

Verifica della replica headend in SD-Access Fabric

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Premesse](#)

[Topologia](#)

[Configurazione](#)

[Verifica del Control Plane](#)

[Il ricevitore multicast invia il rapporto di appartenenza IGMP](#)

[Creazione di strutture condivise in modalità sparsa PIM](#)

[Vicini PIM nella sovrapposizione](#)

[Anycast RP crea \(*,G\)](#)

[Registrazione origine multicast](#)

[Annuncio origine MSDP](#)

[Copertina dell'albero del percorso più breve](#)

[Verifica del piano dati \(indipendente dalla piattaforma\)](#)

[Verifica lato origine](#)

[Registrazione origine](#)

[Verifica lato ricevitore](#)

[Verifica LHR PIM \(*,G\)](#)

[Verifica albero condiviso PIM LHR](#)

[Inoltro MFIB - Verifica lato origine](#)

[Inoltro MFIB - Verifica lato ricevitore](#)

[Verifica del piano dati \(dipendente dalla piattaforma\)](#)

[\(S,G\) Creazione - Percorso punt CPU](#)

[Programmazione hardware Mroute - IOS Mroute](#)

[Programmazione hardware Mroute - MFIB IOS](#)

[Programmazione hardware Mroute - RP MFIB](#)

[Programmazione hardware Mroute - FP MFIB](#)

[Programmazione hardware Mroute - Oggetti Mroute](#)

[Programmazione Mroute Hardware - Oggetti Mlist](#)

[Mroute Programmazione hardware - FED Mroute](#)

Introduzione

In questo documento viene descritto come risolvere i problemi relativi alla replica headend nella

struttura SD-Access (SDA).

Prerequisiti

Requisiti

Cisco raccomanda la conoscenza dei seguenti argomenti:

- Inoltro IP (Internet Protocol)
- Locator/ID Separation Protocol (LISP)
- PIM (Protocol Independent Multicast) in modalità sparse

Componenti usati

- C9000v su Cisco IOS® XE 17.10.1
- Cisco Catalyst Center versione 2.3.5.3

Le informazioni discusse in questo documento fanno riferimento a dispositivi usati in uno specifico ambiente di emulazione. Su tutti i dispositivi menzionati nel documento la configurazione è stata ripristinata ai valori predefiniti. Se la rete è operativa, valutare attentamente eventuali conseguenze derivanti dall'uso dei comandi.

Il presente documento può essere utilizzato anche per le seguenti versioni hardware e software:

- C9200
- C9300
- C9400
- C9500
- C9600
- Cisco IOS® XE 16.12 e versioni successive


Premesse

La replica headend SDA è una forma di multicast di overlay, utilizzata per trasportare il traffico multicast tra i dispositivi fabric, incapsulando il traffico multicast in un'intestazione IP unicast. La replica headend può indirizzare il traffico multicast tra le origini e i ricevitori nella stessa VLAN o in una VLAN diversa (è possibile indirizzare il multicast della stessa VLAN).

Il traffico multicast tra le origini e i ricevitori sullo stesso perimetro del fabric non viene inoltrato utilizzando l'overlay multicast (incapsulamento VXLAN), ma viene instradato localmente dal perimetro del fabric.

