blue_vn | inc Id

```
VRF blue_vn (
```

VRF Id = 2

```
); default RD <not set>; default VPNID <not set>

Edge-1#
```

show platform software ip switch active r0 mfib vrf index 2 group 239.1.1.1/32

```
Route flags:
S - Signal; C - Directly connected;
IA - Inherit A Flag; L - Local;
BR - Bidir route
*,
```

239.1.1.1/32

```
 --> OBJ_INTF_LIST (0x6b)
Obj id:
```

0x6b

```
, Flags: C
OM handle: 0x34803c47f0

Edge-2#
```

show platform software ip switch active r0 mfib vrf index 2 group address 239.1.1.1 10.47.7.2

```
Route flags:
S - Signal; C - Directly connected;
IA - Inherit A Flag; L - Local;
BR - Bidir route
```

239.1.1.1, 10.47.7.2/64

```
 --> OBJ_INTF_LIST (0x21)
Obj id:
```

0x21

```
, Flags: unknown
OM handle: 0x34803c4088
```

Programación de hardware Mroute - FP MFIB

La entrada FMAN RP para las mismas rutas multicast incluye un ID de Asynchronous Object Manager (AOM), este ID de AOM se utiliza para validar la programación adicional.

Utilice el comando "show platform software ip switch active f0 mfib vrf index <VRF Index> group <GDA/32>"

<#root>

```
Edge-1#
```

show platform software ip switch active f0 mfib vrf index 2 group 239.1.1.1/32

```
Route flags:
```

```
S - Signal; C - Directly connected;
IA - Inherit A Flag; L - Local;
BR - Bidir route
*,
```

239.1.1.1/32

```
 --> OBJ_INTF_LIST (0x6b)
Obj id:
```

0x6b

```
, Flags: C
aom id:
```

29154

```
, HW handle: (nil) (created)

Edge-1#
```

show platform software ip switch active f0 mfib vrf index 2 group address 239.1.1.1 10.47.7.2

```
Route flags:
S - Signal; C - Directly connected;
IA - Inherit A Flag; L - Local;
BR - Bidir route
```

239.1.1.1., 10.47.7.2/64

```
 --> OBJ_INTF_LIST (0x21)
Obj id:
```

0x21

```
, Flags: unknown
aom id:
```

36933

```
, HW handle: (nil) (created)
```

Programación De Hardware Mroute: Objetos Mroute

Con los ID. de AOM, verifique el objeto y los objetos padre para (*,G) y (S,G) utilizando los comandos del gestor de objetos. Puede utilizar el comando "show platform software object-manager switch active f0 object <ID de AOM>" o "show platform software object-manager switch active f0 object <ID de AOM> parent"

Cada ruta multicast tiene dos objetos primarios. Uno de los objetos hace referencia a la tabla ipv4_mcast y el otro es una mlist, que se utiliza en los comandos siguientes.

<#root>

```
Edge-1#
```

show platform software object-manager switch active f0 object 29154

Object identifier: 29154
Description:

PREFIX 0.0.0.0 , 239.1.1.1/32

(Table id 2)
Obj type id: 72
Obj type:

mroute-pfx

Status:

Done

, Epoch: 0, Client data: 0xa3e23c48

Edge-1#

show platform software object-manager switch active f0 object 29154 parents

Object identifier: 26509
Description:

ipv4_mcast table 2 (blue_vn

), vrf id 2
Status: Done
Object identifier: 29153
Description:

mlist 107

Status:

Done

Edge-1#

show platform software object-manager switch active f0 object 36933

Object identifier: 36933
Description:

PREFIX 10.47.7.2 , 239.1.1.164

(Table id 2)
Obj type id: 72
Obj type:

mroute-pfx

Status:

Done

, Epoch: 0, Client data: 0xa413c928

Edge-1#

show platform software object-manager switch active f0 object 36933 parents

```
Object identifier: 26509
Description: ipv4_mcast table 2 (blue_vn), vrf id 2
Status:
```

Done

```
Object identifier: 47695
Description:
```

mlist 33

```
Status:
```

Done

Programación De Hardware Mroute - Objetos Mlist

Los objetos MLIST son una combinación de interfaces entrantes y listas de interfaces salientes. Puede utilizar el comando "show platform software mlist switch active f0 index <index>"

