

排除SDA转发东 — 西流量故障

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[相关产品](#)

[背景信息](#)

[拓扑](#)

[配置](#)

[主机登录验证](#)

[IPDT/IP设备跟踪条目](#)

[MAC/ARP条目](#)

[LISP条目](#)

[SDA中的ARP解析](#)

[基本主机可达性在SDA交换矩阵中 \(相同VLAN/相同VN \)](#)

[基本主机可达性SDA交换矩阵 \(不同VLAN/相同VN \)](#)

简介

本文档介绍如何验证作为软件定义访问(SDA)一部分的东 — 西流量。

先决条件

要求

Cisco 建议您了解以下主题：

- Internet协议(IP)转发
- 定位器/ID分离协议(LISP)

使用的组件

本文档中的信息基于以下软件和硬件版本：

- Cisco IOS® XE 17.10.1上的C9000v
- SDA 1.0 (非LISP PubSub)

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始 (默认) 配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

相关产品

本文档也可用于以下硬件和软件版本：

- C9200
- C9300
- C9400
- C9500
- C9600
- Cisco IOS® XE 16.12及更高版本


背景信息

SDA东 — 西流量是指在SDA交换矩阵内的终端希望与同一交换矩阵内的另一个终端进行通信的概念。关于什么是东 — 西流，什么不认为是东 — 西流，有一些注意事项。东 — 西流量可以是以下示例：


- 位于同一子网（172.17.10.2与172.19.10.3通信）中的终端被视为L2LISP扩展
- 位于同一VRF(VN)中的终端（172.19.10.2与172.19.11.2通信，并且两者都位于VRF园区中）这被视为L3 LISP
- 交换矩阵内的终端与连接到L2切换边界的主机进行通信，与L2LISP完全相同

东 — 西流量不参考以下示例：

- 从SDA交换矩阵到交换矩阵外部（即南北向）的流量
- VRF间路由也不视为东 — 西（VRF园区中的终端，IP地址172.19.10.2与VRF访客中的终端通话，IP地址172.19.11.2）
- SD-WAN集成域
- SDA传输
- 边界关联
- 外联网

 注意：平台(fed)命令可能有所不同。命令可以是“show platform fed <active|standby>”和“show platform fed switch <active|standby>”。如果示例中注明的语法未解析出，请尝试变体。

拓扑



在本示例中，C9000v交换机用作交换矩阵边缘和并置边界。所有终端都位于同一个虚拟网络(VN)(red_vn)中。位于10.47.4.2和10.47.4.2的终端位于同一子网中，位于10.47.10.2的终端位于不同子网中，但属于同一VN。

配置

假设使用Cisco DNA-Center来调配具有默认设置的SDA交换矩阵：

- 启用第2层扩展（这会强制根据MAC地址查找而不是IP地址查找来转发流量）。
- 第2层泛洪被禁用（这将在边缘设备上启用ARP抑制并启用LISP辅助ARP学习）。

在正确的主机登录过程之后，接口配置包含以下几个部分：

交换矩阵边缘(10.47.1.12)接口配置：

```
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
```

交换矩阵边缘(10.47.1.12)LISP配置：

```

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

```

交换矩阵边缘(10.47.1.13)接口配置：

```

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

```

交换矩阵边缘(10.47.1.13)LISP配置

```

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

```

主机登录验证

作为主机自注册过程的一部分，会创建以下几种结构：

IPDT/IP设备跟踪条目

主机成功登录后，IP设备跟踪(IPDT)表中存在有效条目，终端主机标记为REACHABLE:

```
<#root>
```

```
Edge-1#
```

```
show device-tracking database interface gi1/0/3
```

```
portDB has 2 entries for interface Gi1/0/3, 2 dynamic
```

```
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol, DHCP - IPv4 DHCP
```

```
Preflevel flags (prlvl):
```

```

0001:MAC and LLA match      0002:Orig trunk          0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access 0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated  0100:Statically assigned

```

| | Network Layer Address | Link Layer Address | Interface | vlan | prlvl | ag |
|-----|-----------------------|--------------------|-----------|------|-------|----|
| DH4 | 10.47.4.2 | 5254.0019.93e9 | Gi1/0/3 | 1026 | 0024 | 3m |

