Troubleshoot SDA Forwarding East-West Traffic Flow

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Related Products

Background Information

Topology

Configuration

Host Onboarding Verification

- IPDT / IP Device Tracking Entry
- MAC / ARP Entry
- LISP Entry

ARP Resolution in SDA

Basic Host Reachabilityin SDA Fabric (Same VLAN / Same VN)

Basic HostReachabilityin SDA Fabric (Different VLANs / Same VN)

Introduction

This document describes how to validate East-West Traffic Flow as part of Software Defined Access (SDA).

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Internet Protocol (IP) Forwarding
- Locator/ID Separation Protocol (LISP)

Components Used

The information in this document is based on these software and hardware versions:

- C9000v on Cisco IOS® XE 17.10.1
- SDA 1.0 (not LISP PubSub)

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.

Related Products

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 East-West Traffic Flow refers to the concept where an endpoint within the SDA Fabric wants to communicate with another endpoint within the same fabric. There are caveats as to what is and is not considered an East-West flow. An East-West traffic flow can be these examples:

- Endpoints that are in the same subnet (172.17.10.2 talking to 172.19.10.3) this is considered L2LISP Extension
- Endpoints that are in the same VRF (VN) (172.19.10.2 talking to 172.19.11.2 and both are in VRF Campus) this is considered L3 LISP
- Endpoint that is within the fabric talking to a host that is connected to an L2 Handoff Border, exactly the same as L2LISP

East-West traffic flows do not refer to these examples:

- Traffic originated from the SDA Fabric to outside the fabric, that is North-South
- Inter-VRF routing is also not considered East-West (An endpoint in VRF Campus, IP address 172.19.10.2 talking to an endpoint in VRF Guest, IP address 172.19.11.2)
- SD-WAN Integrated Domains
- SDA Transit
- Border Affinity
- Extranet

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

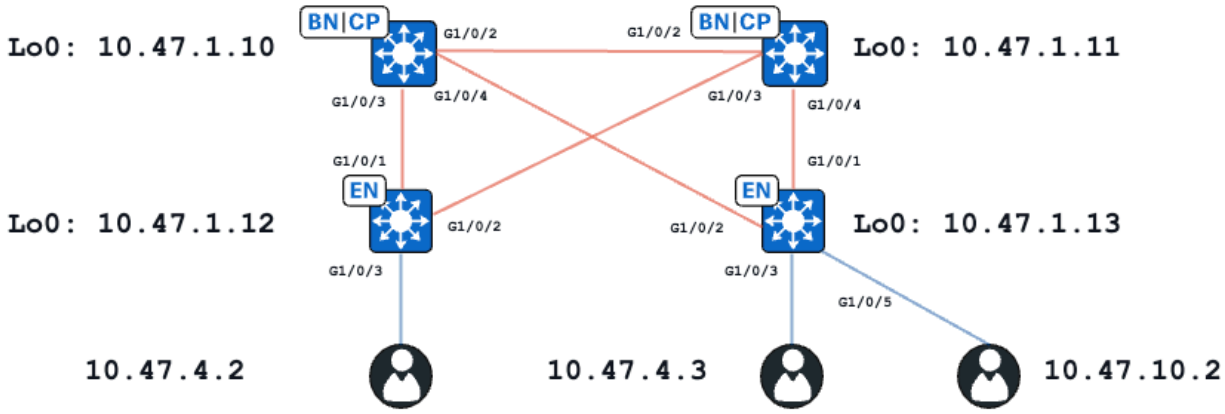

For the purpose of this example, C9000v switches function as the Fabric Edges and Collocated Borders. All endpoints are within the same Virtual Network (VN), red_vn. Endpoints at 10.47.4.2 and 10.47.4.2 are in the same subnet, the endpoint at 10.47.10.2 is in a different subnet but same VN.

Configuration

It is assumed that CIsco DNA-Center is used to provision SDA fabric with the default settings:

- Layer-2 extension is enabled (this forces traffic to be forwarded based on MAC address lookups rather than IP address lookups).
- Layer-2 flooding is disabled (this enables ARP supression on edge devices and LISP-assisted ARP learning).

After proper host onboarding process, the interface configuration contains several sections:

Fabric Edge (10.47.1.12) Interface Configuration:

```
interface GigabitEthernet1/0/3
 switchport access vlan 1026
 switchport mode access
 device-tracking attach-policy IPDT_POLICY
 spanning-tree portfast
 spanning-tree bpduguard enable
end

interface Vlan1026
 description Configured from Cisco DNA-Center
 mac-address 0000.0c9f.f341
 vrf forwarding red_vn
 ip address 10.47.4.1 255.255.255.0
 ip helper-address 10.47.9.9
 no ip redirects
 ip route-cache same-interface
 no lisp mobility liveness test
 lisp mobility red-IPV4
end
```

Fabric Edge (10.47.1.12) LISP Configuration:

```
router lisp
 locator-table default
 locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f
  IPv4-interface Loopback0 priority 10 weight 10
  exit-locator-set
!
instance-id 4099
  remote-rloc-probe on-route-change
  dynamic-eid red-IPV4
   database-mapping 10.47.4.0/24 locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f
   exit-dynamic-eid
  !
  dynamic-eid red-helpdesk-IPV4
   database-mapping 10.47.10.0/24 locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f
   exit-dynamic-eid
  !
  service ipv4
   eid-table vrf red_vn
   map-cache 0.0.0.0/0 map-request
   sgt distribution
   sgt
   exit-service-ipv4
  !
  exit-instance-id
!
 !
 instance-id 8190
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 1026
   database-mapping mac locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f
   dynamic-eid detection multiple-addr bridged-vm
   exit-service-ethernet
  !
  exit-instance-id
 !
 instance-id 8192
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 1028
   database-mapping mac locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f
   dynamic-eid detection multiple-addr bridged-vm
   exit-service-ethernet
  !
  exit-instance-id
```

Fabric Edge (10.47.1.13) Interface Configuration:

```
interface GigabitEthernet1/0/3
 switchport access vlan 1026
 switchport mode access
 device-tracking attach-policy IPDT_POLICY
 spanning-tree portfast
 spanning-tree bpduguard enable
end
!
interface GigabitEthernet1/0/5
```

```
 switchport access vlan 1028
 switchport mode access
 device-tracking attach-policy IPDT_POLICY
 spanning-tree portfast
 spanning-tree bpduguard enable
end
!
interface Vlan1026
 description Configured from Cisco DNA-Center
 mac-address 0000.0c9f.f341
 vrf forwarding red_vn
 ip address 10.47.4.1 255.255.255.0
 ip helper-address 10.47.9.9
 no ip redirects
 ip route-cache same-interface
 no lisp mobility liveness test
 lisp mobility red-IPV4
end
!
interface Vlan1028
 description Configured from Cisco DNA-Center
 mac-address 0000.0c9f.f800
 vrf forwarding red_vn
 ip address 10.47.10.1 255.255.255.0
 ip helper-address 10.47.9.9
 no ip redirects
 ip route-cache same-interface
 no lisp mobility liveness test
 lisp mobility red-helpdesk-IPV4
end
```

Fabric Edge (10.47.1.13) LISP Configuration

```
router lisp
 locator-table default
 locator-set rloc_691b1fe4-5264-44c2-bb1b-0903b3eb2c51
  IPv4-interface Loopback0 priority 10 weight 10
  exit-locator-set
 !
 instance-id 4099
  remote-rloc-probe on-route-change
  dynamic-eid red-IPV4
   database-mapping 10.47.4.0/24 locator-set rloc_691b1fe4-5264-44c2-bb1b-0903b3eb2c51
   exit-dynamic-eid
  !
  dynamic-eid red-helpdesk-IPV4
   database-mapping 10.47.10.0/24 locator-set rloc_691b1fe4-5264-44c2-bb1b-0903b3eb2c51
   exit-dynamic-eid
  !
  service ipv4
   eid-table vrf red_vn
   map-cache 0.0.0.0/0 map-request
   sgt distribution
   sgt
   exit-service-ipv4
  !
  exit-instance-id
!
```

```
 instance-id 8190
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 1026
   database-mapping mac locator-set rloc_691b1fe4-5264-44c2-bb1b-0903b3eb2c51
   dynamic-eid detection multiple-addr bridged-vm
   exit-service-ethernet
  !
  exit-instance-id
 !
 instance-id 8192
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 1028
   database-mapping mac locator-set rloc_691b1fe4-5264-44c2-bb1b-0903b3eb2c51
   dynamic-eid detection multiple-addr bridged-vm
   exit-service-ethernet
  !
  exit-instance-id
```

Host Onboarding Verification

As part of host on-boarding process, several structures are created:

IPDT / IP Device Tracking Entry

After successful host onboarding, there is a valid entry in the IP Device Tracking (IPDT) table and the end-host is be marked as REACHABLE:

<#root>

```
Edge-1#
```

```
show device-tracking database interface g1/0/3
```

```
portDB has 2 entries for interface Gi1/0/3, 2 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol, DH4 - IPv4 DH
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      ag
DH4 10.47.4.2                              5254.0019.93e9         Gi1/0/3    1026       0024       3mr
```

MAC / ARP Entry

When the end-host is successfully onboarded, it can ping the default gateway (or can be pinged from the default gateway if no firewall is installed on the end-point blocking this communication):

<#root>

```
Edge-1#
ping vrf red_vn 10.47.4.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.47.4.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 142/150/161 ms
```

On the Edge Node, there is a MAC address, as well as corresponding ARP entry in the table (in VRF):

```
<#root>

Edge-1#
show mac address-table interface g1/0/3

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
1026    5254.0019.93e9    DYNAMIC     Gi1/0/3
Total Mac Addresses for this criterion: 1

Edge-1#
show ip arp vrf red_vn

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.47.4.1               -   0000.0c9f.f341  ARPA   Vlan1026

Internet  10.47.4.2               1   5254.0019.93e9  ARPA   Vlan1026

Internet  10.47.10.1              -   0000.0c9f.f800  ARPA   Vlan1028
```

Software FED MAC Address Programming**

To check the MAC address in FED use the command **show platform software fed switch active matm macTable vlan <vlan id> mac <mac address>**

```
<#root>

Edge-1#
show platform software fed switch active matm macTable vlan 1026 mac 5254.0019.93e9

VLAN   MAC                 Type  Seq#    EC_Bi  Flags

machandle


siHandle

            riHandle

diHandle

              *a_time  *e_time  ports                                                    Con
```

```
-------------------------------------------------------------------------------------------------------------
1026   5254.0019.93e9        0x1       9      0      0
```

0x7f65ec7bda68

0x7f65ec7c21f8

```
      0x0
```

0x7f65ec6e1368

```
            300        7  GigabitEthernet1/0/3                                          Yes

======platform hardware details ======

Asic: 0
 htm-handle = 0x7f65ec95dc68 MVID = 7 gpn = 1
 SI = 0xc3 RI = 0x25 DI = 0x526e
 DI = 0x526e pmap = 0x00000000 0x00000004 pmap_intf : [GigabitEthernet1/0/3]

Asic: 1
 SI = 0xc3 RI = 0x25 DI = 0x526e
 DI = 0x526e pmap = 0x00000000 0x00000000
```

MAC Address macHandle Programming

Take the macHandle value from the previous command (**0x7f65ec7bda68**) and use in **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <macHandle> 1**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec7bda68 1

```
Handle:0x7f65ec7bda68 Res-Type:ASIC_RSC_HASH_TCAM Res-Switch-Num:0 Asic-Num:255 Feature-ID:AL_FID_L2 Lk
priv_ri/priv_si Handle: (nil)Hardware Indices/Handles: handle [ASIC: 0]: 0x7f65ec95dc68
Features sharing this resource:Cookie length: 12
19 00 54 52 e9 93 07 80 07 00 00 00

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
Number of HTM Entries: 1

Entry 0: (handle 0x7f65ec95dc68)

Absolute Index: 6778
Time Stamp: 4
KEY -
```

vlan:7

mac:0x5254001993e9

```
 l3_if:0
```

gpn:3

```
 epoch:0 static:0 flood_en:0 vlan_lead_wless_flood_en: 0 client_home_asic: 0 learning_peerid 0, learnin
MASK - vlan:0 mac:0x0 l3_if:0 gpn:0 epoch:0 static:0 flood_en:0 vlan_lead_wless_flood_en: 0 client_home_
SRC_AD - need_to_learn:0 lrn_v:0 catchall:0 static_mac:0 chain_ptr_v:0 chain_ptr: 0 static_entry_v:0 au
DST_AD - si:0xb7 bridge:0 replicate:0 blk_fwd_o:0 v4_rmac:0 v6_rmac:0 catchall:0 ign_src_lrn:0 port_mas
============================================================
```