<#root>

This is for (*,G)

```
Edge-1#
```

show platform software mlist switch active f0 index 107

```
Multicast List entries
OCE Flags:
NS - Negate Signalling; IC - Internal copy;
A - Accept; F - Forward;
OCE         Type                   OCE Flags   Interface
-------------------------------------------------------------------------------
0xf8000171  OBJ_ADJACENCY
```

A

```
         LISP0.4100
```

<-- A Flag indicates an Incoming interface for (*,G)

```
0xf80001d1  OBJ_ADJACENCY          NS,
```

F

```
       Vlan1025
```

<-- F Flag indicates an Outgoing interface for (*,G)

```
This is for (S,G)


Edge-1#
show platform software mlist switch active f0 index 33


Multicast List entries
OCE Flags:
NS - Negate Signalling; IC - Internal copy;
A - Accept; F - Forward;
OCE         Type                  OCE Flags      Interface
------------------------------------------------------------------------------------
0x5c        OBJ_ADJACENCY         NS,
F
          LISP0.4100
<-- F Flag indicates an Outgoing interface(s), for (S,G)


0xf80001d1  OBJ_ADJACENCY
A
              Vlan1025
<-- A Flag indicates an Incoming interface, for (S,G)
```

Programación de hardware Mroute - FED Mroute

Para validar la programación FED, utilice el comando "show platform software fed switch active ip mfib vrf <VN Name> <GDA> <unicast source>"

<#root>

```
Edge-1#
show platform software fed switch active ip mfib vrf blue_vn 239.1.1.1 10.47.7.2


Multicast (S,G) Information
VRF : 2
Source Address : 10.47.7.2
HTM Handler : 0x7f45d98c7728
SI Handler : 0x7f45d9a44a28
DI Handler : 0x7f45d9bcb2d8
REP RI handler : 0x7f45d97e7188


Flags :
Packet count : 0
State : 4
RPF :
Vlan1025 A
```

```
OIF :
Vlan1025 A
LISP0.4100 F NS
(Adj: 0x5c )
```

El índice de reescritura proporciona información sobre la encapsulación del tráfico multicast, que es lo que aprovecha la replicación de cabecera. Puede utilizar el comando "show platform hardware fed switch active fwd-asic abstraction print-resource-handle <REP RI Handle> 1"

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f45d97e718

```
Handle:0x7f45d97e7188 Res-Type:ASIC_RSC_RI_REP Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_
priv_ri/priv_si Handle: (nil)Hardware Indices/Handles: index0:0x19 mtu_index/l3u_ri_index0:0x0 ir
Cookie length: 56
00 00 00 00 00 00 00 00 02 00 00 00 03 07 2f 0a fe ff ff ef 00 00 00 00 00 00 00 00 00 00 00 00 0
Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
Replication list RI handle = 7f45d97e7188
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
ASIC [0] Replication Expansion Handle [0x7f45d9b9c048]
Replication list :
Number of RIs = 6
Start RI = 25
Common rewrite = No
Replication REP_RI 0x19 [elements = 1]
[0] ri_list[0]=4 RI_MCAST_BRIDGE_V6 port=88 ri_ref_count:1 dirty=0
RIL first:4 last:4 start:4 ril_total:4 ref_count:0
RI list this:4 num_pairs:4 free:3 next:0 prev:0 ---->
uri1:
```

50

```
 ri_ref_count_1:1 uri0:
```

26

```
 ri_ref_count_0:1 ptr_type:0 last:1 dirty:1
uri1:
```

49151

```
 ri_ref_count_1:0 uri0:49151 ri_ref_count_0:1 ptr_type:1 last:1 dirty:1
uri1:49151 ri_ref_count_1:0 uri0:49151 ri_ref_count_0:0 ptr_type:1 last:1 dirty:0
uri1:49151 ri_ref_count_1:0 uri0:49151 ri_ref_count_0:0 ptr_type:1 last:1 dirty:0
<snip>
```

A continuación, tome los valores URI para validar el intervalo de índice de reescritura.