MAC/ARP条目

当终端主机成功入网时，它可以ping默认网关（或者，如果在阻止此通信的终端未安装防火墙，则可以从默认网关ping通）：

<#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

在边缘节点上，表中有一个MAC地址以及相应的ARP条目（在VRF中）：

<#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 |

软件FED MAC地址编程**

要检查FED中的MAC地址，请使用命令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地址macHandle编程

从上一个命令(0x7f65ec7bda68)获取macHandle值并用于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, learning 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验证****

上一个输出中的数字7是硬件中的映射VLAN ID(MVID)。要验证它们是否与“real”vlan匹配，请使用 `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 0x0000000000000007
```

```
7
```

****全局端口号(GPN)验证****

要将GPN与“实际”接口相关联，请使用命令 `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地址siHandle编程

使用之前命令(0x7f65ec7c21f8)中的siHandle值，并在show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_handle> 1中使用

<#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/13u_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 00 00

Detailed Resource Information (ASIC_INSTANCE# 0)

Station Index (SI) [0xc3] <-- Station Index is comprised of the Rewrite Index (RI) and Destination Index

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

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地址重写索引验证

从上一个命令(0x25)获取RI值，并在show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>中使用

<#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地址目标索引验证

使用上一个命令的DI值(0x526e)，并利用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 0x526e
```

ASIC#0:

Destination index = 0x526e

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

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

端口验证

要关联先前看到的端口，请使用show platform software fed switch active ifm mappings命令并查看Port列。

<#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

****硬件馈送MAC地址验证****

在工作/理想情况下，此输出与macHandle解码器提供的输出匹配。

<#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

- 硬件中的VLAN ID(MVID)为7
- MAC地址:5254.0019.93e9
- GPN:3

LISP条目

在成功主机登录后，终端主机的LISP条目在边缘节点本地创建，并在控制节点 (LISP MSMR -

LISP映射服务器/映射解析器)上注册。对于可以为L2和L3检查的特定实例ID范围，需要进行所有LISP检查：

```
<#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数据库验证****

要检查L2 LISP数据库，请使用命令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地址解析(AR)数据库验证****

要检查LISP L2 AR数据库，请使用命令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 |

****LISP L3数据库验证****

要检查LISP L3数据库，请使用命令show lisp instance-id <LISP L3 IID> ipv4 database <IP地址/子网掩码>

<#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验证****

要检查CEF，请使用命令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

sublocks:

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, hwc]

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

除SDA边缘节点上的本地LISP条目外，SDA控制节点(LISP MS/MR)还包含有关终端的额外信息：

并置边界L2 LISP服务器验证：

要检查L2 LISP服务器，请使用命令show lisp instance-id <L2 LISP ID> 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)
```

并置边界L2 LISP地址解析(AR)服务器验证：

要检查L2 LISP AR服务器，请使用命令show lisp instance-id <LISP L2 IID> ethernet server address-resolution <IP address>

要检查注册历史记录，请使用命令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 |
|----------------------|----------|-------|------|-----|---------------------------------|
| *Sep 29 16:50:27.762 | 8190 | TCP | No | No | 10.47.1.12 |
| | | | | | EID prefix / Locator |
| | | | | | +*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 |

并置边界L3 LISP服务器验证

要检查L3 LISP服务器，请使用命令show lisp instance-id <LISP L3 IID> ipv4 server <IP address>

要检查L3 LISP服务器注册历史记录，请使用命令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 |

SDA中的ARP解析

假设已使用Cisco Catalyst Center来调配具有默认设置的SDA交换矩阵。这意味着第2层扩展已启用，并且交换矩阵内的所有流量（位于同一VLAN/VN中）都基于MAC地址查找/LISP以太网实例（而不是IP地址查找/LISP IP实例）进行转发。

从故障排除的角度来看，在两台主机上配置静态ARP条目以快速检查问题是否与交换矩阵中的通用连接有关（在这种情况下，ping在主机之间不起作用）或仅与ARP解决有关，可能非常有用。

SDA交换矩阵中的ARP过程利用LISP来解析主机的识别和位置，与传统路由/交换环境中的ARP行为不同。

第1步：交换矩阵终端发送ARP请求以确定其他交换矩阵终端的MAC/IP绑定

可以在入口接口上配置数据包捕获，以确认已从主机接收ARP数据包：

```
<#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 - 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: 00000000000000000000000000000000
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

```

第二步：边缘节点使用ARP数据包并生成LISP请求以确定HOST-02的MAC地址。

Edge-1向LISP控制平面（并置边界）发送解析MAC地址10.47.4.3的LISP映射请求：

```
<#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'.

