

# Verify Head-End Replication in SD-Access Fabric

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

[Requirements](#)

[Components Used](#)

### [Background Information](#)

### [Topology](#)

### [Configuration](#)

### [Control Plane Verification](#)

[Multicast Receiver sends IGMP Membership Report](#)

[PIM Sparse-Mode Shared Tree Creation](#)

[PIM Neighbors in the Overlay](#)

[Anycast RP creates \(\\*.G\)](#)

[Multicast Source Registration](#)

[MSDP Source Advertisement](#)

[Shortest Path Tree \(SPT\) Cutover](#)

### [Data Plane Verification \(Platform Independent\)](#)

[Source Side Verification](#)

[Source Registration](#)

[Receiver Side Verification](#)

[LHR PIM \(\\*.G\) Verification](#)

[LHR PIM Shared Tree Verification](#)

[MFIB Forwarding - Source Side Verification](#)

[MFIB Forwarding - Receiver Side Verification](#)

### [Data Plane Verification \(Platform Dependent\)](#)

[\(S,G\) Creation - CPU Punt Path](#)

[Mroute Hardware Programming - IOS Mroute](#)

[Mroute Hardware Programming - IOS MFIB](#)

[Mroute Hardware Programming - RP MFIB](#)

[Mroute Hardware Programming - FP MFIB](#)

[Mroute Hardware Programming - Mroute Objects](#)

[Mroute Hardware Programming - Mlist Objects](#)

[Mroute Hardware Programming - FED Mroute](#)

---

## Introduction

This document describes how to troubleshoot HeadEnd Replication in SD-Access (SDA) fabric.

## Prerequisites

## Requirements

Cisco recommends that you have knowledge of these topics:

- Internet Protocol (IP) Forwarding
- Locator/ID Separation Protocol (LISP)
- Protocol Independent Multicast (PIM) Sparse-Mode

## Components Used

- C9000v on Cisco IOS® XE 17.10.1
- Cisco Catalyst Center Version 2.3.5.3

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

This document can also be used with these hardware and software versions:

- C9200
- C9300
- C9400
- C9500
- C9600
- Cisco IOS® XE 16.12 and later

## Background Information

SDA Headend Replication is a form of overlay multicast, which is used to carry multicast traffic between fabric devices, encapsulating multicast traffic into a unicast IP header. Headend replication can route multicast traffic between source(s) and receiver(s) either in the same VLAN or different VLAN (same-VLAN multicast can be routed).

Multicast traffic between sources and receivers on the same Fabric Edge is not forwarded using overlay multicast (VXLAN encapsulation), but is locally routed by the Fabric Edge.

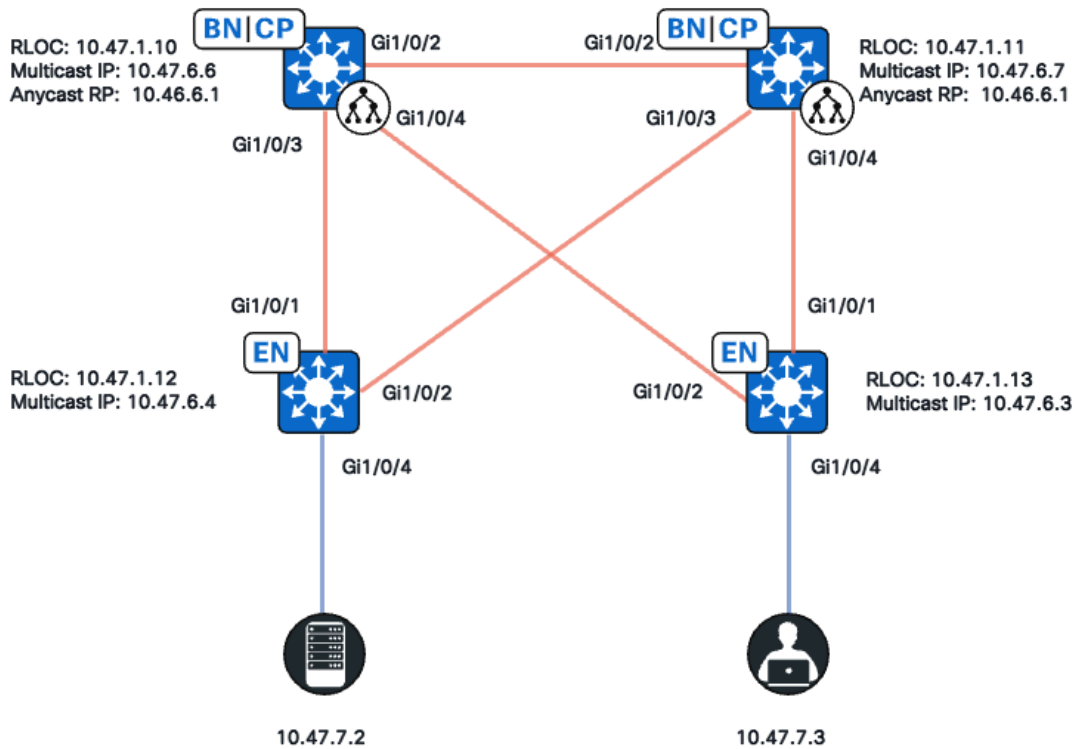
Any form of overlay multicast (headend or native) cannot route multicast traffic for groups in the 224.0.0.0/24 range, or with TTL=1, this is handled via Layer 2 Flooding.