****MVID Verification****

The number 7 in the previous output is the Mapped VLAN ID (MVID) in hardware. To verify that they match up to the "real" vlan, use **show platform software fed switch active vlan <vlan number>**

```
<#root>

Edge-1#

show platform software fed switch active vlan 1026

VLAN Fed Information

Vlan


Id

  IF Id                LE Handle            STP Handle           L3 IF Handle         SVI IF ID

MVID


----------------------------------------------------------------------------------------------------

1026

     0x0000000000420011   0x00007f65ec6a08b8   0x00007f65ec6a1138   0x00007f65ec77e838   0x000000000000

7
```

****Global Port Number (GPN) Verification****

To correlate the GPN with a "real" interface, use the command **show platform software fed switch active ifm mappings gpn**

```
<#root>

Edge-1#

show platform software fed switch active ifm mappings gpn

Mappings Table

GPN         Interface                       IF_ID               IF_TYPE
-----------------------------------------------------------------------------------------------
1           GigabitEthernet1/0/1            0x0000001a          ETHER
2           GigabitEthernet1/0/2            0x0000001b          ETHER
```

```
3


GigabitEthernet1/0/3

             0x0000000b          ETHER

<-- GPN 3 lines up with the expected Egress interface
```

MAC Address siHandle Programming

Take the siHandle value from the previous command (**0x7f65ec7c21f8**) and utilize in **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_handle> 1**

```
<#root>

 Edge-1#

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec7c21f8 1

Handle:0x7f65ec7c21f8 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f65ec7c2498Hardware Indices/Handles: index0:0xc3  mtu_index/l3u_ri_index0:0x
Features sharing this resource:66 (1)]
57 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 02 04 00 00 00 00 00 00 00 00 00 00 07 00 52 54 00 19 93 e9 00 00 00 00 00 00 0


Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------


Station Index (SI) [0xc3] <-- Station Index is comprised of the Rewrite Index (RI) and Destination Index
RI = 0x25 <-- Rewrite Index contains the forwarding information
DI = 0x526e <-- Destination Index contains information related to the outgoing interface


stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0x1

Replication Bitmap: LD <-- Local Data (LD) indicates that the destination is on this ASIC


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------


Station Index (SI) [0xc3] <-- Station Index is comprised of the Rewrite Index (RI) and Destination Index
RI = 0x25 <-- Rewrite Index contains the forwarding information
DI = 0x526e <-- Destination Index contains information related to the outgoing interface


stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
```

```
dejaVuPreCheckEn = 0x1
```

Replication Bitmap: CD <-- Core Data (CD) indicates that the destination is on the same ASIC, different

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

MAC Address Rewrite-Index Verification

Take the RI value from the previous command (**0x25**) and utilize in **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x25 0x25

```
ASIC#:0 RI:37 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr:
```

MAC Addr: 52:54:00:19:93:e9

```
,
 L3IF LE Index 41


ASIC#:0 RI:38 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr: MAC Addr: 01:00:5e:00:00:00,
 L3IF LE Index 40


ASIC#:0 RI:39 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr: MAC Addr: 52:54:00:00:50:17,
 L3IF LE Index 40


ASIC#:1 RI:37 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr:
```

MAC Addr: 52:54:00:19:93:e9

```
,
 L3IF LE Index 41


ASIC#:1 RI:38 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr: MAC Addr: 01:00:5e:00:00:00,
 L3IF LE Index 40


ASIC#:1 RI:39 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 MAC Addr: MAC Addr: 52:54:00:00:50:17,
 L3IF LE Index 40
```

MAC Address Destination-Index Verification

Take the DI value from the previous command (**0x526e**) and utilize in **show platform hardware fed switch**

active fwd-asic resource asic all destination-index range <DI> <DI>

```
<#root>

Edge-1#

show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x526e 0x526

ASIC#0:


Destination index   = 0x526e

pmap                = 0x00000000 0x00000004 <-- Convert decimal 4 to binary, which is 0100. Count this b

pmap_intf : [GigabitEthernet1/0/3]
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
ASIC#1:


Destination index   = 0x526e
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
```

****Port Verification****

To correlate the port that was seen previously, use the command**show platform software fed switch active ifm mappings** and look at the Port column.

```
<#root>

Edge-1#

show platform software fed switch active ifm mappings
```

```
------------------ show platform software fed switch active ifm mappings  ------------------
Interface                IF_ID    Inst Asic Core Port SubPort Mac  Cntx LPN  GPN  Type Active
GigabitEthernet1/0/1     0x1a       0   0   0    0      0      1    0    1    1    NIF  Y
GigabitEthernet1/0/2     0x1b       0   0   0    1      0      2    1    2    2    NIF  Y

GigabitEthernet1/0/3

      0xb        0   0   0

2

      0      3    2    3    3    NIF  Y

<-- Matches port 2 from previous output
```

****Hardware FED MAC Address Verification****

This output in a working/ideal scenario matches what the macHandle decode provided.

```
<#root>

Edge-1#

show platform hardware fed switch active matm macTable vlan 1026 mac 5254.0019.93e9

HEAD: MAC address 5254.0019.93e9 in VLAN 1026
KEY:

vlan 7

,

 mac 0x5254001993e9

, l3_if 0,

gpn 3

, epoch 0, static 0, flood_en 0, vlan_lead_wless_flood_en 0, client_home_asic 0, learning_peerid 0, lea
MASK: vlan 0, mac 0x0, l3_if 0, gpn 0, epoch 0, static 0, flood_en 0, vlan_lead_wless_flood_en 0, clien
SRC_AD: need_to_learn 0, lrn_v 0, catchall 0, static_mac 0, chain_ptr_v 0, chain_ptr 0, static_entry_v
DST_AD: si 0xb7, bridge 0, replicate 0, blk_fwd_o 0, v4_mac 0, v6_mac 0, catchall 0, ign_src_lrn 0, por

Total Mac number of addresses:: 1
```

- **The VLAN ID in hardware (MVID) is 7**
- **MAC Address: 5254.0019.93e9**
- **GPN: 3**

LISP Entry

After successful host onboarding LISP entries for the end-host are created locally on Edge Node, as well as registered on Control Nodes (LISP MSMR - LISP Map Server / Map Resolver). All LISP checks needs to be done in regards to specific instance-id scope that can be checked for L2 and for L3:

```
<#root>

Edge-1#
```

```
show vlan id 1026

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1026 red                              active

L2LI0:8190

, Gi1/0/3

<-- L2 LISP Instance ID tied to VLAN 1026
```

L2 LISP Database Verification

To check the L2 LISP Database use the command **show lisp instance-id <L2 LISP IID> ethernet database <mac address>**

```
<#root>

Edge-1#

show lisp instance-id 8190 ethernet database 5254.0019.93e9

LISP ETR MAC Mapping Database for LISP 0 EID-table Vlan 1026 (IID 8190), LSBs: 0x1
Entries total 1, no-route 0, inactive 0, do-not-register 2
5254.0019.93e9/48, dynamic-eid Auto-L2-group-8190, inherited from default locator-set rloc_222e1707-175
  Uptime: 2d17h, Last-change: 2d17h
  Domain-ID: local
  Service-Insertion: N/A
  Locator     Pri/Wgt  Source     State


10.47.1.12

   10/10   cfg-intf   site-self, reachable

 -----> Our own RLOC

  Map-server       Uptime         ACK  Domain-ID


 10.47.1.10

       1d11h          Yes  0

-----> RLOC of upstream collocated border


10.47.1.11

       2d17h          Yes  0

-----> RLOC of upstream collocated border
```

LISP L2 Address Resolution (AR) Database Verification

To check the LISP L2 AR Database, use the command **show lisp instance-id <LISP L2 IID> ethernet database address-resolution <mac address>**

<#root>

Edge-1#

show lisp instance-id 8190 ethernet database address-resolution 5254.0019.93e9

LISP ETR Address Resolution for LISP 0 EID-table Vlan 1026 (IID 8190)
(*) -> entry being deleted

Hardware Address L3 InstID Host Address
5254.0019.93e9 4099 10.47.4.2/32 <-- Endpoint MAC Address, LISP L3 Instance ID, Endpoint I

****LISP L3 Database Verfication****

To check the LISP L3 Database, use the command **show lisp instance-id <LISP L3 IID> ipv4 database <IP address/Subnet Mask>**

<#root>

Edge-1#

show lisp instance-id 4099 ipv4 database 10.47.4.2/32

LISP ETR IPv4 Mapping Database for LISP 0 EID-table vrf red_vn (IID 4099), LSBs: 0x1
Entries total 1, no-route 0, inactive 0, do-not-register 1

10.47.4.2

/32, dynamic-eid red-IPV4, inherited from default locator-set rloc_222e1707-175d-4019-a783-060404f8bc2f

-----> Endpoint IPv4 Address

Uptime: 2d18h, Last-change: 2d18h
Domain-ID: local
Service-Insertion: N/A
Locator Pri/Wgt Source State

10.47.1.12

10/10 cfg-intf site-self, reachable

-----> Our own RLOC

Map-server Uptime ACK Domain-ID

10.47.1.10

1d11h Yes 0

-----> RLOC of upstream collocated border

10.47.1.11

2d17h Yes 0

-----> RLOC of upstream collocated border

****CEF Verification****

To check CEF, use the command **show ip cef vrf <vrf name> <IP address> internal**

```
<#root>

Edge-1#

show ip cef vrf red_vn 10.47.4.2 internal


10.47.4.2/32, epoch 1, flags [att, sc], RIB[D], refcnt 6, per-destination sharing
  sources: RIB, Adj, IPL
  feature space:
    IPRM: 0x00058000
    Broker: linked, distributed at 3rd priority
  subblocks:
    SC owned,sourced:

LISP local EID

 -
    SC inherited: LISP remote EID - locator status bits 0x00000000
    SC inherited: LISP cfg dyn-EID - LISP configured dynamic-EID
    LISP EID attributes: localEID Yes, c-dynEID Yes, d-dynEID Yes, a-dynEID No
    SC owned,sourced: LISP generalised SMR - [disabled, not inheriting, 0x7F06D0A67E40 locks: 1]
    Adj source:

IP adj out of Vlan1026

,

addr 10.47.4.2

 7F06D300B738
      Dependent covered prefix type adjfib, cover 10.47.4.0/24
    2 IPL sources [no flags]
  ifnums:
    Vlan1026(29): 10.47.4.2
  path list 7F06CEE8D720, 3 locks, per-destination, flags 0x49 [shble, rif, hwcn]
    path 7F06D0A900C8, share 1/1, type attached nexthop, for IPv4
      nexthop 10.47.4.2 Vlan1026, IP adj out of Vlan1026, addr 10.47.4.2 7F06D300B738
  output chain:


IP adj out of Vlan1026, addr 10.47.4.2

 7F06D300B738
```

In addition to local LISP entries on the SDA Edge Node, SDA Control Nodes (LISP MS/MR) also contain extra information about the endpoints:

Collocated Border L2 LISP Server Verification:

To check the L2 LISP Server, use the command **show lisp instance-id <L2 LISP IID> ethernet server <MAC Address>**

```
<#root>

Border-1#

show lisp instance-id 8190 ethernet server 5254.0019.93e9
```

```
LISP Site Registration Information

Site name: site_uci
Description: map-server configured from Cisco DNA-Center
Allowed configured locators: any
Requested EID-prefix:

  EID-prefix:
```

5254.0019.93e9

```
/48 instance-id 8190
```

<-- Endpoint MAC Address

```
    First registered:     2w5d
    Last registered:      3d16h
    Routing table tag:    0
    Origin:               Dynamic, more specific of any-mac
    Merge active:         No
    Proxy reply:          Yes
    Skip Publication:     No
    Force Withdraw:       No
    TTL:                  1d00h
    State:                complete
    Extranet IID:         Unspecified
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
```

