

# EX硬件：ACI信息包转发深潜。

## 目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[方案](#)

[在同样EPG/Same分支的2个EP -交换式帧](#)

[拓扑](#)

[伊拉姆](#)

[在不同的EPG/Same分支的2个EP -路由信息包](#)

[拓扑](#)

[伊拉姆](#)

[在不同的EPG/Different分支的2个EP -路由信息包](#)

[拓扑](#)

[伊拉姆](#)

[1个EP --> L3外的已路由流](#)

[拓扑](#)

[伊拉姆](#)

[1个EP --> 远程EP或SVI -脊椎验证](#)

[拓扑](#)

[逻辑](#)

[集成IP](#)

[结构模块伊拉姆](#)

[额外的方案：获得不在“hal内部端口pi”输出中的Ovector](#)

[拓扑](#)

[逻辑](#)

## 简介

本文描述不同的转发方案使用在应用程序中心基础设施(ACI)的新一代ASIC硬件。它将显示如何验证硬件正确地被编程，并且我们是转发数据包对正确目的地终端(EP)在适合的终端组(EPGs)中。

## [先决条件](#)

### [要求](#)

本文档没有任何特定的要求。

### 使用的组件

本文档中的信息基于下列硬件和软件版本：



我们在软件方面认识FD\_VLAN 30匹配，但是我们能总是验证映射：

```
leaf4# show vlan extended | grep 2268 30 enet CE vlan-2268
```

当然并且，我们能检查硬件确保VLAN 30地图到VLAN 2268作为前面板封装。

```
leaf4# vsh_lc
```

```
module-1# show system internal eltc info vlan 30 vlan_id: 30 ::: hw_vlan_id: 22 vlan_type:
FD_VLAN ::: bd_vlan: 28 access_encap_type: 802.1q ::: access_encap: 2268 fabric_encap_type:
VXLAN ::: fabric_encap: 11960 sclass: 32778 ::: scope: 11 untagged: 0 access_encap_hex: 0x8dc :::
fabric_enc_hex: 0x2eb8 pd_vlan_ft_mask: 0x8 fd_learn_disable: 0 qos_class_id: 0 ::: qos_pap_id:
0 qq_met_ptr: 25 ::: ipmc_index: 0 ingressBdAclLabel: 0 ::: ingBdAclLblMask: 0 egressBdAclLabel:
0 ::: egrBdAclLblMask: 0 qos_map_idx: 0 ::: qos_map_pri: 0 qos_map_dscp: 0 ::: qos_map_tc: 0
vlan_ft_mask: 0xe30 hw_bd_idx: 0 ::: hw_epg_idx: 11267 intf_count: 2 ::: glbl_scp_if_cnt: 2
<SNIPPED>
```

在的情况下EP在软件方面学习，我们能也验证硬件编程了这些EP的L2信息。在新的硬件中，有是硬件的软件状态的硬件抽象层(HAL)。HAL的工作是采取编程请求和推送他们对硬件。

为了查看L2关于终端的硬件信息，我们能查看在HAL的L2表为给的MAC地址：

```
leaf4# vsh_lc
```

```
module-1# show platform internal hal ep l2 mac 0050.56a5.fccc LEGEND: ----- BDI: BD Id BD
Name: BD Name T: EP Type (P: Physical V: Virtual Xr: Remote EP Mac: Mac L2 IfId: L2 Interface
L2 IfName: L2 IfName FDI: FD Id FD Name: FD Name S Class: S Class Age Intvl: Age Interval P A:
Packet Action (F: Forward, T: Trap to CPU, L: Log & Forward, D: Drop, N: None) S T: Static Ep S
E: Secure EP L D: Learn Disable B N D: Bind Notify Disable E N D: Epg Notify Disable B E: Bounce
Enable I D L: IVxlan Dont Learn SPI: Source Policy Incomplete DPI: Dest Policy Incomplete SPA:
Source Policy Applied DPA: Dest Policy Applied DSS: Dest Shared Service IL: Is Local VUB: Vnid
Use Bd SO: SA Only L2 EP Count: 1
```

```
=====
===== B E I S D S D D V BD EP L2 L2 FD S Age P S S L N N B D P P P P S I U S
BdId Name T Mac IfId Ifname FDI: Name Class Intvl A T E D D D E L I I A A S L B O
```

```
=====
===== 1c BD-28 P1 00:50:56:a5:fc:cc 16000002 Po3 1e FD-30 800a 29f F 0 0 0 1 0
0 0 0 0 0 0 1 0 0 module-1# show platform internal hal ep l2 mac 0050.56a5.6794
```

```
=====
===== B E I S D S D D V BD EP L2 L2 FD S Age P S S L N N B D P P P P S I U S
BdId Name T Mac IfId Ifname FDI: Name Class Intvl A T E D D D E L I I A A S L B O
```

```
=====
===== 1c BD-28 P1 00:50:56:a5:67:94 16000003 Po4 1e FD-30 800a 29f F 0 0 0 1 0
0 0 0 0 0 0 1 0 0
```