**Note:** Platform (fed) commands can vary. Command can be "**show platform fed <active|standby>**" versus "**show platform fed switch <active|standby>**". If the syntax noted in the examples do not parse out, please try the variant.

## **Topology**

---



### Network Topology

In this topology:

- 10.47.10 and 10.47.1.11 are Collocated Borders that also function as Anycast Rendezvous Point (RP) with Multicast Source Discovery Protocol (MSDP) between the two in the Virtual Network (VN)/VRF.
- 10.47.1.12 and 10.47.1.13 are Fabric Edge Nodes
- 10.47.7.2 is the multicast source
- 10.47.7.3 is the multicast receiver
- 239.1.1.1 is the multicast Group Destination Address (GDA)

## Configuration

It is assumed that Cisco Catalyst Center is used to provision the SDA fabric with the default settings:

- Replication Implementation is Headend Replication
- Anycast RP with MSDP for an Any Source Multicast (ASM) multicast on the collocated borders

After successful configuration from Catalyst Center, relevant configuration per device contains several sections:

### Fabric Edge (10.47.1.12) Configuration

```
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
```

### **Fabric Edge (10.47.1.13) Configuration**

```
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
```

### **Collocated Border/Anycast RP (10.47.1.10) Configuration**

```
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
```

### **Collocated Border/Anycast RP (10.47.1.11) Configuration**

---

```
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
```

## Control Plane Verification

Next, verify Internet Group Membership Protocol (IGMP) and PIM.

### Multicast Receiver sends IGMP Membership Report

The multicast receiver (10.47.7.3) sends an IGMP Membership Report (MR) or IGMP Join to indicate interest in receiving multicast traffic. You can configure an embedded packet capture (EPC) to confirm that there is an IGMP MR received:

```
<#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

Next, ensure that the Fabric Edge is the PIM Designated Router (DR) for the VLAN the multicast receiver is in. This is also referred to as the Last Hop Router (LHR). You can use the command **"show ip pim vrf <VN Name> interface vlan <vlan> detail | include 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)

Validate that IGMP Snooping has picked up the IGMP MR with the command **"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

---



---

## PIM Sparse-Mode Shared Tree Creation

Edge-2, which is the DR on that segment sends a (\*,G) PIM Join towards the Anycast RP. If the Anycast RP address is not resolved in the LISP Map-Cache, the LISP EID Watch process is responsible for triggering LISP Map-Requests. You can use the command "**show lisp instance-id <LISP L3 IID> 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: LISPO.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
```

Validate the (\*,G) entry on Edge-2 with the command "**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: 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
```

---

---

## PIM Neighbors in the Overlay

Once the RPF neighbor that is represented by its routing locator (RLOC) and is reachable via the LISP interface, it is added as a PIM neighbor in the VRF/VN.

Several things to be aware of:

- RPF check, which is used to send the PIM (\*,G) Join triggers PIM neighbor creation with a two minute expiration timer. If no PIM Join messages are sent for 2 minutes, the neighbor times out.
- PIM has to explicitly create a neighbor structure for the corresponding RLOC as PIM Hello messages are not sent in the SDA Overlay

<#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 /
```

## Anycast RP creates (\*,G)

Based on the PIM (\*,G) Join received from Edge-2, Border-1 creates (\*,G) with the OIL towards Edge-2's RLOC

<#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 - 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), 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
```

## Multicast Source Registration

Multicast source 10.47.7.2 sends multicast traffic, which ingresses Edge-1. Edge-1 punts the packet up to the CPU to create (S,G) state, and Edge-1 registers the source to the 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
```

---

---

Once Border-1 receives the multicast packet via PIM Registration, Border-1 has (S,G) and advertises this to Border-2 via 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:

## MSDP Source Advertisement

Use the command "**show ip msdp vrf <VN name> sa-cache**" to view the Source Active Cache. You can use the command "**show ip msdp vrf <VN name> summary**" to see the MSDP peer

<#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/ Downtime	Reset Count	SA Count	Peer Name
--------------	----	-------	---------------------	----------------	-------------	-----------

10.47.6.7

23456

Up

1w1d	0	1
------	---	---

Border-2 receives (S,G) information from Border-1 via MSDP announcement. If Border-2 had received a PIM (\*,G) Join from Edge-2, Border-2 creates a (S,G) entry and inherits the LISP OIL from the (\*,G) pointing to Edge-2's RLOC. The rule of thumb is that MSDP SA entries are only installed in the multicast routing information base (MRIB) if an (\*,G) exists.