ETR

10.47.1.12

```
:21038, last registered 3d16h, proxy-reply, map-notify
```

<-- Egress Tunnel Router (Fabric Edge IP address)

```
                          TTL 1d00h, no merge, hash-function sha1
                          state complete, no security-capability
                          nonce 0xB60C4314-0x97BB332D
                          xTR-ID 0xAB3179F6-0xC774F22C-0x00F2C82E-0x3A66738D
                          site-ID unspecified
                          Domain-ID local
                          Multihoming-ID unspecified
                          sourced by reliable transport
      Locator     Local  State      Pri/Wgt  Scope
```

10.47.1.12

```
  yes    up          10/10   IPv4 none
```

<--(Fabric Edge IP address)

Collocated Border L2 LISP Address Resolution (AR) Server Verification:

To check the L2 LISP AR Server use the command **show lisp instance-id <LISP L2 IID> ethernet server address-resolution <IP address>**

To check the registration history use the command **show lisp instance-id <LISP L2 IID> ethernet server address-resolution <IP address> registration-history**

<#root>

```
Border-1#

show lisp instance-id 8190 ethernet server address-resolution 10.47.4.2

Address-resolution data for router lisp 0 instance-id 8190

Site name: site_uci

Host Address:

10.47.4.2

/32
Hardware Address:

5254.0019.93e9

First registered:     2w5d
Last registered:      3d16h
Registration errors:
  Authentication failures:   0
 ETR

10.47.1.12

:21038
  Last registered:       3d16h
  TTL:                   1d00h
  xTR-ID:                0xAB3179F6-0xC774F22C-0x00F2C82E-0x3A66738D
  Site-ID:               unspecified
  Registered addr:       5254.0019.93e9
  L3 Instance ID:        4099

Border-1#

show lisp instance-id 8190 ethernet server address-resolution 10.47.4.2 registration-history


Map-Server registration history
Roam = Did host move to a new location?
WLC = Did registration come from a Wireless Controller?
Prefix qualifier: + = Register Event, - = Deregister Event, * = AR register event

Timestamp (UTC)      Instance Proto Roam WLC Source
                                            EID prefix / Locator
*Sep 29 16:50:27.762     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
*Oct  1 21:05:11.086     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
*Oct  2 06:51:11.882     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
*Oct  3 00:56:33.642     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
*Oct  3 01:53:45.934     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
*Oct  6 04:36:08.685     8190 TCP   No   No  10.47.1.12
                                           +*10.47.4.2/32 / 5254.0019.93e9
```

Collocated Border L3 LISP Server Verification

To check the L3 LISP Server, use the command **show lisp instance-id <LISP L3 IID> ipv4 server <IP address>**

To check the L3 LISP Server registration history, use the command **show lisp instance-id <LISP L3 IID> ipv4 server <IP address> registration-history**

<#root>

```
Border-1#
```

show lisp instance-id 4099 ipv4 server 10.47.4.2

```
LISP Site Registration Information

Site name: site_uci
Description: map-server configured from Cisco DNA-Center
Allowed configured locators: any
Requested EID-prefix:

  EID-prefix:
```

10.47.4.2

```
/32 instance-id 4099
    First registered:     2w5d
    Last registered:      02:39:39
    Routing table tag:    0
    Origin:               Dynamic, more specific of 10.47.4.0/24
    Merge active:         No
    Proxy reply:          Yes
    Skip Publication:     No
    Force Withdraw:       No
    TTL:                  1d00h
    State:                complete
    Extranet IID:         Unspecified
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
    ETR
```

10.47.1.12

```
:21038, last registered 02:39:39, proxy-reply, map-notify
                          TTL 1d00h, no merge, hash-function sha1
                          state complete, no security-capability
                          nonce 0x128CB668-0xF7B85F77
                          xTR-ID 0xAB3179F6-0xC774F22C-0x00F2C82E-0x3A66738D
                          site-ID unspecified
                          Domain-ID local
                          Multihoming-ID unspecified
                          sourced by reliable transport
      Locator     Local  State      Pri/Wgt  Scope
```

10.47.1.12

```
  yes    up          10/10   IPv4 none

Border-1#
```

show lisp instance-id 4099 ipv4 server 10.47.4.2/32 registration-history

```
Map-Server registration history
Roam = Did host move to a new location?
WLC = Did registration come from a Wireless Controller?
Prefix qualifier: + = Register Event, - = Deregister Event, * = AR register event

Timestamp (UTC)      Instance Proto Roam WLC Source
                                            EID prefix / Locator
*Oct  6 04:36:01.548     4099 UDP   No   No  10.47.1.12
                                          + 10.47.4.2/32
*Oct  6 04:36:08.686     4099 TCP   No   No  10.47.1.12
                                          + 10.47.4.2/32
*Oct  9 18:35:48.058     4099 TCP   No   No  10.47.1.12
                                          + 10.47.4.2/32
```

ARP Resolution in SDA

It is assumed that CIsco Catalyst Center has been used to provision SDA fabric with default settings. This means that Layer-2 extension is enabled and that all traffic within the Fabric (in same VLAN / VN) is forwarded based on MAC address lookups / LISP Ethernet instance, rather than IP address lookups / LISP IP instance.

From troubleshooting perspective, it can be useful to configure static ARP entries on both hosts to quickly check if the issue is with generic connectivity in the fabric (in such case ping does not work between hosts) or only with ARP resolution.

ARP process in SDA Fabric leverages LISP to resolve the identify and location of hosts and is different from ARP behavior in traditional Routing/Switching environments.

Step 1: Fabric Endpoint sends an ARP request to determine the MAC/IP binding for the other Fabric Endpoint

Packet capture can be configured on the ingress interface to confirm that ARP packet is received from the host:


```
Edge-1#
monitor capture 1 interface g1/0/3 in match any
Edge-1#
mon cap 1 start
Started capture point : 1
Edge-1#
mon cap 1 stop
Capture statistics collected at software:
    Capture duration - 22 seconds
    Packets received - 13
    Packets dropped - 0
    Packets oversized - 0
Number of Bytes dropped at asic not collected
```

```
Capture buffer will exists till exported or cleared

Stopped capture point : 1
Edge-1#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    1   0.000000 52:54:00:19:93:e9 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 10.47.4.3? Tell 10.47.4.2
    2   1.028893 52:54:00:19:93:e9 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 10.47.4.3? Tell 10.47.4.2
    3   2.058244 52:54:00:19:93:e9 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 10.47.4.3? Tell 10.47.4.2

Edge-1#
```

show monitor capture 1 buffer display-filter arp detailed

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface /tmp/epc_ws/wif_to_ts_p
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 10, 2023 14:52:03.659290000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1696949523.659290000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 60 bytes (480 bits)
    Capture Length: 60 bytes (480 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:arp]
Ethernet II, Src: 52:54:00:19:93:e9 (
```

52:54:00:19:93:e9

```
), Dst:
```

ff:ff:ff:ff:ff:ff

```
 (ff:ff:ff:ff:ff:ff)
```

<-- SMAC/DMAC respectively

```
    Destination: ff:ff:ff:ff:ff:ff (ff:ff:ff:ff:ff:ff)
        Address: ff:ff:ff:ff:ff:ff (ff:ff:ff:ff:ff:ff)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
    Source: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        Address: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: ARP (
```

0x0806

```
)
    Padding: 000000000000000000000000000000000000
Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
```

```
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address:
```

52:54:00:19:93:e9

```
 (52:54:00:19:93:e9)
    Sender IP address:
```

10.47.4.2

```
    Target MAC address:
```

00:00:00:00:00:00

```
 (00:00:00:00:00:00)
    Target IP address:
```

10.47.4.3

Step 2. Edge Node consumes ARP packet and generate LISP request to determine MAC address of HOST-02.

Edge-1 sends a LISP Map-Request to resolve the MAC Address of 10.47.4.3 to the LISP Control-Planes (Collocated Borders):

<#root>

```
Edge-1#
```

debug lisp control-plane all

```
Edge-1#
```

debug l2lisp all

```
LISP[REMT ]-0: Map Request: Delay is over for IID 8190 EID 10.47.4.3/32, requester 'AR'.
LISP[REMT ]-0 IID 8190: Schedule processing of Map-Requests from 'remote EID prefix' in IPv4.
LISP[REMT ]-0: Map Request: Sending request for IID 8190 EID 10.47.4.3/32, requester 'AR'.
```

Step 3. Control Node receives LISP request for IP/MAC mapping and send a response back to SDA Edge Node

LISP Map-Request is received from the Fabric Edge and responds with a LISP Map-Reply with the MAC address that ties to 10.47.4.3

<#root>

```
Border-1#
```

debug lisp control-plane all

```
Border-1#
```

debug l2lisp all

```
LISP[TRNSP]-0: Processing received Map-Request(1) message on GigabitEthernet1/0/3 from 10.47.4.3:4342 t
LISP[MR   ]-0: Received Map-Request with 1 records, first EID IID 8190 10.47.4.3/32, source EID UNSPEC,
LISP[MR   ]-0 IID 8190 Eth-ARP: MS EID 10.47.4.3/32: Sending proxy reply to 10.47.1.12.
```

The LISP Control-Plane responds with a Proxy Reply based on the address-resolution entry stored in its local database

<#root>

```
Border-1#
```

show lisp instance-id 8190 ethernet server address-resolution 10.47.4.3

```
Address-resolution data for router lisp 0 instance-id 8190

Site name: site_uci

Host Address:
```

10.47.4.3

```
/32
Hardware Address:
```

5254.001e.ad00

```
First registered:     21:11:17
Last registered:      21:11:17
Registration errors:
  Authentication failures:   0
 ETR 10.47.1.13:16056
  Last registered:       21:11:17
  TTL:                   1d00h
  xTR-ID:                0x8CEE6478-0x9358E248-0xE935FF07-0x8C3C5450
  Site-ID:               unspecified
  Registered addr:
```

5254.001e.ad00

```
  L3 Instance ID:
```

4099

Step 4. Edge Node receives LISP reply with MAC address of 10.47.4.3

LISP Proxy Reply is received by the Fabric Edge Node:

```
LISP[REMT ]-0: Processing Map-Reply mapping record for IID 8190 MAC 5254.001e.ad00/48 LCAF 2, ttl 1440,
LISP[REMT ]-0: Processing mapping information for EID prefix IID 8190 5254.001e.ad00/48.
```

Step 5. Edge Node sends LISP Map-Request packet to determine RLOC location for MAC address

After successful completion of first three steps, Edge Node knows MAC address of 10.47.4.3 for which ARP was initially generated for. As Layer-2 extension is enabled, Edge Node does not reply back with this information to 10.47.4.2, but rather use it to determine RLOC location of Egress Node Edge, so that it can forward ARP towards 10.47.4.3 as in a traditional Layer-2 network.

For this reason, Edge Node generates yet another LISP Map Request packet in Ethernet Instance, this time requesting RLOC information for 10.47.4.2's MAC address:

```
<#root>

Edge-1#

debug lisp control-plane all


Edge-1#

debug l2lisp all


*Oct 10 17:01:41.430: LISP[REMT ]-0 IID 8190: Schedule processing of Map-Requests from 'remote EID pref
*Oct 10 17:01:41.430: LISP[REMT ]-0: Map Request: Sending request for IID 8190 EID 5254.001e.ad00/48, r
```

Step 6: LISP Map-Request packet is received by Control Node to determine RLOC location for MAC address

Control Node receives the LISP packet and reply to it based on its local database state

```
<#root>

Border-1#

debug lisp control-plane all

Border-1#

debug l2lisp all

*Oct 10 16:04:42.055: LISP[MR   ]-0 IID 8190 Eth-ARP: MS EID 10.47.4.3/32: Sending proxy reply to 10.47
*Oct 10 16:04:42.407: LISP[MR   ]-0: Received Map-Request with 1 records, first EID IID 8190 5254.001e.
*Oct 10 16:04:42.408: LISP[MR   ]-0 IID 8190 MAC: MS EID 5254.001e.ad00/48: Sending proxy reply to 10.4
```

Step 7: LISP Map-Reply is received by Edge Node

LISP Map Reply generated by Control Node is received by Edge Node:

```
<#root>

Edge-1#

debug lisp control-plane all


Edge-1#

debug l2lisp all
```

```
*Oct 10 17:44:00.181: LISP[TRNSP]-0: Processing received Map-Reply(2) message on GigabitEthernet1/0/2 f
*Oct 10 17:44:00.181: LISP[REMT ]-0: Received Map-Reply with nonce 0xF954EC80-0x039D7E4A, 1 records.
*Oct 10 17:44:00.181: LISP[REMT ]-0: Map-Reply nonce matches pending request for IID 8190 EID 5254.001e
*Oct 10 17:44:00.181: LISP[REMT ]-0: Processing Map-Reply mapping record for IID 8190 MAC 5254.001e.ad0
*Oct 10 17:44:00.181: LISP[REMT ]-0: Map Request: Received reply with rtt 560ms.
*Oct 10 17:44:00.181: LISP[REMT ]-0: Processing mapping information for EID prefix IID 8190 5254.001e.a
```

This finally creates an entry in LISP Ethernet instance map-cache and allows ARP packet to be forwarded towards Edge-2 where 10.47.4.3 is connected to

<#root>

```
Edge-1#
```

`show lisp instance-id 8190 ethernet map-cache 5254.001e.ad00`

```
LISP MAC Mapping Cache for LISP 0 EID-table Vlan 1026 (IID 8190), 1 entries

5254.001e.ad00/48, uptime: 00:04:11, expires: 23:55:48, via map-reply, complete
  Sources: map-reply
  State: complete, last modified: 00:04:11, map-source: 10.47.1.13
  Active, Packets out: 8(0 bytes), counters are not accurate (~ 00:00:04 ago)
  Encapsulating dynamic-EID traffic
  Locator     Uptime    State  Pri/Wgt     Encap-IID
  10.47.1.13  00:04:11  up      10/10        -
    Last up-down state change:         00:04:11, state change count: 1
    Last route reachability change:    00:04:11, state change count: 1
    Last priority / weight change:     never/never
    RLOC-probing loc-status algorithm:
      Last RLOC-probe sent:            00:04:11 (rtt 560ms)
```

Step 8. ARP is encapsulated in VXLAN and send towards HOST-02

All of the LISP related steps were required to determine where 10.47.4.3 is located, so that Edge Node can sent original ARP (broadcast) packet as unicast towards proper Edge Node. The original ARP request is cached (not dropped) by Edge Node CPU until all of the steps are completed allowing proper ARP resolution even when single ARP packet was sent from 10.47.4.2.

ARP packet is encapsulated in VXLAN, as seen in the example:

<#root>

```
Edge-2#
```

`show monitor capture 1 buffer display-filter arp brief`

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
   67  15.149181 52:54:00:19:93:e9 -> 52:54:00:1e:ad:00 ARP 110 Who has 10.47.4.3? Tell 10.47.4.2
   68  15.155511 52:54:00:19:93:e9 -> 52:54:00:1e:ad:00 ARP 110 Who has 10.47.4.3? Tell 10.47.4.2
```

The ARP request has been encapsulated in VXLAN and also been converted from a broadcast ARP request to a unicast ARP request.

<#root>

```
Frame 68: 110 bytes on wire (880 bits), 110 bytes captured (880 bits) on interface /tmp/epc_ws/wif_to_t
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 10, 2023 17:56:43.256570000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1696960603.256570000 seconds
    [Time delta from previous captured frame: 0.006330000 seconds]
    [Time delta from previous displayed frame: 0.006330000 seconds]
    [Time since reference or first frame: 15.155511000 seconds]
    Frame Number: 68
    Frame Length: 110 bytes (880 bits)
    Capture Length: 110 bytes (880 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:udp:vxlan:eth:ethertype:arp]
Ethernet II, Src: 52:54:00:0a:42:11 (52:54:00:0a:42:11), Dst: 52:54:00:17:fe:65 (52:54:00:17:fe:65)
    Destination: 52:54:00:17:fe:65 (52:54:00:17:fe:65)
        Address: 52:54:00:17:fe:65 (52:54:00:17:fe:65)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 52:54:00:0a:42:11 (52:54:00:0a:42:11)
        Address: 52:54:00:0a:42:11 (52:54:00:0a:42:11)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.1.12

```
, Dst:
```

10.47.1.13 <-- 10.47.1.12 is Edge-1 RLOC, 10.47.1.13 is Edge-2 RLOC

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 96
    Identification: 0x1781 (6017)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 253
    Protocol: UDP (17)
    Header checksum: 0x4f95 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.1.12
    Destination: 10.47.1.13
User Datagram Protocol, Src Port: 65354, Dst Port: 4789
    Source Port: 65354
    Destination Port: 4789
    Length: 76
    [Checksum: [missing]]
    [Checksum Status: Not present]
    [Stream index: 0]
    [Timestamps]
        [Time since first frame: 15.155511000 seconds]
        [Time since previous frame: 0.006330000 seconds]
```

```
Virtual eXtensible Local Area Network
    Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
        1... .... .... .... = GBP Extension: Defined
        .... .... .0.. .... = Don't Learn: False
        .... 1... .... .... = VXLAN Network ID (VNI): True
        .... .... .... 0... = Policy Applied: False
        .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 0

    VXLAN Network Identifier (VNI): 8190 <-- L2 LISP IID

    Reserved: 0
Ethernet II, Src:

52:54:00:19:93:e9

 (52:54:00:19:93:e9), Dst:

52:54:00:1e:ad:00

 (52:54:00:1e:ad:00)

<--Unicast ARP Request

    Destination: 52:54:00:1e:ad:00 (52:54:00:1e:ad:00)
        Address: 52:54:00:1e:ad:00 (52:54:00:1e:ad:00)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        Address: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: ARP (

0x0806

)
    Trailer: 000000000000000000000000000000000000
Address Resolution Protocol (

request

)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
    Sender IP address: 10.47.4.2
    Target MAC address: 00:00:00:00:00:00 (00:00:00:00:00:00)
    Target IP address: 10.47.4.3
```

Step 9. ARP Reply is generated by 10.47.4.3 and sent towards 10.47.4.2

<#root>

```
Edge-2#

show monitor capture 1 buffer display-filter arp brief

Starting the packet display ........ Press Ctrl + Shift + 6 to exit
```

```
    1   0.000000 52:54:00:1e:ad:00 -> 52:54:00:19:93:e9 ARP 60 10.47.4.3 is at 52:54:00:1e:ad:00
    2   0.069429 52:54:00:1e:ad:00 -> 52:54:00:19:93:e9 ARP 60 10.47.4.3 is at 52:54:00:1e:ad:00
   11   5.960508 52:54:00:1e:ad:00 -> 52:54:00:19:93:e9 ARP 60 Who has 10.47.4.2? Tell 10.47.4.3
```

At this time packet is destined not to broadcast address (as original ARP request) but to 10.47.4.2's MAC address, when it reaches ingress Edge Node (Edge-2), normal LISP operation is triggered. Initially MAC address of 10.47.4.2 is missing in LISP Ethernet instance of Edge Node, packet is punted to CPU to generate LISP Map Request to determine RLOC for HOST-01. This behavior is exactly the same as described in other sections in this document and allow to create LISP Map Cache entry for 10.47.4.2 on Edge-2:

<#root>

```
Edge-2#
```

show lisp instance-id 8190 ethernet map-cache 5254.0019.93e9

```
LISP MAC Mapping Cache for LISP 0 EID-table Vlan 1026 (IID 8190), 1 entries

5254.0019.93e9/48, uptime: 03:18:28, expires: 20:41:32, via map-reply, complete
  Sources: map-reply
  State: complete, last modified: 03:18:28, map-source: 10.47.1.12
  Active, Packets out: 386(0 bytes), counters are not accurate (~ 00:00:12 ago)
  Encapsulating dynamic-EID traffic
  Locator     Uptime    State  Pri/Wgt     Encap-IID
```

10.47.1.12

```
  03:18:28  up      10/10        -
    Last up-down state change:         03:18:28, state change count: 1
    Last route reachability change:    03:18:28, state change count: 1
    Last priority / weight change:     never/never
    RLOC-probing loc-status algorithm:
      Last RLOC-probe sent:            03:18:28 (rtt 710ms)
```

The entry allows ARP Reply to be successfully sent towards Edge-1 in VXLAN Encapsulation and forwarded further to 10.47.4.2 competing whole ARP resolution process.

Basic Host Reachability in SDA Fabric (Same VLAN / Same VN)

It is assumed that ARP resolution completed successfully and both hosts 10.47.4.2 and 10.47.4.3 have proper ARP entries for each other.

From troubleshooting perspective, it is very useful to configure static ARP entries on both hosts to quickly check if the issue is with generic connectivity in the fabric (in such case ping does not work between hosts) or only with ARP process.

10.47.4.2 generates an ICMP Request towards 10.47.4.3:

<#root>

```
Edge-1#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    1   0.000000    10.47.4.2 -> 10.47.4.3    ICMP 98 Echo (ping) request  id=0x0040, seq=3/768, ttl=64

Edge-1#
```

show monitor capture 1 buffer detail

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 98 bytes on wire (784 bits), 98 bytes captured (784 bits) on interface /tmp/epc_ws/wif_to_ts_p
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 10, 2023 18:21:21.484694000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1696962081.484694000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 98 bytes (784 bits)
    Capture Length: 98 bytes (784 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:icmp:data]
Ethernet II, Src:
```

52:54:00:19:93:e9

```
 (52:54:00:19:93:e9), Dst:
```

52:54:00:1e:ad:00

```
 (52:54:00:1e:ad:00)
```

<-- Endpoint MAC, Anycast GW MAC respectively

```
    Destination: 52:54:00:1e:ad:00 (52:54:00:1e:ad:00)
        Address: 52:54:00:1e:ad:00 (52:54:00:1e:ad:00)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        Address: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.4.2

```
, Dst:
```

10.47.4.3

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 84
    Identification: 0x7321 (29473)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
```

```
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 64
    Protocol: ICMP (1)
    Header checksum: 0xab25 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.4.2
    Destination: 10.47.4.3
Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x02ea [correct]
    [Checksum Status: Good]
    Identifier (BE): 64 (0x0040)
    Identifier (LE): 16384 (0x4000)
    Sequence number (BE): 3 (0x0003)
    Sequence number (LE): 768 (0x0300)
    Data (56 bytes)

0000  68 95 8c 3d 00 00 00 00 00 00 00 00 00 00 00 00   h..=............
0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00                           ........
        Data: 68958c3d0000000000000000000000000000000000000000b^@&
        [Length: 56]
```

ICMP Packet is sent towards 10.47.4.3 to the Edge Node specified in Locator field: 10.47.1.13 (Edge-2) and can be captured via embedded packet capture.

As packet is received in VLAN where L2-extension is enabled, the lookup is done in LISP Ethernet instance:

<#root>

```
Edge-1#
```

show lisp instance-id 8190 ethernet map-cache 5254.001e.ad00

```
LISP MAC Mapping Cache for LISP 0 EID-table Vlan 1026 (IID 8190), 1 entries

5254.001e.ad00/48, uptime: 00:22:29, expires: 23:37:32, via map-reply, complete
  Sources: map-reply
  State: complete, last modified: 00:22:29, map-source: 10.47.1.13
  Active, Packets out: 42(0 bytes), counters are not accurate (~ 00:00:58 ago)
  Encapsulating dynamic-EID traffic
  Locator     Uptime    State  Pri/Wgt     Encap-IID
```

10.47.1.13

```
  00:22:29  up      10/10        -
    Last up-down state change:         00:22:29, state change count: 1
    Last route reachability change:    00:22:29, state change count: 1
    Last priority / weight change:     never/never
    RLOC-probing loc-status algorithm:
      Last RLOC-probe sent:            00:22:28 (rtt 1609ms)
```

Check the MAC address of the remote endpoint, it points the L2LI0, which is expected

```
<#root>

Edge-1#

show mac add add 5254.001e.ad00

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
1026    5254.001e.ad00    CP_LEARN    L2LI0
Total Mac Addresses installed by LISP: REMOTE: 1
```

Check the MAC address in FED, additional information can be gleaned

```
<#root>

Edge-1#

show platform software fed sw active matm macTable vlan 1026 mac 5254.001e.ad00

VLAN   MAC                  Type  Seq#    EC_Bi  Flags

machandle


siHandle


riHandle

            diHandle              *a_time  *e_time  ports
-------------------------------------------------------------------------------------------------------

1026


5254.001e.ad00

  0x1000001       0      0     64

0x7f65ecfdd3a8


0x7f65ecfdd1f8


0x7f65ecfdd048

      0x0                         0        2  RLOC 10.47.1.13 adj_id 97

======platform hardware details ======

Asic: 0
```

```
 htm-handle = 0x7f65ecc4d188 MVID = 7 gpn = 1
 SI = 0xc7 RI = 0x12 DI = 0x5012
Asic: 1
 SI = 0xc7 RI = 0x12 DI = 0x5013
```

MAC Address macHandle Decode

Take the macHandle (**0x7f65ecfdd3a8)** from the previous command and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <macHandle> 1**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ecfdd3a8 1