既然我们映射硬件，请执行伊拉姆和发现数据包哪里应该是。

## 伊拉姆

```
leaf4# vsh_lc
```

```
module-1# debug platform internal tah elam asic 0 module-1(DBG-TAH-elam)# trigger reset module-
1(DBG-TAH-elam)# trigger init in-select 6 out-select 0 module-1(DBG-TAH-elam-insel6)# set outer
l2 src_mac 0050.56a5.fccc dst_mac 0050.56a5.6794 module-1(DBG-TAH-elam-insel6)# start module-
1(DBG-TAH-elam-insel6)# stat ELAM STATUS ===== Asic 0 Slice 0 Status Armed Asic 0 Slice 1
Status Triggered module-1(DBG-TAH-elam-insel6)# report | grep ovec
sug_elam_out_sidebnd_no_spare_vec.ovector_idx: 0x9E
```

极大，Leaf4如此接收在Asic的帧0片式1。使用是非常重要的新的硬件的伊拉姆，有新字段，当排除故障时：ovector\_idx。此索引是物理端口索引应该转发帧/数据包在外面。一旦有ovector\_idx，我们能使用此命令查找什么端口映射到：

```
module-1(DBG-TAH-elam-insel6)# show platform internal hal l2 port gpd Legend: ----- IfId:
Interface Id IfName: Interface Name I P: Is PC Mbr IfId: Interface Id Uc PC Cfg: UcPcCfg Idx Uc
PC MbrId: Uc Pc Mbr Id As: Asic AP: Asic Port Sl: Slice Sp: Slice Port Ss: Slice SrcId Ovec:
Ovector (slice | srcid) L S: Local Slot Reprogram: L3: Is L3 P: PifTable Xla Idx: Xlate Idx RP:
```

```
Rw PifTable OvX Idx: OXlate Idx IP: If Profile Table N L3: Num. of L3 Ifs RS: Rw SrcId Table NI
L3: Num. of Infra L3 Ifs DP: DPort Table Vif Tid: Vif Tid SP: SrcPortState Table RwV Tid: RwVif
Tid RSP: RwSrcPortstate Table Ing Lbl: Ingress Acl Label UC: UCPcCfg Egr Lbl: Egress Acl Label
UM: UCPcMbr Reprogram: PROF ID: Lport Profile Id VS: VifStateTable HI: LportProfile Hw Install
RV: Rw VifTable Num. of Sandboxes: 1 Sandbox_ID: 0, BMP: 0x0 Port Count: 8
```

```
=====
===== Uc Uc | Reprogram | | Rep | I PC Pc L | R I R D R U
U X | L Xla OvX N NI Vif RwV Ing Egr | V R | PROF H IfId Ifname P Cfg MbrID As AP Sl Sp Ss Ovec
S | P P P S P Sp Sp C M L | 3 Idx Idx L3 L3 Tid Tid Lbl Lbl | S V | ID I
```

```
=====
===== 1a004000 Eth1/5 1 0 1d 0 d 0 c 18 18 1 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 - - 800 0 0 1 0 0 1a005000 Eth1/6 1 0 b 0 e 0 d 1a 1a 1 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 - - 800 0 0 1 0 0 1a006000 Eth1/7 0 26 5 0 f 0 e 1c 1c 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
D-256 - 800 0 0 1 e 0 1a007000 Eth1/8 0 2e 7 0 10 0 f 1e 1e 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 D-84
- 800 0 0 1 30 0 1a01e000 Eth1/31 1 0 2d 0 37 1 e 1c 9c 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 - - 0 0
0 1 0 0 1a01f000 Eth1/32 1 0 3d 0 38 1 f 1e 9e 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 - - 0 0 0 1 0 0
1a030000 Eth1/49 0 2 1 0 49 1 20 38 b8 1 0 0 0 0 0 0 0 0 0 0 0 1 8 6 2 2 D-24d - 400 0 0 0 1 0
1a031000 Eth1/50 0 3 3 0 29 1 0 0 80 1 0 0 0 0 0 0 0 0 0 0 0 1 9 7 2 2 D-350 - 400 0 0 0 1 0
```