<#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 - encaps-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 sends a PIM (S,G) Join towards the source 10.47.7.2 to attract the multicast traffic natively, which allows Edge-1 to update the (S,G) OIL

<#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

Multicast traffic from 10.47.7.2 to 239.1.1.1 is forwarded out 10.47.6.6 (Border-1) via Unicast VXLAN encapsulation. Border-1 de-encapsulates the VXLAN traffic, and re-encapsulate it to 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

## Shortest Path Tree (SPT) Cutover

Once the Last Hop Router (LHR) Edge-2, receives the multicast packet along the (\*,G) tree, it attempts to do SPT cutover and send a PIM (S,G) join to 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 - encaps-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
```

The FHR (Edge-1) has (S,G) pointing directly towards Edge-2's RLOC

```
<#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: LISPO.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:
```

```
LISPO.4100,
```

```
10.47.1.13
```

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

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

## Data Plane Verification (Platform Independent)

There can be various issues that can prevent the multicast source or multicast receiver from sending/receiving the traffic. This section focuses on validation of issues that can impact both the multicast source and multicast receiver, with emphasis on issues that do not related to hardware programming.

### Source Side Verification

To validate the multicast source and the FHR ability to create (S,G), validate Switch Integrated

---

---

Security Feature (SISF), LISP, Cisco Express Forwarding (CEF), and then RPF.

The multicast source must be in SISF/IP Device-Tracking (IPDT) which drives the rest of LISP, CEF, and ultimately, RPF.

You can use the command "**show device-tracking database address <IP address>**" to ensure that the multicast source has a valid IPDT entry.

```
<#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
```

Next, ensure that the LISP database on the FHR has an entry for the multicast source. Use the command "**show lisp instance-id <LISP L3 IID> 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 creates an entry based on LISP, CEF points to a /32 host entry, not 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
```

Next, RPF is derived from 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
```

If there is not a valid entry in SISF/IPDT, which results in no LISP database mapping on the FHR, which results in CEF and RPF pointing to the Borders. If the multicast source sends traffic RPF points to the incorrect interface, which results in RPF failure, (S,G) is not formed.

<#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 - 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 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
```

To prevent this, treat the multicast source as a silent host, where IP Directed Broadcast, Flooding, Static SISF/IPDT bindings can overcome this issue.

## Source Registration

PIM registration is a unicast packet flow, which uses LISP/VXLAN like any other unicast packet. There are several requisite checks to validate that the FHR can properly register the multicast source to the Anycast RP.

First, ensure that the Anycast RP is configured correctly for the 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
```

Ensure the PIM Register tunnel is formed.

<#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)
```

Ensure IP reachability to the Anycast RP

<#root>

Edge-1#

```
show ip cef vrf blue_vn 10.47.6.1
```

```
10.47.6.1/32
```

```
nexthop
```

```
10.47.1.10
```

```
LISPO.4100
```

```
<-- RLOC of Border-1
```

```
nexthop
```

```
10.47.1.11
```

```
LISPO.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

## Receiver Side Verification

- Ensure that the multicast receiver is sending an IGMP MR
- Ensure that IGMP Snooping is enabled. L2 only VNs are the only type of VN that do not have IGMP Snooping enabled
- Ensure there is no Port ACL, VLAN ACL, Routed Port ACL configured that would drop the IGMP MR
- Validate the version of the IGMP MR, by default it is IGMPv2, if the multicast receiver is IGMPv3, that requires "ip igmp version 3"
- Ensure "ip option drop" is not configured

## LHR PIM (\*,G) Verification

- Ensure that the LHR is the PIM DR for the receiver subnet/segment
- Ensure that there is not "ip multicast group-range" is configured
- Ensure there is no Port ACL, VLAN ACL, Routed Port ACL configured that would drop the IGMP MR
- Ensure there is no high CPU or Control-Plane Policing (CoPP) dropping the IGMP MR

## LHR PIM Shared Tree Verification

Ensure that the RP is configured for the 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), 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:

Ensure that RPF to the Anycast RP is correct

<#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
```