```
Handle:0x7f65ecfdd3a8 Res-Type:ASIC_RSC_HASH_TCAM Res-Switch-Num:0 Asic-Num:255 Feature-ID:AL_FID_L2_WI
priv_ri/priv_si Handle: (nil)Hardware Indices/Handles: handle [ASIC: 0]: 0x7f65ecc4d188
Features sharing this resource:Cookie length: 12
1e 00 54 52 00 ad 07 80 07 00 00 00


Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
Number of HTM Entries: 1

Entry 0: (handle 0x7f65ecc4d188)

Absolute Index: 4706
Time Stamp: 14
KEY -
```

vlan:7

mac:0x5254001ead00

```
 l3_if:0
```

gpn:3401

```
 epoch:0 static:0 flood_en:0 vlan_lead_wless_flood_en: 0 client_home_asic: 0 learning_peerid 0, learnin
MASK - vlan:0 mac:0x0 l3_if:0 gpn:0 epoch:0 static:0 flood_en:0 vlan_lead_wless_flood_en: 0 client_home
SRC_AD - need_to_learn:0 lrn_v:0 catchall:0 static_mac:0 chain_ptr_v:0 chain_ptr: 0 static_entry_v:0 au
DST_AD - si:0xc7 bridge:0 replicate:0 blk_fwd_o:0 v4_rmac:0 v6_rmac:0 catchall:0 ign_src_lrn:0 port_mas


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

MAC Address siHandle Decode

Take the siHandle (**0x7f65ecfdd1f8**) from the previous command and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <siHandle> 1**

<#root>

```
Edge-1#
```

```
show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ecfdd1f8 1
```

```
Handle:0x7f65ecfdd1f8 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L2_WIRELES
priv_ri/priv_si Handle: 0x7f65ecfdd048Hardware Indices/Handles: index0:0xc7  mtu_index/l3u_ri_index0:0x
Features sharing this resource:58 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 02 04 00 00 00 00 00 00 00 00 00 00 07 00 61 00 00 00 00 00 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------

Station Index (SI) [0xc7] <-- Contains the Rewrite Index (RI) and Outgoing Interface Information (DI)
RI = 0x12 <-- Rewrite Index = Contains information for forwarding
DI = 0x5012 <-- Destination Index = Outgoing Interface

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD

Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------

Station Index (SI) [0xc7] <-- Contains the Rewrite Index (RI) and Outgoing Interface Information (DI)
RI = 0x12 <-- Rewrite Index = Contains information for forwarding
DI = 0x5013 <-- Destination Index = Outgoing Interface

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD

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

Rewrite Index Decode

Take the RI (**0x12**) and use in the command **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

<#root>

```
Edge-1#
```

```
show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x12 0x12
```

```
ASIC#:0 RI:18 Rewrite_type:AL_RRM_REWRITE_L2_PAYLOAD_L2LISP_ENCAP(115) Mapped_rii:LVX_L2_ENCAP_L2_PAYLO
 Src IP:
```

```
10.47.1.12 <-- Local RLOC
```

```
 Dst IP:
```

```
10.47.1.13 <-- Remote RLOC

 iVxlan dstMac:  0x5254:0x01c:0x7de0
 iVxlan srcMac:  0x00:0x00:0x00
 IPv4 TTL:       0
 iid present:    1
 lisp iid:       0
 lisp flags:     0
 dst Port:       4789
 update only l3if:       0
 is Sgt:         1
 is TTL Prop:    0
 L3if LE:        0 (0)
 Port LE:        0 (0)
 Vlan LE:        7 (0)

ASIC#:1 RI:18 Rewrite_type:AL_RRM_REWRITE_L2_PAYLOAD_L2LISP_ENCAP(115) Mapped_rii:LVX_L2_ENCAP_L2_PAYLO
 Src IP:

10.47.1.12 <-- Local RLOC

 Dst IP:

10.47.1.13 <-- Remote RLOC

 iVxlan dstMac:  0x5254:0x01c:0x7de0
 iVxlan srcMac:  0x00:0x00:0x00
 IPv4 TTL:       0
 iid present:    1
 lisp iid:       0
 lisp flags:     0
 dst Port:       4789
 update only l3if:       0
 is Sgt:         1
 is TTL Prop:    0
 L3if LE:        0 (0)
 Port LE:        0 (0)
 Vlan LE:        7 (0)
```

Destination-Index Decode

Take the DI (**0x5012**) and use in the command **show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>**

<#root>

```
Edge-1#

show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x5012 0x501

ASIC#0:

Destination index   = 0x5012

DI_RCP_PORT1 <-- Recirculation port for VXLAN imposition

pmap                = 0x00000000 0x00000000
cmi                 = 0x0
rcp_pmap            = 0x1
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
ASIC#1:

Destination index   = 0x5012
```

DI_RCP_PORT1 <-- Recirculation port for VXLAN imposition

```
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
```

siHandle Decode

Take the siHandle (**0x7f65ecfdd048)** and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <siHandle> 1**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ecfdd048 1

```
Handle:0x7f65ecfdd048 Res-Type:ASIC_RSC_RI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L2_WIRELES
priv_ri/priv_si Handle: 0x7f65ecfdcd78Hardware Indices/Handles: index0:0x12  mtu_index/l3u_ri_index0:0x
Features sharing this resource:58 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 02 04 00 00 00 00 00 00 00 00 00 00 07 00 61 00 00 00 00 00 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
ASIC#:0 RI:18 Rewrite_type:AL_RRM_REWRITE_L2_PAYLOAD_L2LISP_ENCAP(115) Mapped_rii:LVX_L2_ENCAP_L2_PAYLO
 Src IP:
```

10.47.1.12 <-- Local RLOC

```
 Dst IP:
```

10.47.1.13 <-- Remote RLOC

```
 iVxlan dstMac:     0x610:0x00:0x00
 iVxlan srcMac:     0x00:0x00:0x00
```

```
 IPv4 TTL:     0
 iid present:     1
 lisp iid:     0
 lisp flags:     0
 dst Port:     4789
 update only l3if:     0
 is Sgt:     1
 is TTL Prop:     0
 L3if LE:     0 (0)
 Port LE:     279 (0)
 Vlan LE:     7 (0)


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
ASIC#:1 RI:18 Rewrite_type:AL_RRM_REWRITE_L2_PAYLOAD_L2LISP_ENCAP(115) Mapped_rii:LVX_L2_ENCAP_L2_PAYLO
Src IP:

10.47.1.12 <-- Local RLOC


Dst IP:

10.47.1.13 <-- Remote RLOC

 iVxlan dstMac:     0x610:0x00:0x00
 iVxlan srcMac:     0x00:0x00:0x00
 IPv4 TTL:     0
 iid present:     1
 lisp iid:     0
 lisp flags:     0
 dst Port:     4789
 update only l3if:     0
 is Sgt:     1
 is TTL Prop:     0
 L3if LE:     0 (0)
 Port LE:     279 (0)
 Vlan LE:     7 (0)

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

Underlay Route Verification

Traffic is encapsulated in VXLAN with IID 8190 using 10.47.1.12 and has the ability to be load-balanced out Gig1/0/1 and G1/0/2

```
<#root>

Edge-1#

show ip route 10.47.1.13


Routing entry for 10.47.1.13/32
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 10.47.1.4 on GigabitEthernet1/0/2, 2d22h ago
  Routing Descriptor Blocks:
    10.47.1.4, from 10.47.1.13, 2d22h ago, via GigabitEthernet1/0/2
      Route metric is 30, traffic share count is 1
```

```
  * 10.47.1.0, from 10.47.1.13, 2d22h ago, via GigabitEthernet1/0/1
      Route metric is 30, traffic share count is 1

Edge-1#
```

`show ip cef 10.47.1.13`

```
10.47.1.13/32
  nexthop 10.47.1.0 GigabitEthernet1/0/1
  nexthop 10.47.1.4 GigabitEthernet1/0/2
```

To get si_hdl, ri_hdl information, use the command **show platform software fed switch active ip adj**

<#root>

```
Edge-1#
```

`show platform software fed switch active ip adj`

```
IPV4 Adj entries
dest                                    if_name                dst_mac          si_hdl          r
----                                    -------                -------          ------          -
225.0.0.0                               GigabitEthernet1/0/1   0100.5e00.0000   0x7f65ec958128 0
10.47.1.10                              LISP0.4100             4500.0000.0000   0x7f65ec895ed8 0
225.0.0.0                               GigabitEthernet1/0/2   0100.5e00.0000   0x7f65ec958f68 0

10.47.1.4                               GigabitEthernet1/0/2   5254.001c.7de0   0x7f65ec8a5458 0x

225.0.0.0                               Null0                  f800.0011.0000   0x7f65ec3740c8 0

10.47.1.0                               GigabitEthernet1/0/1   5254.000a.42f3   0x7f65ec8b8468 0x
```

Underlay Next-Hop si_hdl Decode

To check the si_hdl (**0x7f65ec8a5458)** use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_hdl> 1**

<#root>

```
Edge-1#
```

`show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec8a5458 1`

```
Handle:0x7f65ec8a5458 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f65ec8a4eb8Hardware Indices/Handles: index0:0xbc  mtu_index/l3u_ri_index0:0x0
Features sharing this resource:66 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 26 00 00 00 00 00 00 00 00 00 00 00 08 00 52 54 00 1c 7d e0 00 00 00 00 00 00 0
Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
```

`Station Index (SI) [0xbc]     -----> Contains RI and DI information`
`RI = 0x1a             -----> Rewrite Index = MAC address rewrite information for L3 forwarding to the ne`
`DI = 0x526d            -----> Destination Index = Outgoing Interface`

```
stationTableGenericLabel = 0
```

```
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
```

Replication Bitmap: LD ----> Local Data, indicating that this ASIC is directly connected to the

```
Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
```

Station Index (SI) [0xbc] -----> Contains RI and DI information
RI = 0x1a -----> Rewrite Index = MAC address rewrite information for L3 forwarding to the ne
DI = 0x526d -----> Destination Index = Outgoing Interface

```
stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
```

Replication Bitmap: CD -----> Core Data, indicating that this instance of the ASIC is on the same

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

Underlay Next-Hop Rewrite Index Decode

To decode the RI (**0x1a)** use in the command **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x1a 0x1a

```
ASIC#:0
```

RI:26

```
 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
```

-----> Decimal 26 is hex 0x1a

```
MAC Addr: MAC Addr: 52:54:00:1c:7d:e0,
```

-----> MAC address 5254.001c.7de0 for the next-hop adjacency

```
L3IF LE Index 38
ASIC#:1 RI:26 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
MAC Addr: MAC Addr: 52:54:00:1c:7d:e0,
```

-----> MAC address 5254.001c.7de0 for the next-hop adjacency

```
L3IF LE Index 38
```

Underlay Next-Hop Destination Index Decode

To decode the DI (**0x526d)** use in **show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>**

<#root>

Edge-1#

show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x526d 0x526

ASIC#0:

Destination index   = 0x526d

pmap   = 0x00000000 0x00000002 <-- Convert decimal 2 to binary, which is 0010. Count this b

pmap_intf : [GigabitEthernet1/0/2]
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
ASIC#1:

Destination index   = 0x526d
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

Edge-1#

show platform software fed switch active ifm mappings

Interface   IF_ID   Inst Asic Core

Port

 SubPort Mac Cntx LPN GPN Type Active
GigabitEthernet1/0/1   0x1a   0   0   0   0   0   1   0   1   1   NIF Y

GigabitEthernet1/0/2

 0x1b   0   0   0

1

     0   2   1   2   2   NIF Y

<-- Port 1 lines up to G1/0/2

GigabitEthernet1/0/3   0xb   0   0   0   2   0   3   2   3   3   NIF Y

```
GigabitEthernet1/0/4      0xc         0   0   0    3      0      4    3    4    4    NIF  Y
GigabitEthernet1/0/5      0xd         0   0   0    4      0      5    4    5    5    NIF  Y
GigabitEthernet1/0/6      0xe         0   0   0    5      0      6    5    6    6    NIF  Y
GigabitEthernet1/0/7      0xf         0   0   0    6      0      7    6    7    7    NIF  Y
GigabitEthernet1/0/8      0x10        0   0   0    7      0      8    7    8    8    NIF  Y
```