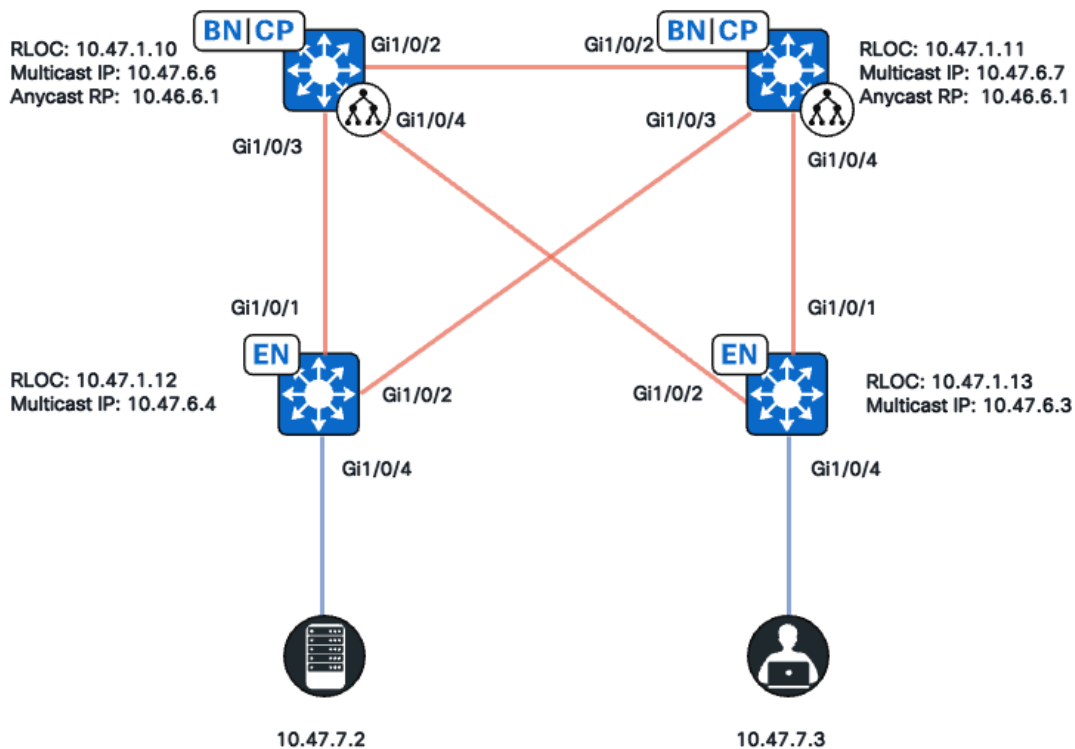
Qualsiasi forma di multicast di overlay (headend o nativo) non può instradare il traffico multicast per gruppi nell'intervallo 224.0.0.0/24 o, con TTL=1, viene gestita tramite Flooding di layer 2

Nota: indica che il lettore prende nota. Le note contengono utili suggerimenti o riferimenti a materiale non trattato nel documento.



Nota: i comandi della piattaforma (feed) possono variare. Il comando può essere "show platform fed <active|standby>" oppure "show platform fed switch <active|standby>". Se la sintassi indicata negli esempi non viene analizzata, provare a utilizzare la variante.

Topologia



Topologia della rete

In questa topologia:

- Le versioni 10.47.10 e 10.47.1.11 sono bordi ripartiti che fungono anche da punto di ripristino (RP) di Anycast con protocollo MSDP (Multicast Source Discovery Protocol) tra i due sulla rete virtuale (VN)/VRF.
- 10.47.1.12 e 10.47.1.13 sono nodi Fabric Edge
- 10.47.7.2 è la fonte multicast
- 10.47.7.3 è il ricevitore multicast
- 239.1.1.1 è l'indirizzo di destinazione del gruppo multicast (GDA)

Configurazione

Si presume che Cisco Catalyst Center venga utilizzato per effettuare il provisioning del fabric SDA con le impostazioni predefinite:

- L'implementazione della replica è la replica headend
- Anycast RP con MSDP per un multicast Any Source Multicast (ASM) sui bordi riposizionati

Dopo aver completato la configurazione dal Catalyst Center, la configurazione pertinente per ciascun dispositivo contiene diverse sezioni:

Configurazione di 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
```

Configurazione di 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
```

Configurazione Border/Anycast RP (10.47.1.10) collocata

```
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
```

Configurazione Border/Anycast RP (10.47.1.11) collocata

```
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
```

Verifica del Control Plane

Verificare quindi che i protocolli IGMP (Internet Group Membership Protocol) e PIM siano corretti.