交换机认为应该转发数据包在以太网接口1/32外面。该PO4我们了解mac address的地方？

```
leaf4# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
       F - Configuration failed
```

```
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
1      Po1(SU)     Eth       LACP      Eth1/5(P)
2      Po2(SU)     Eth       LACP      Eth1/6(P)
3      Po3(SU)     Eth       LACP      Eth1/31(P)
4      Po4(SU)     Eth LACP  Eth1/32(P)
```

是，因此数据包增殖比转发在接口1/32外面对目的地主机。

## 在不同的EPG/Same分支的2个EP -路由信息包

### 拓扑



```

common*rewall Pl 10.6.112.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - -
0.0.0.0 common*rewall Pl 10.6.114.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 -
- - 0.0.0.0 common*rewall Pl 10.6.114.129 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 -
00:00:00:00:00:00 - - - 0.0.0.0 common*efault Pl 100.100.101.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0
- L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Pl 192.168.1.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0
1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Xr 192.168.1.100 8013 128 0 0 0 1 0 0
0 0 0 0 0 0 1 0 - L3 - 00:0c:0c:0c:0c:0c Tunnel2 Tunnel2 - 0.0.0.0 Joey-T*ternal2 Pl
192.168.3.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-
T*ternal Pl 192.168.20.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - -
0.0.0.0 Joey-T*ternal Pl 192.168.20.2 800a 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 - L2 BD-28
00:50:56:a5:fc:cc - Po3 FD-30 - Joey-T*ternal Pl 192.168.21.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0
- L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Pl 192.168.21.2 800c 0 0 0 0 0 0 0 0 0 0 0
0 0 1 0 0 - L2 BD-7 00:50:56:a5:0c:11 - Po4 FD-8 - Joey-T*ternal Pl 2001:0:0:100::1 1 0 1 0 0 0
0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0

```

HAL第3层(I3)表是非常useful，因为提供我们I3了解的EP的VLAN/Port信息。我们知道目的地存在Po4，因此应该转发数据包在Po4的所有端口外面。

请运行伊拉姆和发现什么我们获得!

### 伊拉姆

```

leaf4# vsh_lc
module-1# debug platform internal tah elam asic 0 module-1(DBG-TAH-elam)# trigger init in-select
6 out-select 0 module-1(DBG-TAH-elam-insel6)# set outer ipv4 src_ip 192.168.20.2 dst_ip
192.168.21.2 module-1(DBG-TAH-elam-insel6)# start module-1(DBG-TAH-elam-insel6)# stat ELAM
STATUS ===== Asic 0 Slice 0 Status Armed Asic 0 Slice 1 Status Armed module-1(DBG-TAH-
elam-insel6)# stat ELAM STATUS ===== Asic 0 Slice 0 Status Armed Asic 0 Slice 1 Status
Triggered module-1(DBG-TAH-elam-insel6)# report | grep ovec
sug_elam_out_sidebnd_no_spare_vec.ovector_idx: 0x9E