```

第三步：控制节点接收IP/MAC映射的LISP请求，并将响应发送回SDA边缘节点

从交换矩阵边缘接收LISP映射请求，并使用与10.47.4.3关联的MAC地址进行LISP映射回复响应

```
<#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 to 10.47.1.12  
LISP[MR ]-0: Received Map-Request with 1 records, first EID IID 8190 10.47.4.3/32, source EID UNSPECIFIED, destination EID UNSPECIFIED  
LISP[MR ]-0 IID 8190 Eth-ARP: MS EID 10.47.4.3/32: Sending proxy reply to 10.47.1.12.
```

LISP控制平面根据本地数据库中存储的地址解析条目使用代理应答进行响应

```
<#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
```

第四步：边缘节点收到的LISP回复的MAC地址为10.47.4.3

交换矩阵边缘节点收到LISP代理应答：

```
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.
```

第五步：边缘节点发送LISP映射请求数据包，以确定MAC地址的RLOC位置

成功完成前三个步骤后，边缘节点知道最初为其生成ARP的MAC地址10.47.4.3。启用第2层扩展后，边缘节点不会用此信息回复到10.47.4.2，而是使用它来确定出口节点边缘的RLOC位置，以便像在传统第2层网络中那样将ARP转发到10.47.4.3。

因此，边缘节点在以太网实例中生成另一个LISP映射请求数据包，这次请求10.47.4.2的MAC地址的RLOC信息：

```
<#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
```

第六步：控制节点接收LISP映射请求数据包，以确定MAC地址的RLOC位置

控制节点接收LISP数据包，并根据其本地数据库状态回复该数据包

```
<#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
```


第7步：边缘节点收到LISP映射应答

由控制节点生成的LISP映射应答由边缘节点接收：

```
<#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
```

这最终在LISP以太网实例映射缓存中创建一个条目，并允许将ARP数据包转发到连接10.47.4.3的边缘2

```
<#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)
```

步骤 8 ARP封装在VXLAN中并发送到HOST-02

需要所有与LISP相关的步骤来确定10.47.4.3的位置，以便边缘节点能够将原始ARP（广播）数据包作为单播发送到正确的边缘节点。即使从10.47.4.2发送单个ARP数据包，原始的ARP请求也会被边缘节点CPU缓存（不会丢弃），直到完成所有步骤，才能进行正确的ARP解析。

ARP数据包封装在VXLAN中，如示例所示：

```
<#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
```

ARP请求已封装在VXLAN中，并且已从广播ARP请求转换为单播ARP请求。

```
<#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 default)
```

```
.... ..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 default)
```

```
.... ..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: 00000000000000000000000000000000
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
```

步骤 9 ARP应答由10.47.4.3生成并发送至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
```

此时，数据包的目的地不是广播地址（作为原始ARP请求），而是10.47.4.2的MAC地址，当它到达入口边缘节点(Edge-2)时，会触发正常的LISP操作。最初，边缘节点的LISP以太网实例中缺少MAC地址10.47.4.2，数据包被传送到CPU以生成LISP映射请求，从而确定HOST-01的RLOC。此行为与本文档其他部分中描述的完全相同，并且允许为Edge-2上的10.47.4.2创建LISP映射缓存条目：