Il ricevitore multicast invia il rapporto di appartenenza IGMP

Il ricevitore multicast (10.47.7.3) invia un rapporto di appartenenza IGMP (MR) o un join IGMP per indicare l'interesse a ricevere il traffico multicast. È possibile configurare un'acquisizione pacchetto incorporata (EPC) per confermare che è stato ricevuto un MR

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 exit  
145 63.730527 10.47.7.4 -> 239.1.1.1 IGMPv2 60 Membership Report group 239.1.1.1
```

Quindi, verificare che Fabric Edge sia il PIM Designated Router (DR) per la VLAN su cui si trova il ricevitore multicast. Questo router è anche noto come LHR (Last Hop Router). È possibile usare il comando "show ip pim vrf <nome VN> interface vlan <vlan> detail" | includere 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)
```

Verificare che lo snooping IGMP abbia selezionato IGMP MR con il 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
```

Creazione di strutture condivise in modalità sparsa PIM

Edge-2, che è il DR su quel segmento, invia un (*,G) PIM Join all'Anycast RP. Se l'indirizzo RP dell'Anycast non viene risolto nella cache delle mappe LISP, il processo LISP EID Watch è responsabile dell'attivazione delle richieste delle mappe LISP. È possibile utilizzare il comando "show lisp instance-id <LISP L3 ID> ipv4/ipv6 eid-watch | begin 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

Convalidare la voce (*,G) su Edge-2 con il comando "show ip mroute vrf <nome VN>
<gruppo 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), 4d05h/00:02:12, RP

10.47.6.1

, flags: SC

<-- Anycast RP IP address 10.47.6.1

Incoming interface: LISPO.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

Vicini PIM nella sovrapposizione

Una volta che il router adiacente RPF è rappresentato dal relativo RLOC (Routing Locator) ed è raggiungibile tramite l'interfaccia LISP, viene aggiunto come router adiacente PIM nel VRF/VN.

Di seguito sono riportate alcune delle cose di cui tenere conto:

- Controllo RPF, utilizzato per inviare il PIM (*,G) Join attiva la creazione di router adiacenti PIM con un timer di scadenza di due minuti. Se non vengono inviati messaggi di aggiunta PIM per 2 minuti, si verifica il timeout del router adiacente.
- PIM deve creare in modo esplicito una struttura adiacente per la RLOC corrispondente in quanto i messaggi PIM Hello non vengono inviati nella sovrapposizione 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	LISPO.4100	4d23h/00:01:37	v2	0 /

Anycast RP crea (*,G)

In base all'unione PIM (*,G) ricevuta da Edge-2, Border-1 crea (*,G) con l'olio verso l'RLOC di Edge-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
```

Registrazione origine multicast

La sorgente multicast 10.47.7.2 invia il traffico multicast, che entra in Edge-1. Edge-1 porta il pacchetto allo stato S,G (CPU to create) e Edge-1 registra l'origine sull'Anycast RP.

```
<#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
```

Quando Border-1 riceve il pacchetto multicast tramite la registrazione PIM, Border-1 lo ha (S,G) e lo annuncia a Border-2 tramite 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:

Annuncio origine MSDP

Utilizzare il comando "show ip msdp vrf <nome VN> sa-cache" per visualizzare la cache attiva di origine. È possibile utilizzare il comando "show ip msdp vrf <nome VN> summary" per visualizzare il 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 riceve (S,G) informazioni dal Border-1 tramite annuncio MSDP. Se Border-2 ha ricevuto un'unione PIM (*,G) da Edge-2, Border-2 crea una voce (S,G) ed eredita LISP OIL da (*,G) che punta alla RLOC di Edge-2. La regola pratica è che le voci SA di MSDP vengono installate nel database MRIB (Multicast Routing Information Base) solo se è presente un asterisco (*,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, then
```

Border-1 invia un PIM (S,G) Join verso la sorgente 10.47.7.2 per attirare il traffico multicast nativo, che consente a Edge-1 di aggiornare l'OLIO (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:

LISPO.4100,

10.47.1.10

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

<-- RLOC of Border-1

Il traffico multicast da 10.47.7.2 a 239.1.1.1 viene inoltrato alla versione 10.47.6.6 (Bordo-1) tramite incapsulamento VXLAN unicast. Il bordo 1 decapsula il traffico VXLAN e lo incapsula nuovamente sul bordo 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
```