Underlay Next-Hop ri_hdl Decode

To decode the ri_hdl (**0x7f65ec8a4eb8**) use in **show platform hardware fed switch active fwd-asic abstraction print-resource-handle (ri_hdl) 1**

<#root>

```
Edge-1#
show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec8a4eb8 1

Handle:0x7f65ec8a4eb8 Res-Type:ASIC_RSC_RI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f65ec903b28Hardware Indices/Handles: index0:0x1a  mtu_index/l3u_ri_index0:0x0
Features sharing this resource:66 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 26 00 00 00 00 00 00 00 00 00 00 00 08 00 52 54 00 1c 7d e0 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
ASIC#:0
RI:26
 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
<-- Decimal 26 is 0x1a in hex
    MAC Addr: MAC Addr:
52:54:00:1c:7d:e0
,
<-- MAC address 5254.001c.7de0 for the next-hop adjacency
    L3IF LE Index 38


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
ASIC#:1
RI:26
 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
<-- Decimal 26 is 0x1a in hex
 MAC Addr: MAC Addr:
52:54:00:1c:7d:e0
,
    MAC Addr: MAC Addr:
```

```
52:54:00:1c:7d:e0

,

<-- MAC address 5254.001c.7de0 for the next-hop adjacency


    L3IF LE Index 38


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

Basic Host Reachability in SDA Fabric (Different VLANs / Same VN)

In this section, communication between 10.47.4.2 and 10.47.10.2 is examined. As these hosts belong to different VLANs, both need to have default-gateway configured that points to default gateway. For 10.47.4.2 it is 10.47.4.1 and 10.47.10.2 it is 10.47.10.1.

Step 1. Confirm that connectivity between the endpoint and the default gateway works:

<#root>

```
Edge-1#

ping vrf red_vn 10.47.4.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.47.4.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 155/164/181 ms
```

<#root>

```
Edge-2#

ping vrf red_vn 10.47.10.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.47.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 41/46/62 ms
```

Step 2. Confirm that packet from 10.47.4.2 is received successfully by Edge-1:

Packet can be captured on the ingress interface facing 10.47.4.2:

<#root>

```
Edge-1#

monitor capture 1 interface g1/0/3 in match any

Edge-1#
```

mon cap 1 start

```
Started capture point : 1
Edge-1#
```

mon cap 1 stop

```
Capture statistics collected at software:
    Capture duration - 12 seconds
    Packets received - 9
    Packets dropped - 0
    Packets oversized - 0

Number of Bytes dropped at asic not collected

Capture buffer will exists till exported or cleared

Stopped capture point : 1

Edge-1#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    1   0.000000    10.47.4.2 -> 10.47.10.2   ICMP 98 Echo (ping) request  id=0x0041, seq=0/0, ttl=64
    2   0.023447    10.47.4.2 -> 10.47.10.2   ICMP 98 Echo (ping) request  id=0x0041, seq=0/0, ttl=64

Edge-1#
```

show monitor capture 1 buffer detailed

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 98 bytes on wire (784 bits), 98 bytes captured (784 bits) on interface /tmp/epc_ws/wif_to_ts_p
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 11, 2023 15:27:46.033825000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1697038066.033825000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 98 bytes (784 bits)
    Capture Length: 98 bytes (784 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:icmp:data]
Ethernet II, Src: 52:54:00:19:93:e9 (
```

52:54:00:19:93:e9

```
), Dst: 00:00:0c:9f:f3:41 (
```

00:00:0c:9f:f3:41

```
)
```

<-- SMAC and DMAC respectively

```
    Destination: 00:00:0c:9f:f3:41 (00:00:0c:9f:f3:41)
        Address: 00:00:0c:9f:f3:41 (00:00:0c:9f:f3:41)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
```

```
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        Address: 52:54:00:19:93:e9 (52:54:00:19:93:e9)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.4.2

```
, Dst:
```

10.47.10.2

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 84
    Identification: 0x395e (14686)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 64
    Protocol: ICMP (1)
    Header checksum: 0xdee9 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.4.2
    Destination: 10.47.10.2
Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x248a [correct]
    [Checksum Status: Good]
    Identifier (BE): 65 (0x0041)
    Identifier (LE): 16640 (0x4100)
    Sequence number (BE): 0 (0x0000)
    Sequence number (LE): 0 (0x0000)
    Data (56 bytes)

0000  2a 46 a8 ee 00 00 00 00 00 00 00 00 00 00 00 00   *F..............
0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00                           ........
        Data: 2a46a8ee0000000000000000000000000000000000000000b^@&
        [Length: 56]
```

Step 3 - LISP Lookup

Ingress Edge Node must determine the location (RLOC) of HOST-03 to which it sends packets to. As in this case the end-host HOST-03 is in different VLAN (but same VN / VRF: USERS), LISP IPv4 instance is used as the lookup is based on IP address (MAC address belong to Edge Node itself).

<#root>

```
Edge-1#
```

```
debug lisp control-plane all

LISP[REMT ]-0: Map Request: Sending request for IID 4099 EID 10.47.10.2/32, requester 'remote EID prefi
LISP[REMT ]-0: Map-Reply nonce matches pending request for IID 4099 EID 10.47.10.2/32, requester 'remot
```

LISP Map Request reaches Control Node (LISP Map Server) Border-1:

```
<#root>

Border-1#

debug lisp control-plane all

LISP[TRNSP]-0: Processing received Map-Request(1) message on GigabitEthernet1/0/3 from 10.47.10.2:4342
LISP[MR   ]-0: Received Map-Request with 1 records, first EID IID 4099 10.47.10.2/32, source EID 10.47.
LISP[MR   ]-0 IID 4099 IPv4: MS EID 10.47.10.2/32: Sending proxy reply to 10.47.1.12.
```

LISP Map-Reply reaches Edge Node:

```
LISP[REMT ]-0: Processing Map-Reply mapping record for IID 4099 IPv4 10.47.10.2/32 LCAF 2, ttl 1440, ac
LISP[REMT ]-0: Processing mapping information for EID prefix IID 4099 10.47.10.2/32.
```

The Fabric Edge queries for the RLOC for 10.47.10.2 and process the Map-Reply

```
LISP[REMT ]-0: Map Request: Sending request for IID 4099 EID 10.47.10.2/32, requester 'remote EID RLOC'
LISP[REMT ]-0: Processing Map-Reply mapping record for IID 4099 IPv4 10.47.10.2/32 LCAF 2, ttl 1440, ac
LISP[REMT ]-0: Processing mapping information for EID prefix IID 4099 10.47.10.2/32.
```

In case entry does not exist, debugs needs to be collected from LISP process perspective. There is also a tool, called LIG (LISP Grouper) which can be used to trigger manually LISP process (this is very effective way to test redundant Control Node configuration and database consistency between both Control Nodes):

```
<#root>

Edge-1#

lig instance-id 4099 10.47.10.2 to 10.47.1.10

Mapping information for EID 10.47.10.2 from 10.47.1.10 with RTT 334 msecs
10.47.10.2/32, uptime: 00:00:00, expires: 23:59:59, via map-reply, complete
  Locator     Uptime    State  Pri/Wgt     Encap-IID
  10.47.1.13  00:00:00  up      10/10        -

Edge-1#

lig instance-id 4099 10.47.10.2 to 10.47.1.11

Mapping information for EID 10.47.10.2 from 10.47.1.11 with RTT 327 msecs
10.47.10.2/32, uptime: 00:00:06, expires: 23:59:59, via map-reply, complete
  Locator     Uptime    State  Pri/Wgt     Encap-IID
```

```
  10.47.1.13  00:00:06  up      10/10        -
```

Route Verification

CEF uses LISP, and LISP utilizes the map-cache entry it has received

```
<#root>

Edge-1#

show ip cef vrf red_vn 10.47.10.2

10.47.10.2/32
  nexthop 10.47.1.13 LISP0.4099

Edge-1#

show ip route 10.47.1.13

Routing entry for 10.47.1.13/32
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 10.47.1.4 on GigabitEthernet1/0/2, 3d19h ago
  Routing Descriptor Blocks:
    10.47.1.4, from 10.47.1.13, 3d19h ago, via GigabitEthernet1/0/2
      Route metric is 30, traffic share count is 1
  * 10.47.1.0, from 10.47.1.13, 3d19h ago, via GigabitEthernet1/0/1
      Route metric is 30, traffic share count is 1

Edge-1#

show lisp instance-id 4099 ipv4 map-cache  10.47.10.2

LISP IPv4 Mapping Cache for LISP 0 EID-table vrf red_vn (IID 4099), 1 entries


10.47.10.2

/32, uptime: 00:08:48, expires: 23:51:17, via map-reply, complete
  Sources: map-reply
  State: complete, last modified: 00:08:48, map-source: 10.47.1.11
  Active, Packets out: 51(29376 bytes), counters are not accurate (~ 00:00:15 ago)
  Encapsulating dynamic-EID traffic
  Locator     Uptime    State  Pri/Wgt     Encap-IID


10.47.1.13

  00:08:48  up      10/10        -
    Last up-down state change:         00:08:48, state change count: 1
    Last route reachability change:    22:07:12, state change count: 1
    Last priority / weight change:     never/never
    RLOC-probing loc-status algorithm:
      Last RLOC-probe sent:            00:08:48 (rtt 931ms)
```

LISP Next-Hop Verification

Since this packet is VXLAN encapsulated, verification of the LISP next-hop needs to occur. Use the command **show platform software fed switch active ip adj** to get additional information about 10.47.1.13,

the LISP next-hop

```
<#root>

Edge-1#

show platform software fed switch active ip adj


IPV4 Adj entries
dest                                    if_name                dst_mac          si_hdl          r
----                                    -------                -------          ------          -
10.47.1.10                              LISP0.4100             4500.0000.0000   0x7f65ec895ed8 0
10.47.1.4                               GigabitEthernet1/0/2   5254.001c.7de0   0x7f65ec8a5458 0
10.47.1.0                               GigabitEthernet1/0/1   5254.000a.42f3   0x7f65ec8b8468 0
10.47.4.2                               Vlan1026               5254.0019.93e9   0x7f65ec7c21f8 0

10.47.1.13                              LISP0.4099             4500.0000.0000   0x7f65ed00f668 0x
```

LISP Next-Hop si_hdl Decode

Take the si_hdl (**0x7f65ed00f668)** and use in **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_hdl> 1**

```
<#root>

Edge-1#

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ed00f668 1

Handle:0x7f65ed00f668 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_LISP Lkp-f
priv_ri/priv_si Handle: 0x7f65ed00fd58Hardware Indices/Handles: index0:0xc8  mtu_index/l3u_ri_index0:0x
Features sharing this resource:109 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 38 5f 84 ec 0a 2f 01 0d ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00


Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------


Station Index (SI) [0xc8] <-- Contains the RI and DI
RI = 0x2c <-- Rewrite Index contains information for L3 Forwarding
DI = 0x5012 <-- Destination Index contains information for the destination port

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0xc
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------

Station Index (SI) [0xc8] <-- Contains the RI and DI
RI = 0x2c <-- Rewrite Index contains information for L3 Forwarding
```

```
DI = 0x5013 <-- Destination Index contains information for the destination port

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0xc
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD

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

LISP Next-Hop RI Decode

Take the RI (**0x2c**) and use in **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

<#root>

```
Edge-1#

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x2c 0x2c

ASIC#:0 RI:44 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_P
 Dst Mac:     MAC Addr: ba:25:cd:f4:ad:38,
 Src IP:

10.47.1.12 <-- Local RLOC

 Dst IP:

10.47.1.13 <-- RLOC of Edge-2

 IPv4 TTL:     0
 LISP INSTANCEID:     0
 L3IF LE Index:     46

ASIC#:1 RI:44 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_P
 Dst Mac:     MAC Addr: ba:25:cd:f4:ad:38,
 Src IP:

10.47.1.12 <-- Local RLOC


 Dst IP:

10.47.1.13 <-- RLOC of Edge-2


 IPv4 TTL:     0
 LISP INSTANCEID:     0
 L3IF LE Index:     46
```

LISP Next-Hop DI Decode

Take the DI (0x5012) and use in show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x5012 0x501

```
ASIC#0:

Destination index   = 0x5012
```

DI_RCP_PORT1 <-- Expected, this means the packet is recirculated for VXLAN imposition