```

极大，因此我们触发了数据包和我们发现“ovector\_idx”是0x9E。ovector索引是流出的physical接口索引应该转发数据包在外面。请发现什么端口有该索引：

```

module-1(DBG-TAH-elam-insel6)# show platform internal hal l2 port gpd Legend: ----- IfId:
Interface Id IfName: Interface Name I P: Is PC Mbr IfId: Interface Id Uc PC Cfg: UcPcCfg Idx Uc
PC MbrId: Uc Pc Mbr Id As: Asic AP: Asic Port Sl: Slice Sp: Slice Port Ss: Slice SrcId Ovec:
Ovector (slice | srcid) L S: Local Slot Reprogram: L3: Is L3 P: PifTable Xla Idx: Xlate Idx RP:
Rw PifTable Ovx Idx: OXlate Idx IP: If Profile Table N L3: Num. of L3 Ifs RS: Rw SrcId Table NI
L3: Num. of Infra L3 Ifs DP: DPort Table Vif Tid: Vif Tid SP: SrcPortState Table RwV Tid: RwVif
Tid RSP: RwSrcPortstate Table Ing Lbl: Ingress Acl Label UC: UCPcCfg Egr Lbl: Egress Acl Label
UM: UCPcMbr Reprogram: PROF ID: Lport Profile Id VS: VifStateTable HI: LportProfile Hw Install
RV: Rw VifTable Num. of Sandboxes: 1 Sandbox_ID: 0, BMP: 0x0 Port Count: 8
=====
===== Uc Uc | Reprogram | | Rep | I PC Pc L | R I R D R U
U X | L Xla Ovx N NI Vif RwV Ing Egr | V R | PROF H IfId Ifname P Cfg MbrID As AP Sl Sp Ss Ovec
S | P P P S P Sp Sp C M L | 3 Idx Idx L3 L3 Tid Tid Lbl Lbl | S V | ID I
=====
===== 1a004000 Eth1/5 1 0 1d 0 d 0 c 18 18 1 0 0 0 0 0 0
0 0 0 0 0 0 0 0 - - 800 0 0 1 0 0 1a005000 Eth1/6 1 0 b 0 e 0 d 1a 1a 1 0 0 0 0 0 0 0 0 0
0 0 0 0 - - 800 0 0 1 0 0 1a006000 Eth1/7 0 26 5 0 f 0 e 1c 1c 1 0 0 0 0 0 0 0 0 0 0 0 0
D-256 - 800 0 0 1 c 0 1a007000 Eth1/8 0 2f 7 0 10 0 f 1e 1e 1 0 0 0 0 0 0 0 0 0 0 0 0 0
D-199 - 800 0 0 1 2e 0 1a01e000 Eth1/31 1 0 2d 0 37 1 e 1c 9c 1 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 1 0 0 1a01f000 Eth1/32 1 0 3d 0 38 1 f 1e 9e 1 0 0 0 0 0 0 0 0 0 0 0 0 0 - - 0 0 0 1 0
0 1a030000 Eth1/49 0 2 1 0 49 1 20 38 b8 1 0 0 0 0 0 0 0 0 0 0 1 6 4 2 2 D-24d - 400 0 0 0 1 0
1a031000 Eth1/50 0 3 3 0 29 1 0 0 80 1 0 0 0 0 0 0 0 0 0 0 1 5 3 2 2 D-350 - 400 0 0 0 1 0

```

查找，如我们应该传送它波尔特1/32，是更正？

```

leaf4# show port-channel summary
Flags: D - Down P - Up in port-channel (members)
I - Individual H - Hot-standby (LACP only)
s - Suspended r - Module-removed

```



```

module-1# show platform internal hal ep l3 all LEGEND: ----- VrfName: Vrf Name T: Type (Pl:
Physical, Vl: Virtual, Xr: Remote) EP IP: Endpoint IP S Class: S Class Age Intvl: Age Interval S
T: Static Ep S E: Secure EP L D: Learn Disable B N D: Bind Notify Disable E N D: Epg Notify
Disable B E: Bounce Enable I D L: IVxlan Dont Learn SPI: Source Policy Incomplete DPI: Dest
Policy Incomplete SPA: Source Policy Applied DPA: Dest Policy Applied DSS: Dest Shared Service
IL: Is Local VUB: Vnid Use Bd SO: SA Only EP NH L3IfName: EP Next Hop L3 If Name NHT: Next Hop
Type (L2: L2 Entry L3: L3 Next Hop) BD Name: L2 NH BD Name EP Mac: EP Mac L3 IfName: L3 NH If
Name L2 IfName: L2 If Name FD Name: L2 Entry FD Name IP: L3 NH IP L3 EP Count: 12
=====
===== B E
I S D S D D V EP-NH N | Vrf EP S Age S S L N N B D P P P S I U S L3 H | BD EP L3 L2 FD Name T
IP Class Intvl T E D D D E L I I A A S L B O IfName T | Name Mac IfName Ifname Name IP
=====
=====
common*rewall Pl 10.6.112.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - -
0.0.0.0 common*rewall Pl 10.6.114.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 -
- - 0.0.0.0 common*rewall Pl 10.6.114.129 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 -
00:00:00:00:00:00 - - - 0.0.0.0 common*efault Pl 100.100.101.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0
- L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Pl 192.168.1.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 0
1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Xr 192.168.1.100 8013 128 0 0 0 1 0 0
0 0 0 0 0 0 1 0 - L3 - 00:0c:0c:0c:0c:0c Tunnel2 Tunnel2 - 0.0.0.0 Joey-T*ternal2 Pl
192.168.3.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-
T*ternal Pl 192.168.20.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - -
0.0.0.0 Joey-T*ternal Pl 192.168.20.2 800a 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 - L2 BD-28
00:50:56:a5:fc:cc - Po3 FD-30 - Joey-T*ternal Pl 192.168.21.1 1 0 1 0 0 0 0 0 1 1 0 0 0 0 1 0 0
- L3 - 00:00:00:00:00:00 - - - 0.0.0.0 Joey-T*ternal Pl 192.168.21.2 800c 0 0 0 0 0 0 0 0 0 0 0
0 0 1 0 0 - L2 BD-7 00:50:56:a5:0c:11 - Po4 FD-8 - Joey-T*ternal Pl 2001:0:0:100::1 1 0 1 0 0 0
0 0 1 1 0 0 0 0 1 0 0 - L3 - 00:00:00:00:00:00 - - - 0.0.0.0