Copertina dell'albero del percorso più breve

Una volta che l'ultimo router hop (LHR) Edge-2 riceve il pacchetto multicast lungo l'albero (*,G), tenta di eseguire il cutover SPT e invia un join PIM (S,G) a 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: LISPO.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: LISPO.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
```

L'FHR (Edge-1) ha (S,G) che punta direttamente verso l'RLOC di 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 - encaps-helper tunnel flag, 1 - 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
```

Verifica del piano dati (indipendente dalla piattaforma)

È possibile che si verifichino diversi problemi che impediscono alla sorgente multicast o al ricevitore multicast di inviare/ricevere il traffico. In questa sezione vengono illustrati i problemi di convalida che possono influire sia sull'origine multicast che sul ricevitore multicast, con particolare attenzione ai problemi non correlati alla programmazione hardware.

Verifica lato origine

Per convalidare l'origine multicast e la capacità FHR di creare (S,G), convalidare Switch Integrated Security Feature (SISF), LISP, Cisco Express Forwarding (CEF) e quindi RPF.

L'origine multicast deve trovarsi in SISF/IP Device-Tracking (IPDT) che gestisce il resto di LISP, CEF e, in ultima analisi, RPF.

È possibile utilizzare il comando "show device-tracking database address <indirizzo IP>" per assicurarsi che l'origine multicast abbia una voce IPDT valida.

```
<#root>
```

```
Edge-1#
```

```
show device-tracking database address 10.47.7.2
```

Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol, DHCP - DHCP
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

Verificare quindi che il database LISP sull'FHR disponga di una voce per l'origine multicast. Utilizzare il 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 voce basata su LISP, CEF punta a una voce host /32, non su 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
```

Successivamente, RPF deriva da 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
```

Se non esiste una voce valida in SISF/IPDT, non viene eseguito alcun mapping del database LISP sull'FHR, con il risultato che l'ECF e l'RPF puntano ai confini. Se la sorgente multicast invia traffico, i punti RPF all'interfaccia errata provocando un errore RPF, (S,G) non viene formato.

<#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 - DHCP  
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 vln 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 interface
```

```
nexthop 10.47.1.11
```

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

```
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
```

Per evitare questo problema, trattare la sorgente multicast come un host silenzioso, dove il binding IP Directed Broadcast, Flooding, Static SISF/IPDT possono risolvere il problema.

Registrazione origine

La registrazione PIM è un flusso di pacchetto unicast, che utilizza LISP/VXLAN come qualsiasi altro pacchetto unicast. Sono necessari diversi controlli per verificare che l'FHR possa registrare correttamente la sorgente multicast nell'RP di Anycast.

Innanzitutto, verificare che l'Anycast RP sia configurato correttamente per 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
```

Verificare che il tunnel del registro PIM sia formato.

```
<#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_10.47.6.1"
```

```
Source :
```