```
pmap                = 0x00000000 0x00000000
cmi                 = 0x0
rcp_pmap            = 0x1
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
ASIC#1:

Destination index   = 0x5012
```

DI_RCP_PORT1 <-- Expected, this means the packet is recirculated for VXLAN imposition

```
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
```

LISP Next-Hop ri_hdl Decode

Take the ri_hdl (**0x7f65ed00fd58**) and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <ri_hdl> 1**

<#root>

```
Edge-1#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ed00fd58 1

```
Handle:0x7f65ed00fd58 Res-Type:ASIC_RSC_RI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_LISP Lkp-f
priv_ri/priv_si Handle: 0x7f65ed00b618Hardware Indices/Handles: index0:0x2c  mtu_index/l3u_ri_index0:0x
Features sharing this resource:109 (1)]
Cookie length: 56
```

```
00 00 00 00 00 00 00 00 2e 00 00 00 0a 2f 01 0d ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
ASIC#:0 RI:44 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_PA
 Dst Mac:     MAC Addr: ba:25:cd:f4:ad:38,
 Src IP:
10.47.1.12 <-- Local RLOC
 Dst IP:
10.47.1.13 <-- Edge-2 RLOC
 IPv4 TTL:     0
 LISP INSTANCEID:     0
 L3IF LE Index:     46

Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
ASIC#:1 RI:44 Rewrite_type:AL_RRM_REWRITE_IPV4_VXLAN_INNER_IPV4_ENCAP(110) Mapped_rii:LVX_L3_ENCAP_L2_PA
 Dst Mac:     MAC Addr: ba:25:cd:f4:ad:38,
 Src IP:
10.47.1.12 <-- Local RLOC
 Dst IP:
10.47.1.13 <-- Edge-2 RLOC
 IPv4 TTL:     0
 LISP INSTANCEID:     0
 L3IF LE Index:     46

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

Underlay Next-Hop Verification

To get to the LISP Next-Hop, there are two possible paths in the underlay, verification occurs for one, the same logic applies in verification of the other underlay next-hop.

```
<#root>

Edge-1#
show ip route 10.47.1.13

Routing entry for 10.47.1.13/32
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 10.47.1.4 on GigabitEthernet1/0/2, 3d19h ago
  Routing Descriptor Blocks:
10.47.1.4
, from 10.47.1.13, 3d19h ago, via GigabitEthernet1/0/2
      Route metric is 30, traffic share count is 1
  *
```

```
10.47.1.0

, from 10.47.1.13, 3d19h ago, via GigabitEthernet1/0/1
      Route metric is 30, traffic share count is 1
```

To get more information about the next-hops, use **show platform software fed switch active ip adj**

<#root>

```
Edge-1#

show platform software fed switch active ip adj

IPV4 Adj entries
dest                                    if_name                dst_mac          si_hdl         r
----                                    -------                -------          ------         -

10.47.1.4                               GigabitEthernet1/0/2   5254.001c.7de0   0x7f65ec8a5458 0x

10.47.1.0                               GigabitEthernet1/0/1   5254.000a.42f3   0x7f65ec8b8468 0x

<snip>
```

Underlay Next-Hop si_hdl Decode

Take the si_hdl (**0x7f65ec8a5458**) and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_hdl> 1**

<#root>

```
Edge-1#

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec8a5458 1

Handle:0x7f65ec8a5458 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f65ec8a4eb8Hardware Indices/Handles: index0:0xbc  mtu_index/l3u_ri_index0:0x0
Features sharing this resource:66 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 26 00 00 00 00 00 00 00 00 00 00 00 08 00 52 54 00 1c 7d e0 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------


Station Index (SI) [0xbc] <-- Contains the RI and DI
RI = 0x1a <-- Rewrite index contains information for L3 Forwarding
DI = 0x526d <-- Destination index contains information for the destination port

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: LD


Detailed Resource Information (ASIC_INSTANCE# 1)
```

```
---------------------------------------

Station Index (SI) [0xbc] <-- Contains the RI and DI
RI = 0x1a <-- Rewrite index contains information for L3 Forwarding
DI = 0x526d <-- Destination index contains information for the destination port

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0
Replication Bitmap: CD

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

Underlay Next-Hop RI Decode

Take the RI (**0x1a**) and use in the command **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

```
<#root>

Edge-1#

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x1a 0x1a

ASIC#:0

RI:26

 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)

<-- Decimal 26 is hex 0x1a

    MAC Addr: MAC Addr:

52:54:00:1c:7d:e0

,

<-- MAC Address 5254.001c.7de0 corresponds to the next-hop

    L3IF LE Index 38


ASIC#:1

RI:26

 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)

<-- Decimal 26 is hex 0x1a

    MAC Addr: MAC Addr:

 52:54:00:1c:7d:e0

,

<-- MAC Address 5254.001c.7de0 corresponds to the next-hop

    L3IF LE Index 38
```

Underlay Next Hop DI Decode

Take the DI (**0x526d**) and use in the command **show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>**

<#root>

```
Edge-1#

show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x526d 0x526

ASIC#0:

Destination index   = 0x526d

pmap                = 0x00000000 0x00000002 <-- Take decimal 2 and convert to binary, so 0010, and then

pmap_intf : [GigabitEthernet1/0/2]
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
ASIC#1:

Destination index   = 0x526d

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

Edge-1#

show platform software fed  switch active ifm mappings

Interface                IF_ID    Inst Asic Core

Port

 SubPort Mac  Cntx LPN  GPN  Type Active
GigabitEthernet1/0/1      0x1a       0   0   0    0      0      1    0    1    1    NIF  Y

GigabitEthernet1/0/2
```

```
 0x1b       0   0   0

1

      0      2    1    2    2    NIF  Y

<-- Port 1 maps to Gig1/0/2

GigabitEthernet1/0/3      0xb        0   0   0    2      0      3    2    3    3    NIF  Y
GigabitEthernet1/0/4      0xc        0   0   0    3      0      4    3    4    4    NIF  Y
GigabitEthernet1/0/5      0xd        0   0   0    4      0      5    4    5    5    NIF  Y
GigabitEthernet1/0/6      0xe        0   0   0    5      0      6    5    6    6    NIF  Y
GigabitEthernet1/0/7      0xf        0   0   0    6      0      7    6    7    7    NIF  Y
GigabitEthernet1/0/8      0x10       0   0   0    7      0      8    7    8    8    NIF  Y
```

Underlay Next-Hop ri_hdl Decode

Take the ri_hdl (**0x7f65ec8b8158**) and use in the command **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <ri_hdl> 1**

<#root>

```
Edge-1#

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f65ec8b8158 1

Handle:0x7f65ec8b8158 Res-Type:ASIC_RSC_RI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f65ec7a6338Hardware Indices/Handles: index0:0x1b  mtu_index/l3u_ri_index0:0x0
Features sharing this resource:66 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 25 00 00 00 00 00 00 00 00 00 00 00 08 00 52 54 00 0a 42 f3 00 00 00 00 00 00 0

Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------
ASIC#:0 RI:27 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
    MAC Addr: MAC Addr:

52:54:00:0a:42:f3

,
    L3IF LE Index 37


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------
ASIC#:1 RI:27 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
    MAC Addr: MAC Addr:

52:54:00:0a:42:f3

,
    L3IF LE Index 37


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

Packets are encapsulated in VXLAN and sent based on load-balancing rules. Embedded Packet Capture

(EPC) can be used to capture traffic on all interfaces at the same time. Remember at this point the packet is VXLAN encapsulated, the EPC filter must be against RLOC to RLOC, not inner IPv4 addresses.

```
<#root>

Edge-1#

monitor capture 1 interface range g1/0/1-2 out match ipv4 host 10.47.1.12 host 10.47.1.13

Edge-1#

monitor capture 1 start

Started capture point : 1
Edge-1#
Edge-1#

monitor capture 1 stop

Capture statistics collected at software:
    Capture duration - 18 seconds
    Packets received - 4
    Packets dropped - 0
    Packets oversized - 0

Number of Bytes dropped at asic not collected

Capture buffer will exists till exported or cleared

Stopped capture point : 1

Edge-1#

show monitor capture 1 buffer brief

Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    1   0.000000    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0046, seq=0/0, ttl=63
    2   0.980849    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0046, seq=1/256, ttl=6
    3   1.984077    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0046, seq=2/512, ttl=6
    4   2.999989    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0046, seq=3/768, ttl=6

Edge-1#

show monitor capture 1 buffer detailed


Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 148 bytes on wire (1184 bits), 148 bytes captured (1184 bits) on interface /tmp/epc_ws/wif_to_
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 11, 2023 16:50:52.262553000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1697043052.262553000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 148 bytes (1184 bits)
    Capture Length: 148 bytes (1184 bits)
    [Frame is marked: False]
```

```
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:udp:vxlan:eth:ethertype:ip:icmp:data]
Ethernet II, Src:
```

00:00:00:00:00:00

```
 (00:00:00:00:00:00), Dst:
```

00:00:00:00:00:00

```
 (00:00:00:00:00:00)
```

<-- EPC does not capture L3 rewrite on egress properly, this is OK

```
    Destination: 00:00:00:00:00:00 (00:00:00:00:00:00)
        Address: 00:00:00:00:00:00 (00:00:00:00:00:00)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 00:00:00:00:00:00 (00:00:00:00:00:00)
        Address: 00:00:00:00:00:00 (00:00:00:00:00:00)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.1.12

```
, Dst:
```

10.47.1.13 <-- RLOC to RLOC

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 134
    Identification: 0x1d6f (7535)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 64
    Protocol: UDP (17)
    Header checksum: 0x0682 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.1.12
    Destination: 10.47.1.13
User Datagram Protocol, Src Port: 65354, Dst Port: 4789
    Source Port: 65354
    Destination Port: 4789
    Length: 114
    [Checksum: [missing]]
    [Checksum Status: Not present]
    [Stream index: 0]
    [Timestamps]
        [Time since first frame: 0.000000000 seconds]
        [Time since previous frame: 0.000000000 seconds]
Virtual eXtensible Local Area Network
    Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
        1... .... .... .... = GBP Extension: Defined
        .... .... .0.. .... = Don't Learn: False
        .... 1... .... .... = VXLAN Network ID (VNI): True
        .... .... .... 0... = Policy Applied: False
```

```
        .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 0
    VXLAN Network Identifier (VNI):
```

4099 <-- LISP L3 IID

```
    Reserved: 0
Ethernet II, Src: 00:00:00:00:61:00 (
```

00:00:00:00:61:00

```
), Dst: ba:25:cd:f4:ad:38 (
```

ba:25:cd:f4:ad:38

```
)
```

<-- Dummy Ethernet header for VXLAN

```
    Destination: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
        Address: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 00:00:00:00:61:00 (00:00:00:00:61:00)
        Address: 00:00:00:00:61:00 (00:00:00:00:61:00)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.4.2

```
, Dst:
```

10.47.10.2 <-- True IPv4 addresses

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 84
    Identification: 0x92f6 (37622)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 63
    Protocol: ICMP (1)
    Header checksum: 0x8651 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.4.2
    Destination: 10.47.10.2
Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0xa383 [correct]
    [Checksum Status: Good]
    Identifier (BE): 70 (0x0046)
    Identifier (LE): 17920 (0x4600)
    Sequence number (BE): 0 (0x0000)
    Sequence number (LE): 0 (0x0000)
    Data (56 bytes)
```

```
0000  78 1e dc 17 00 00 00 00 00 00 00 00 00 00 00 00   x...............
0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00                           ........
        Data: 781edc17000000000000000000000000000000000000000b^@&
        [Length: 56]
```

Encapsulated VXLAN packet reaches Edge-2:

```
<#root>

Edge-2#

monitor capture 1 interface range g1/0/1-2 in match ipv4 host 10.47.1.12 host 10.47.1.13

Edge-2#

monitor capture 1 start

Started capture point : 1
Edge-2#

monitor capture 1 stop

Capture statistics collected at software:
    Capture duration - 7 seconds
    Packets received - 6
    Packets dropped - 0
    Packets oversized - 0

Number of Bytes dropped at asic not collected

Capture buffer will exists till exported or cleared

Stopped capture point : 1

Edge-2#

show monitor capture 1 buffer brief

Starting the packet display ........ Press Ctrl + Shift + 6 to exit

    1   0.000000    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=0/0, ttl=63
    2   0.007826    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=0/0, ttl=63
    3   0.086345    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=1/256, ttl=6
    4   0.097490    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=1/256, ttl=6
    5   1.150969    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=2/512, ttl=6
    6   1.163817    10.47.4.2 -> 10.47.10.2   ICMP 148 Echo (ping) request  id=0x0047, seq=2/512, ttl=6