```
<#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)
```

该条目允许在VXLAN封装中向Edge-1成功发送ARP应答，并进一步转发到竞争整个ARP解析过程10.47.4.2。

SDA交换矩阵中的基本主机可达性（相同VLAN/相同VN）

假设ARP解析成功完成，并且主机10.47.4.2和10.47.4.3彼此具有适当的ARP条目。

从故障排除的角度来看，在两台主机上配置静态ARP条目以快速检查问题是否与交换矩阵中的通用连接有关（在这种情况下，主机之间的ping不起作用）或仅与ARP进程有关，是非常有用的。

10.47.4.2生成指向10.47.4.3的ICMP请求：

```
<#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
```



```
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)
```

检查远程终端的MAC地址，它指向预期的L2L10

```
<#root>
```

```
Edge-1#
```

```
show mac add add 5254.001e.ad00
```

```
Mac Address Table
```

```
-----  
Vlan    Mac Address      Type      Ports  
----    -  
1026    5254.001e.ad00  CP_LEARN  L2L10
```

```
Total Mac Addresses installed by LISP: REMOTE: 1
```

检查FED中的MAC地址，可以收集其他信息

```
<#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地址macHandle解码

从上一个命令获取macHandle(0x7f65ecfdd3a8)，并在命令show platform hardware fed switch active fwd-asic abstraction print-resource-handle <macHandle> 1中使用

<#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

dejaVuPreCheckEn = 0
Replication Bitmap: LD

=====

重写索引解码

采用RI(0x12)并在命令show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>中使用

<#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_PAYLOAD
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_PAYLOAD
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

采用DI(0x5012)并在命令show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>中使用

<#root>

Edge-1#

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

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解码


```
is Sgt:      1
is TTL Prop: 0
L3if LE:     0 (0)
Port LE:     279 (0)
Vlan LE:     7 (0)
```

底层路由验证

使用10.47.1.12将流量封装在IID为8190的VXLAN中，并且能够对Gig1/0/1和G1/0/2进行负载均衡

```
<#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
```

要获取si_hdl、ri_hdl信息，请使用命令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    LISPO.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
```

Underlay Next-Hop si_hdl Decode

要检查si_hdl(0x7f65ec8a5458)，请使用命令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/13u_ri_index0:0x
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 00
```

```
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
```

```
=====
```

底层下一跳重写索引解码

要解码RI(0x1a)，请在show platform hardware fed switch active fwd-asic resource ASIC all rewrite-index range <RI> <RI>命令中使用

<#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

底层下一跳目标索引解码

要解码show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>中使用的DI(0x526d)

<#root>

Edge-1#

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

ASIC#0:

Destination index = 0x526d

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

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

```

底层下一跳ri_hdl解码

解码show platform hardware fed switch active fwd-asic abstraction print-resource-handle(ri_hdl)1中使用的ri_hdl(0x7f65ec8a4eb8)

<#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/13u_ri_index0:0x
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 08 00 52 54 00 1c 7d e0 00 00 00 00 00 00 00

```



```

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

```

SDA交换矩阵中的基本主机可达性 (不同VLAN/相同VN)

在本节中，将检查10.47.4.2和10.47.10.2之间的通信。由于这些主机属于不同的VLAN，因此两台主机都需要配置指向默认网关的默认网关。10.47.4.2是10.47.4.1,10.47.10.2是10.47.10.1。

步骤1:确认终端和默认网关之间的连接工作正常：

```
<#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
```

第二步：确认Edge-1已成功收到来自10.47.4.2的数据包：

数据包可在面向10.47.4.2的入口接口上捕获：

<#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 .....
0020 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]

```

第3步 — LISP查找

入口边缘节点必须确定它向其发送数据包的主机-03的位置(RLOC)。在本例中，终端主机HOST-03位于不同的VLAN中（但具有相同的VN/VRF:USERS），LISP IPv4实例用作基于IP地址的查找（MAC地址属于边缘节点本身）。

```
<#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 prefir
LISP[REMT ]-0: Map-Reply nonce matches pending request for IID 4099 EID 10.47.10.2/32, requester 'remot

```

LISP映射请求到达控制节点（LISP映射服务器）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映射应答到达边缘节点：

```
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.
```

交换矩阵边缘查询RLOC以获取10.47.10.2并处理映射应答

```
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.
```

如果条目不存在，则需要从LISP进程的角度收集调试。还有一种称为LIG (LISP组) 的工具，可用于手动触发LISP进程 (这是测试冗余控制节点配置和两个控制节点之间的数据库一致性的非常有效的方法) ：