## MFIB Forwarding - Source Side Verification

You can use the command "**show ip mfib vrf <VN Name> <multicast group> <unicast source> verbose**" to get additional information about packet forwarding

<#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: 4500000000004000001164770A2F010D0A2F010C000012B5000000000840000000100400

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

---

---

```
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
```

An EPC can be used to validate VXLAN encapsulation of the multicast packet

```
<#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
```

## MFIB Forwarding - Receiver Side Verification

The underlay network routes this packet from Edge-1 to Edge-2 using unicast routing.

```
<#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:

With the command "**show ip mfib vrf <VN Name> <group address> <unicast source> counters**" you can ensure that Hardware Forwarding counters increase

<#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
```

You can use egress multicast counters to validate whether or not the multicast traffic has left the LHR, towards the multicast receiver. Use the command "**show controllers 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  
5733 Unicast frames  
5733 Unicast bytes

100813570 Total bytes  
949456 Unicast frames  
93573016 Unicast bytes

4388569

Multicast frames

32348 Multicast frames

4388569

Multicast bytes

7236641 Multicast bytes

Another way to validate multicast traffic leaving the LHR is to do an EPC towards the multicast receiver.

<#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

## Data Plane Verification (Platform Dependent)

### (S,G) Creation - CPU Punt Path

In order for the FHR to create (S,G) state, a couple of the multicast packets sent from the multicast source are punted up to the CPU to be processed by the MFIB. The multicast packets are sent to the FED queue "CPU\_Q\_MCAST\_DATA"

<#root>

Edge-1#

show platform software fed switch active punt cpug 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
-----

```

Additionally, the CoPP queue for MCAST Data must not have any loss. Use the command "**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	P1cIdx	Queue	Name	Enabled	Rate	Rate
30	9	MCAST	Data	No	500	400

If the traffic comes from a directly connected source, it is processed by the Linux Shared Memory Punt Interface (LSMPI) queue for "directly connected source" if it is from an (S,G) Join, it is "Mcast PIM Signaling"

Use the command "**show platform software infrastructure lsmpi punt | include Cause|Mcast**"

<#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	0	0	0	0
Mcast IPv4 Options data packet	0	0	0	0	0	0	0	0
Mcast Internal Copy	0	0	0	0	0	0	0	0
Mcast IGMP Unroutable	0	0	0	0	0	0	0	0
<b>Mcast PIM signaling</b>								
	0	0	0	0	0	0	0	0
Mcast punt to RP	0	0	0	0	0	0	0	0
Mcast UDLR	0	0	0	0	0	0	0	0

Next, a FED Punject packet-capture can be done to see multicast packets from the source and group ion the CPU, which confirms the incoming interface and CPU queue.

```
<#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.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
```

## Mroute Hardware Programming - IOS Mroute

Hardware programming of the (S,G) uses the same structure as any other programming path: IOS to FMAN RP to FMAN FP, to 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 - encaps-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: LISPO.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
```

## Mroute Hardware Programming - IOS MFIB

Multicast routes are then added to the Multicast Forwarding Information Base (MFIB), which is similar to how the Routing Information Base (RIB) is added to Cisco Express Forwarding (CEF), the multicast equivalent is the 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: 4500000000004000001164770A2F010D0A2F010C000012B5000000000840000000100400  
Pkts: 0/0/0 Rate: 0 pps

## Mroute Hardware Programming - RP MFIB

Use the command "**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

## Mroute Hardware Programming - FP MFIB

The FMAN RP entry for the same mroutes include an Asynchronous Object Manager (AOM) ID, this AOM ID is used to validate further programming.

Use the command "**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)
```

## Mroute Hardware Programming - Mroute Objects

With the AOM IDs, check the object and parent objects for both (\*,G) and (S,G) using the object-manager commands. You can use the command "**show platform software object-manager switch active f0 object <AOM ID>**" or "**show platform software object-manager switch active f0 object <AOM ID> parents**"

Each mroute has two parent objects. One of the objects references the ipv4\_mcast table, the other is an mlist, which is used in subsequent commands.

```
<#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

## Mroute Hardware Programming - Mlist Objects

The MLIST objects are a combination of both Incoming Interfaces and Outgoing Interface Lists. You can use the command "**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)

---



---

```
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>
```

Next, take the URI values to validate the rewrite-index range. Use the command "**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

Next, take the RI from the previous command for further verification. Use the command "**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 )
```

Use the command "**show platofmr 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
```

---

---

On the LHR, you can validate the destination index to see where the multicast packet is forwarded to, which is the multicast receiver(s). You can use the command "**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 )
```

Take the DI Handler and use in the command "**show platform hardware fed switch active fwd-asic abstraction print-resource-handle <DI 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
```

---

---

a1\_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  
a1\_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  
=====

---