Edge-2#

show monitor capture 1 buffer detailed


Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 148 bytes on wire (1184 bits), 148 bytes captured (1184 bits) on interface /tmp/epc_ws/wif_to_
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 11, 2023 16:58:12.702159000 UTC
    [Time shift for this packet: 0.000000000 seconds]
```

```
    Epoch Time: 1697043492.702159000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 148 bytes (1184 bits)
    Capture Length: 148 bytes (1184 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:udp:vxlan:eth:ethertype:ip:icmp:data]
Ethernet II, Src: 52:54:00:0a:42:11 (
```

52:54:00:0a:42:11

```
), Dst: 52:54:00:17:fe:65 (
```

52:54:00:17:fe:65

```
)
```

<-- True MAC addresses post L3 rewrite

```
    Destination: 52:54:00:17:fe:65 (52:54:00:17:fe:65)
        Address: 52:54:00:17:fe:65 (52:54:00:17:fe:65)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 52:54:00:0a:42:11 (52:54:00:0a:42:11)
        Address: 52:54:00:0a:42:11 (52:54:00:0a:42:11)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.1.12

```
, Dst:
```

10.47.1.13 <-- RLOC to RLOC

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 134
    Identification: 0x1d7b (7547)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 62
    Protocol: UDP (17)
    Header checksum: 0x0876 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.1.12
    Destination: 10.47.1.13
User Datagram Protocol, Src Port: 65354, Dst Port: 4789
    Source Port: 65354
    Destination Port: 4789
    Length: 114
    [Checksum: [missing]]
    [Checksum Status: Not present]
    [Stream index: 0]
    [Timestamps]
```

```
        [Time since first frame: 0.000000000 seconds]
        [Time since previous frame: 0.000000000 seconds]
Virtual eXtensible Local Area Network
    Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
        1... .... .... .... = GBP Extension: Defined
        .... .... .0.. .... = Don't Learn: False
        .... 1... .... .... = VXLAN Network ID (VNI): True
        .... .... .... 0... = Policy Applied: False
        .000 .000 0.00 .000 = Reserved(R): 0x0000
    Group Policy ID: 0
    VXLAN Network Identifier (VNI):
```

4099 <-- LISP L3 IID

```
    Reserved: 0
Ethernet II, Src: 00:00:00:00:61:00 (
```

00:00:00:00:61:00

```
), Dst: ba:25:cd:f4:ad:38 (
```

ba:25:cd:f4:ad:38

```
)
```

<-- Dummy Ethernet header for VXLAN

```
    Destination: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
        Address: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Source: 00:00:00:00:61:00 (00:00:00:00:61:00)
        Address: 00:00:00:00:61:00 (00:00:00:00:61:00)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.4.2

```
, Dst:
```

10.47.10.2

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 84
    Identification: 0x1abb (6843)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 63
    Protocol: ICMP (1)
    Header checksum: 0xfe8c [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.4.2
    Destination: 10.47.10.2
Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
```

```
    Checksum: 0x044f [correct]
    [Checksum Status: Good]
    Identifier (BE): 71 (0x0047)
    Identifier (LE): 18176 (0x4700)
    Sequence number (BE): 0 (0x0000)
    Sequence number (LE): 0 (0x0000)
    Data (56 bytes)

0000  e8 37 0b 32 00 00 00 00 00 00 00 00 00 00 00 00   .7.2............
0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00                           ........
        Data: e8370b320000000000000000000000000000000000000000b^@&
        [Length: 56]
```

Edge-2 decapsulates the VXLAN header and consult its ARP table to forward the ICMP Request towards 10.47.10.2

<#root>

```
Edge-2#
```

show ip cef vrf red_vn 10.47.10.2

```
10.47.10.2/32
  nexthop 10.47.10.2 Vlan1028

Edge-2#
```

show platform software fed switch active ip adj

```
IPV4 Adj entries
dest                                        if_name                dst_mac          si_hdl         r
----                                        -------                -------          ------         -

10.47.10.2                                  Vlan1028               5254.0002.cbf5   0x7f5744f89988 0x

<snip>
```

Endpoint si_hdl decode

Take the si_hdl (**0x7f5744f89988**) and use in **show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_hdl> 1**

<#root>

```
Edge-2#
```

show platform hardware fed switch active fwd-asic abstraction print-resource-handle 0x7f5744f89988 1

```
Handle:0x7f5744f89988 Res-Type:ASIC_RSC_SI Res-Switch-Num:255 Asic-Num:255 Feature-ID:AL_FID_L3_UNICAST_
priv_ri/priv_si Handle: 0x7f5744f8afa8Hardware Indices/Handles: index0:0xc8  mtu_index/l3u_ri_index0:0x0
Features sharing this resource:66 (1)]
57 (1)]
Cookie length: 56
00 00 00 00 00 00 00 00 04 04 00 00 00 00 00 00 00 00 00 00 07 00 52 54 00 02 cb f5 00 00 00 00 00 00 0
```

```
Detailed Resource Information (ASIC_INSTANCE# 0)
----------------------------------------


Station Index (SI) [0xc8] <-- Station Index contains RI and DI
RI = 0x2c <-- Rewrite Index contains information for L2 Forwarding
DI = 0x526e <-- Rewrite Index contains destination port information

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0x1
Replication Bitmap: LD


Detailed Resource Information (ASIC_INSTANCE# 1)
----------------------------------------

Station Index (SI) [0xc8] <-- Station Index contains RI and DI
RI = 0x2c <-- Rewrite Index contains information for L2 Forwarding
DI = 0x526e <-- Rewrite Index contains destination port information

stationTableGenericLabel = 0
stationFdConstructionLabel = 0x7
lookupSkipIdIndex = 0
rcpServiceId = 0
dejaVuPreCheckEn = 0x1
Replication Bitmap: CD

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

Endpoint RI Decode

Take the RI (**0x2c**) and use in the command **show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>**

```
<#root>

Edge-2#

show platform hardware fed switch active fwd-asic resource asic all rewrite-index range 0x2c 0x2c

ASIC#:0

RI:44

 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)

 <-- Decimal 44 is hex 0x2c

    MAC Addr: MAC Addr:

52:54:00:02:cb:f5

,

<-- MAC Address 5254.0002.cbf5 is 10.47.10.2

    L3IF LE Index 50
```

```
ASIC#:1 RI:44 Rewrite_type:AL_RRM_REWRITE_L3_UNICAST_IPV4_SHARED(1) Mapped_rii:L3_UNICAST_IPV4(9)
 <-- Decimal 44 is hex 0x2c
    MAC Addr: MAC Addr:
52:54:00:02:cb:f5
,
<-- MAC Address 5254.0002.cbf5 is 10.47.10.2
    L3IF LE Index 50
```

Endpoint DI Decode

Take the DI (**0x526e**) and use in **show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>**

```
<#root>
Edge-2#
show platform hardware fed switch active fwd-asic resource asic all destination-index range 0x526e 0x526
ASIC#0:

Destination index   = 0x526e
pmap                = 0x00000000 0x00000010 <-- Convert 10 into binary, 0001 and 0000, so 00010000, and

pmap_intf : [GigabitEthernet1/0/5]
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
ASIC#1:

Destination index   = 0x526e
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

Edge-2#
```

show platform software fed switch active ifm mappings

```
Interface                IF_ID    Inst Asic Core Port SubPort Mac  Cntx LPN  GPN  Type Active
GigabitEthernet1/0/1     0x1a       0   0   0    0      0      1    0    1    1    NIF  Y
GigabitEthernet1/0/2     0x1b       0   0   0    1      0      2    1    2    2    NIF  Y
GigabitEthernet1/0/3     0xb        0   0   0    2      0      3    2    3    3    NIF  Y
GigabitEthernet1/0/4     0xc        0   0   0    3      0      4    3    4    4    NIF  Y
```

GigabitEthernet1/0/5

```
      0xd        0   0   0
```

4

```
     0      5    4    5    5    NIF  Y
```

<-- Port 4 corresponds to Gig1/0/5

```
GigabitEthernet1/0/6     0xe        0   0   0    5      0      6    5    6    6    NIF  Y
GigabitEthernet1/0/7     0xf        0   0   0    6      0      7    6    7    7    NIF  Y
GigabitEthernet1/0/8     0x10       0   0   0    7      0      8    7    8    8    NIF  Y
```

Edge-2 decapsulates the packet and send it towards the egress interface where HOST-03 is connected to:

<#root>

```
Edge-2#
```

monitor capture 1 interface g1/0/5 out match ipv4 host 10.47.4.2 host 10.47.10.2

```
Edge-2#
```

monitor capture 1 start

```
Started capture point : 1
Edge-2#
```

monitor capture 1 stop

```
Capture statistics collected at software:
    Capture duration - 6 seconds
    Packets received - 3
    Packets dropped - 0
    Packets oversized - 0

Number of Bytes dropped at asic not collected

Capture buffer will exists till exported or cleared

Stopped capture point : 1

Edge-2#
```

show monitor capture 1 buffer brief

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit
```

```
    1   0.000000    10.47.4.2 -> 10.47.10.2   ICMP 106 Echo (ping) request  id=0x0048, seq=0/0, ttl=62
    2   0.984985    10.47.4.2 -> 10.47.10.2   ICMP 106 Echo (ping) request  id=0x0048, seq=1/256, ttl=6
    3   1.985357    10.47.4.2 -> 10.47.10.2   ICMP 106 Echo (ping) request  id=0x0048, seq=2/512, ttl=6

Edge-2#
```

show monitor capture 1 buffer detailed

```
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

Frame 1: 106 bytes on wire (848 bits), 106 bytes captured (848 bits) on interface /tmp/epc_ws/wif_to_ts
    Interface id: 0 (/tmp/epc_ws/wif_to_ts_pipe)
        Interface name: /tmp/epc_ws/wif_to_ts_pipe
    Encapsulation type: Ethernet (1)
    Arrival Time: Oct 11, 2023 17:22:20.730331000 UTC
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1697044940.730331000 seconds
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 106 bytes (848 bits)
    Capture Length: 106 bytes (848 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:cmd:ethertype:ip:icmp:data]
Ethernet II, Src:
```

00:00:00:00:61:00

```
 (00:00:00:00:61:00), Dst:
```

ff:ff:ff:ff:ff:ff

```
 (ff:ff:ff:ff:ff:ff)
```

<-- Dummy Ethernet header, EPC does not capture it properly

```
    Destination: ff:ff:ff:ff:ff:ff (ff:ff:ff:ff:ff:ff)
        Address: ff:ff:ff:ff:ff:ff (ff:ff:ff:ff:ff:ff)
        .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory d
        .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
    Source: 00:00:00:00:61:00 (00:00:00:00:61:00)
        Address: 00:00:00:00:61:00 (00:00:00:00:61:00)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: CiscoMetaData (0x8909)
Cisco MetaData
    Version: 1
    Length: 1
    Options: 0x0001
    SGT: 0
    Type: IPv4 (0x0800)
Internet Protocol Version 4, Src:
```

10.47.4.2

```
, Dst:
```

10.47.10.2 <-- True IP addresses

```
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
```

```
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
        0000 00.. = Differentiated Services Codepoint: Default (0)
        .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 84
    Identification: 0x35e4 (13796)
    Flags: 0x4000, Don't fragment
        0... .... .... .... = Reserved bit: Not set
        .1.. .... .... .... = Don't fragment: Set
        ..0. .... .... .... = More fragments: Not set
    Fragment offset: 0
    Time to live: 62
    Protocol: ICMP (1)
    Header checksum: 0xe463 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.47.4.2
    Destination: 10.47.10.2
Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x2693 [correct]
    [Checksum Status: Good]
    Identifier (BE): 72 (0x0048)
    Identifier (LE): 18432 (0x4800)
    Sequence number (BE): 0 (0x0000)
    Sequence number (LE): 0 (0x0000)
    Data (56 bytes)

0000  69 9c 67 88 00 00 00 00 00 00 00 00 00 00 00 00   i.g.............
0010  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0020  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00                           ........
        Data: 699c6788000000000000000000000000000000000000000000000000b^@&
        [Length: 56]
```