```

硬件认为EP在通道2.存在。 什么是通道的2目的地？

```

module-1# show system internal eltc info interface tunnel2 IfInfo: interface: Tunnel2 :::
ifindex: 402718722 iod: 66 ::: state: up Mod: 0 ::: Port: 0 Tunnel Index: 0 ::: Tunnel Dst ip:
0xc0a87843 Tunnel Encap: ivxlan ::: Tunnel VPC Peer: 0 Tunnel Dst ip str: 192.168.120.67 :::
Tunnel ept: 0x1 [SDK Info]: tunnl_name: vrf_id: 2 ::: if_index: 0x18010002 hwencapidx: 0 :::
encaptype: 1 mac_proxy: 0 ::: v4_proxy: 0 v6_proxy: 0 ::: ip_addr_type: 0 ipv4_address:
0xc0a87843 [SDB INFO]: iod: 66 pc_if_index: 0 fab_if_index: 0 sv_if: 0 src_idx: 0 int_vlan: 0
encap_vlan: 0 mod_port_status: 0x41620003 v6_tbl_id: 0x80000002 v4_tbl_id: 0x2
router_mac:00.00.00.00.00.00 unnumbered: 0 trunk_id: 0 tunnel_mod: 0 tunnel_port: 0 tep_ip:
0xc0a87843 ip_if_mode: 0 sdk_vrf_id: 2 mtu: 9366 ::: ipmtu_id: 0 is_fex_fabric: 0

```

因为目的地存在vPC，该目的地IP应该是远程分支的vPC虚拟IP。 请检查一个远程分支和看：

```

leaf1# show system internal epm vpc Local TEP IP : 192.168.160.95 Peer TEP IP : 192.168.160.93
vPC configured : Yes vPC VIP : 192.168.120.67 MCT link status : Up Local vPC version bitmap :
0x7 Peer vPC version bitmap : 0x7 Negotiated vPC version : 3 Peer advertisement received : Yes
Tunnel to vPC peer : Up

```

完善，因此它学习从远程vPC对的目的地EP。 请发现什么伊拉姆看到，并且请验证我们转发数据包correctly：

## 伊拉姆

```

module-1# debug platform internal tah elam ASIC 0
module-1(DBG-TAH-elam)# trigger init in-select 6 out-select 0 module-1(DBG-TAH-elam-insel6)# set
outer ipv4 src_ip 192.168.20.2 dst_ip 192.168.1.100 module-1(DBG-TAH-elam-insel6)# start module-
1(DBG-TAH-elam-insel6)# stat ELAM STATUS ===== ASIC 0 Slice 0 Status Armed ASIC 0 Slice 1
Status Triggered

```

现在，与在EX硬件的远程目的地，有是非常重要的，当排除故障数据包流时的2个伊拉姆值。  
ovector\_idx喜欢前面和encap\_idx：

```

module-1(DBG-TAH-elam-insel6)# report | grep ovec sug_elam_out_sidebnd_no_spare_vec.ovector_idx:
0xB8 module-1(DBG-TAH-elam-insel6)# report | grep encap sug_lurw_vec.encap_l2_idx: 0x0

```



sug\_lurw\_vec.encap\_pcid: 0x0 sug\_lurw\_vec.encap\_idx: 0x6 sug\_lurw\_vec.encap\_vld: 0x1