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

```
State : UP
```

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

Garantire la raggiungibilità IP dell'Anycast RTP

```
<#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

Verifica lato ricevitore

- Verificare che il ricevitore multicast stia inviando un messaggio IGMP MR.
- Verificare che lo snooping IGMP sia abilitato. Solo le VN L2 sono l'unico tipo di VN per cui non è abilitato lo snooping IGMP
- Accertarsi che non vi siano ACL della porta, ACL della VLAN e ACL della porta di routing configurati in grado di eliminare il file MR IGMP.
- Convalidare la versione di IGMP MR; per impostazione predefinita, IGMPv2, se il ricevitore multicast è IGMPv3, richiede "ip igmp versione 3"
- Verificare che "ip option drop" non sia configurato

Verifica LHR PIM (*,G)

- Assicurarsi che LHR sia il DR PIM per la subnet/segmento del ricevitore
- verificare che non sia configurato ip multicast group-range
- Accertarsi che non vi siano ACL della porta, ACL della VLAN e ACL della porta di routing configurati in grado di eliminare il file MR IGMP.
- Accertarsi che non vi sia un CPU elevato o che il Control-Plane Policing (CoPP) non scarti il MR IGMP.

Verifica albero condiviso PIM LHR

Verificare che RP sia configurato per il gruppo 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:

Accertarsi che la RPF all'Anycast RP sia corretta

<#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
```

Inoltro MFIB - Verifica lato origine

Per ulteriori informazioni sull'inoltro dei pacchetti, usare il comando "show ip mfib vrf <nome VN> <gruppo multicast> <origine unicast> verbose"

<#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
LISPO.4100,
```

```
10.47.1.13
```

```
Flags: RF F NS
```

```
<-- RLOC of Edge-2
```

```
CEF: Adjacency with MAC: 4500000000004000001164770A2F010D0A2F010C000012B5000000000840000000100400
```

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

```
Edge-1#
```

```
show adjacency lisp0.4100
```

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

```
10.47.1.13
```

```
(8)
```

```
Edge-2#
```

```
show adjacency lisp0.4100 10.47.1.13 detail
```

```
Protocol Interface Address
IP LISPO.4100
```

```
10.47.1.13
```

```
(8)
```

```
0 packets, 0 bytes
```

```
epoch 0
sourced in sev-epoch 14
Encap length 50
4500000000004000001164770A2F010D
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
```

Un EPC può essere usato per convalidare l'incapsulamento VXLAN del pacchetto multicast

<#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
```

Inoltro MFIB - Verifica lato ricevitore

La rete sottostante instrada questo pacchetto da Edge-1 a Edge-2 utilizzando il routing 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: LISPO.4100, RPF nbr

10.47.1.12

Outgoing interface list:

vlan1025

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

Con il comando "show ip mfib vrf <nome VN> <indirizzo gruppo> <origine unicast>" potete aumentare i contatori di inoltro hardware

<#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
```

È possibile utilizzare i contatori multicast di uscita per verificare se il traffico multicast ha lasciato o meno il LHR verso il ricevitore multicast. Usare il comando "show controller ethernet-controller <interface> | include Multicast|Transmit"

<#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

```

Un altro modo per convalidare il traffico multicast in uscita da LHR è quello di eseguire un EPC verso il ricevitore 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

```

Verifica del piano dati (dipendente dalla piattaforma)

(S,G) Creazione - Percorso punt CPU

Per consentire alla FHR di creare lo stato (S,G), una coppia di pacchetti multicast inviati dall'origine multicast vengono trasmessi alla CPU per essere elaborati dal MFIB. I pacchetti multicast vengono inviati alla coda 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

Inoltre, la coda CoPP per i dati MCAST non deve presentare alcuna perdita. Usare il comando "show platform hardware fed active qos queue stats internal cpu policer" | include MCAST Data|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

Se il traffico proviene da un'origine connessa direttamente, viene elaborato dalla coda LSMPI (Linux Shared Memory Punt Interface) per la "origine connessa direttamente" se proviene da un join (S,G), è "Mcast PIM Signaling"

Usare il comando "show platform software infrastructure lsmapi punt | include Cause|Mcast"

<#root>

Edge-1#

```
show platform software infrastructure lsmapi 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	
Mcast punt to RP	0	0	0	0	0
Mcast UDLR	0	0	0	0	0

Successivamente, è possibile eseguire un'acquisizione dei pacchetti FED Punject per visualizzare i pacchetti multicast provenienti dall'origine e dal gruppo e ionizzare la CPU, che conferma l'interfaccia in entrata e la coda della 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], pa1:

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.ffffe, 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
```

Programmazione hardware Mroute - IOS Mroute

La programmazione hardware di (S,G) utilizza la stessa struttura di qualsiasi altro percorso di programmazione: da 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

Programmazione hardware Mroute - MFIB IOS

Le route multicast vengono quindi aggiunte al database MFIB (Multicast Forwarding Information Base), in modo simile all'aggiunta del database RIB (Routing Information Base) al file CEF (Cisco Express Forwarding), mentre l'equivalente multicast è il file MFIB.

<#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
```

```
<-- 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: 450000000004000001164770A2F010D0A2F010C000012B5000000000840000000100400
Pkts: 0/0/0 Rate: 0 pps
```

Programmazione hardware Mroute - RP MFIB

Usare il 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
```

Programmazione hardware Mroute - FP MFIB

La voce RP FMAN per le stesse route include un ID AOM (Asynchronous Object Manager), che viene utilizzato per convalidare ulteriori programmazioni.

Usare il 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)

Programmazione hardware Mroute - Oggetti Mroute

Con gli ID AOM, controllate l'oggetto e gli oggetti padre sia per (*,G) che per (S,G) utilizzando i comandi di gestione degli oggetti. È possibile usare il comando "show platform software object-manager switch active f0 object <ID AOM>" o "show platform software object-manager switch active f0 object <ID AOM>"

Ogni route ha due oggetti padre. Uno degli oggetti fa riferimento alla tabella ipv4_mcast, l'altro è un mlist, che viene utilizzato nei comandi successivi.

<#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
```

Programmazione Mroute Hardware - Oggetti Mlist

Gli oggetti MLIST sono una combinazione di interfacce in ingresso ed elenchi di interfacce in uscita. È possibile usare il 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
```

```
      LISPO.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)

Mroute Programmazione hardware - FED Mroute

Per convalidare la programmazione FED, usare il 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
LISPO.4100 F NS
(Adj: 0x5c)

L'indice di riscrittura fornisce informazioni sull'incapsulamento del traffico multicast, che è ciò che utilizza la replica headend. È possibile usare il comando "show platform hardware feed 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/13u_ri_index0:0x0 in
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 00 00
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>
```

Quindi, accettare i valori URI per convalidare l'intervallo di riscrittura dell'indice. Utilizzare il comando "show platform hardware feed 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

Quindi, eseguire l'RI dal comando precedente per un'ulteriore verifica. Usare il 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
LISPO.4100 F NS

(Adj: 0x5c )
```

Usare il 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	LISPO.4100	4500.0000.0000	0x7f45d9a4a5e8	0x7f45d9a4a798	0x60		

```
0x5c
```

```
2024/08/21 16:18:58.948
```

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

Sull'LHR, è possibile convalidare l'indice di destinazione per vedere a dove viene inoltrato il pacchetto multicast, ossia i ricevitori multicast. È possibile usare il 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 :
LISPO.4100 A
OIF :
Vlan1025 F NS
LISPO.4100 A
(Adj: 0xf8000171 )
```

Utilizzare il gestore DI nel comando "show platform hardware fed switch active fwd-asic abstraction print-resource-handle <ID handle> 1"

```
<#root>
```

```
Edge-2#
```

```
show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f0efdad7668
```

```
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/13u_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 00
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

=====

Informazioni su questa traduzione

Cisco ha tradotto questo documento utilizzando una combinazione di tecnologie automatiche e umane per offrire ai nostri utenti in tutto il mondo contenuti di supporto nella propria lingua. Si noti che anche la migliore traduzione automatica non sarà mai accurata come quella fornita da un traduttore professionista. Cisco Systems, Inc. non si assume alcuna responsabilità per l'accuratezza di queste traduzioni e consiglia di consultare sempre il documento originale in inglese (disponibile al link fornito).