```
<#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 msec
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 msec
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        -
```

路由验证

CEF使用LISP，而LISP使用其已接收的映射缓存条目

```
<#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下一跳验证

由于此数据包是VXLAN封装的，因此需要验证LISP下一跳。使用命令show platform software fed switch active ip adj获取有关10.47.1.13 (LISP下一跳) 的其他信息

<#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 | 0 |

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下一跳DI解码

采用DI(0x5012)并用于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 0x5012
```

```
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
```


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_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 <-- Edge-2 RLOC

IPv4 TTL: 0

LISP INSTANCEID: 0

L3IF LE Index: 46

=====

底层下一跳验证

要到达LISP下一跳，底层有两个可能的路径，其中一个路径进行验证，另一个底层下一跳的验证应用相同的逻辑。

<#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

要获取有关下一跳的详细信息，请使用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

采用si_hdl(0x7f65ec8a5458)并在命令show platform hardware fed switch active fwd-asic abstraction print-resource-handle <si_hdl> 1中使用

```
<#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/13u_ri_index0:0x
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 00
```

```
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
```

```
=====
```

底层下一跳RI解码

采用RI(0x1a)并在命令show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>中使用

```
<#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
```

底层下一跳DI解码

采用DI(0x526d)并在命令show platform hardware fed switch active fwd-asic resource asic all destination-index range <DI> <DI>中使用

```
<#root>
```

```
Edge-1#
```

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

```
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

a1_rsc_cmi

CPU Map Index (CMI) [0]

ctiLo0 = 0

ctiLo1 = 0

ctiLo2 = 0

cpuQNum0 = 0

cpuQNum1 = 0

cpuQNum2 = 0

npuIndex = 0

stripSeg = 0

copySeg = 0

ASIC#1:

Destination index = 0x526d

pmap = 0x00000000 0x00000000

cmi = 0x0

rcp_pmap = 0x0

a1_rsc_cmi

CPU Map Index (CMI) [0]

ctiLo0 = 0

ctiLo1 = 0

ctiLo2 = 0

cpuQNum0 = 0

cpuQNum1 = 0

cpuQNum2 = 0

npuIndex = 0

stripSeg = 0

copySeg = 0

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

底层下一跳ri_hdl解码

采用ri_hdl(0x7f65ec8b8158)并在命令show platform hardware fed switch active fwd-asic abstraction print-resource-handle <ri_hdl> 1中使用

<#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/13u_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 00 00
```

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
```

数据包封装在VXLAN中并根据负载均衡规则发送。嵌入式数据包捕获(EPC)可用于同时捕获所有接口上的流量。请记住，此时数据包是VXLAN封装的，EPC过滤器必须针对RLOC到RLOC，而不是内部IPv4地址。

<#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=63 |
| 3 | 1.984077 | 10.47.4.2 -> 10.47.10.2 | ICMP 148 Echo (ping) request id=0x0046, seq=2/512, ttl=63 |
| 4 | 2.999989 | 10.47.4.2 -> 10.47.10.2 | ICMP 148 Echo (ping) request id=0x0046, seq=3/768, ttl=63 |

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

```


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

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

Edge-2解封VXLAN报头并查询其ARP表，以向10.47.10.2转发ICMP请求

<#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>

端点si_hdl解码

采用si_hdl(0x7f5744f89988)并用于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/13u_ri_index0:0xc8
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 00 00
```

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
```

=====

终端RI解码

采用RI(0x2c)并在命令show platform hardware fed switch active fwd-asic resource asic all rewrite-index range <RI> <RI>中使用

```
<#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.cb f5 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.cb f5 is 10.47.10.2
```

```
  L3IF LE Index 50
```

终端DI解码

采用DI(0x526e)并用于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 0x526e
```

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解封数据包并将其发送到连接HOST-03的出口接口：

<#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=62
3  1.985357  10.47.4.2 -> 10.47.10.2  ICMP 106 Echo (ping) request id=0x0048, seq=2/512, ttl=62
```

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 default)

.... ...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

关于此翻译

思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。