Utilice el comando "show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <URI> <URI>"

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 50 50

```
ASIC#:0
```

RI:50

```
 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_PAYLOAD(
Dst Mac: MAC Addr: ba:25:cd:f4:ad:38,
```

Src IP: 10.47.1.12 <-- RLOC of Edge-1

Dst IP: 10.47.1.13 <--

RLOC of Edge-2

```
IPv4 TTL: 0
LISP INSTANCEID: 0
L3IF LE Index: 49
ASIC#:1
```

RI:50

```
 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_PAYLOAD(
Dst Mac: MAC Addr: ba:25:cd:f4:ad:38,
```

Src IP: 10.47.1.12 <-- RLOC of Edge-1

Dst IP: 10.47.1.13 <-- RLOC of Edge-2

```
IPv4 TTL: 0
LISP INSTANCEID: 0
L3IF LE Index: 49
```

A continuación, tome el RI del comando anterior para una verificación adicional. Utilice el comando "show platform software fed switch active ip mfib vrf <VN Name> <GDA> <source>"

<#root>

```
Edge-1#
```

show platform software fed switch active ip mfib vrf blue_vn 239.1.1.1 10.47.7.2

```
Multicast (S,G) Information
VRF : 2
Source Address : 10.47.7.2
HTM Handler : 0x7f45d98c7728
SI Handler : 0x7f45d9a44a28
DI Handler : 0x7f45d9bcb2d8
REP RI handler : 0x7f45d97e7188
Flags :
Packet count : 0
State : 4
RPF :
Vlan1025 A
OIF :
Vlan1025 A
LISP0.4100 F NS
```

(Adj: 0x5c)

Utilice el comando "show platform software fed switch active ip adj | include <destination RLOC>"

<#root>

```
Edge-1#
```

show platform software fed switch active ip adj 10.47.1.12

```
IPV4 Adj entries
dest       if_name    dst_mac        si_hdl         ri_hdl          pd_flags adj_id Last-modified
----       -------    -------        ------         ------          -------- ------ -------------
10.47.1.12 LISP0.4100 4500.0000.0000 0x7f45d9a4a5e8 0x7f45d9a4a798  0x60
```

0x5c

2024/08/21 16:18:58.948

<-- 0x5c matches the Adj in the previous command

En el LHR, puede validar el índice de destino para ver a dónde se reenvía el paquete multicast, que son los receptores multicast. Puede utilizar el comando "show platform software fed switch active ip mfib vrf <VN Name> <GDA> <source>"

<#root>

```
Edge-2#
```

show platform software fed switch active ip mfib vrf blue_vn 239.1.1.1 10.47.7.2

```
Multicast (S,G) Information
VRF : 2
Source Address : 10.47.7.2
HTM Handler : 0x7f0efdad33a8
```

```
SI Handler : 0x7f0efdad2648
```

DI Handler : 0x7f0efdad7668

```
REP RI handler : 0x7f0efdad4858
Flags :
Packet count : 0
State : 4
RPF :
LISP0.4100 A
OIF :
Vlan1025 F NS
LISP0.4100 A
(Adj: 0xf8000171 )
```

Tome el controlador DI y utilice el comando "show platform hardware fed switch active fwd-asic abstraction print-resource-handle <identificador DI> 1"

<#root>

```
Edge-2#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f0efdad766

```
Handle:0x7f0efdad7668 Res-Type:ASIC_RSC_DI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_M
priv_ri/priv_si Handle: (nil)Hardware Indices/Handles: index0:0x527c mtu_index/l3u_ri_index0:0x0
Cookie length: 56
00 00 00 00 00 00 00 00 02 00 00 00 03 07 2f 0a fe ff ff ef 00 00 00 00 00 00 00 00 00 00 00 00 0
Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
Destination index = 0x527c
```

pmap = 0x00000000 0x00000010

pmap_intf : [GigabitEthernet1/0/4]

```
cmi = 0x0
rcp_pmap = 0x0
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0
Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
Destination index = 0x527c
pmap = 0x00000000 0x00000000
```

```
cmi = 0x0
rcp_pmap = 0x0
al_rsc_cmi
CPU Map Index (CMI) [0]
ctiLo0 = 0
ctiLo1 = 0
ctiLo2 = 0
cpuQNum0 = 0
cpuQNum1 = 0
cpuQNum2 = 0
npuIndex = 0
stripSeg = 0
copySeg = 0
================================================================
```

Acerca de esta traducción
Cisco ha traducido este documento combinando la traducción automática y los recursos humanos a fin de ofrecer a nuestros usuarios en todo el mundo contenido en su propio idioma.

Tenga en cuenta que incluso la mejor traducción automática podría no ser tan precisa como la proporcionada por un traductor profesional.

Cisco Systems, Inc. no asume ninguna responsabilidad por la precisión de estas traducciones y recomienda remitirse siempre al documento original escrito en inglés (insertar vínculo URL).