在EX硬件上，我们有能力驱动应该转发在外面数据包的目的地端口。前面，我们通常检查encap idx并且验证目的地idx是正确通道。此处我们能验证对8B的什么端口映射：

```
module-1(DBG-TAH-elam-insel6)# show platform internal hal l2 port gpd Legend: ----- IfId:
Interface Id IfName: Interface Name I P: Is PC Mbr IfId: Interface Id Uc PC Cfg: UcPcCfg Idx Uc
PC MbrId: Uc Pc Mbr Id As: Asic AP: Asic Port Sl: Slice Sp: Slice Port Ss: Slice SrcId Ovec:
Ovector (slice | srcid) L S: Local Slot Reprogram: L3: Is L3 P: PifTable Xla Idx: Xlate Idx RP:
Rw PifTable Ovx Idx: OXlate Idx IP: If Profile Table N L3: Num. of L3 Ifs RS: Rw SrcId Table NI
L3: Num. of Infra L3 Ifs DP: DPort Table Vif Tid: Vif Tid SP: SrcPortState Table RwV Tid: RwVif
Tid RSP: RwSrcPortstate Table Ing Lbl: Ingress Acl Label UC: UCPcCfg Egr Lbl: Egress Acl Label
UM: UCPcMbr Reprogram: PROF ID: Lport Profile Id VS: VifStateTable HI: LportProfile Hw Install
RV: Rw VifTable Num. of Sandboxes: 1 Sandbox_ID: 0, BMP: 0x0 Port Count: 8
=====
===== Uc Uc | Reprogram | | Rep | I PC Pc L | R I R D R U
U X | L Xla Ovx N NI Vif RwV Ing Egr | V R | PROF H IfId Ifname P Cfg MbrID As AP Sl Sp Ss Ovec
S | P P P S P Sp Sp C M L | 3 Idx Idx L3 L3 Tid Tid Lbl Lbl | S V | ID I
=====
===== 1a004000 Eth1/5 1 0 1d 0 d 0 c 18 18 1 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 - - 800 0 0 1 0 0 1a005000 Eth1/6 1 0 b 0 e 0 d 1a 1a 1 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 - - 800 0 0 1 0 0 1a006000 Eth1/7 0 26 5 0 f 0 e 1c 1c 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
D-256 - 800 0 0 1 c 0 1a007000 Eth1/8 0 2f 7 0 10 0 f 1e 1e 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 D-
199 - 800 0 0 1 2e 0 1a01e000 Eth1/31 1 0 2d 0 37 1 e 1c 9c 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 - -
0 0 0 1 0 0 1a01f000 Eth1/32 1 0 3d 0 38 1 f 1e 9e 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 - - 0 0 0 1 0
0 1a030000 Eth1/49 0 2 1 0 49 1 20 38 b8 1 0 0 0 0 0 0 0 0 0 0 1 6 4 2 2 D-24d - 400 0 0 0 1 0
1a031000 Eth1/50 0 3 3 0 29 1 0 0 80 1 0 0 0 0 0 0 0 0 0 0 1 5 3 2 2 D-350 - 400 0 0 0 1 0
```

交换机认为应该转发它到在接口Eth1/49的脊椎。但是如何能验证encap是正确？

我们首先需要查看关于通道的硬件信息。我们能通过运行此HAL命令执行此：

```
module-1(DBG-TAH-elam-insel6)# show platform internal hal tunnel rtep pi Non-Sandbox Mode
LEGEND: ----- Tun Ifid: Tunnel Ifid IfName: Tunnel If Name Lid: Logical Id ET: Encap Type V:
Vxlan I: IVxlan N: NVGRE VrfId: Vrf Id Vrf Name: Vrf Name IP: Tunnel's IP Hw Enc: Hw Encap Idx
IVP: Is VPC Peer IL: Is Local P4: Proxy for v4 P6: Proxy for V6 PM: Proxy for Mac II: Is Ingress
Only IC: Is Copy Service C OBD: Copy Service Outer Bd U D: Use DF NBT: Next Base Type E: ECMP N:
Next-Hop NB Id: Next Base Id NH cnt: Next Hop Count VrfId: Vrf Id Vrf Name: Vrf Name IP: IP
Address Mac: Mac L3 IfId: L3 IfId L3IfName: L3 If Name L2 IfId: L2 IfId L2IfName: L2 If Name
Num. of Sandboxes: 1 Sandbox_ID: 0, BMP: 0x0 Remote Tep Count: 15
=====
===== I N N | E Vrf Hw V I P P P I I C U B B NH | Vrf L3 L3 L2 L2 IfId Ifname T Lid VrfId Name
IP Enc P L 4 6 M I C OBD D T Id Cnt | VrfId Name IP Mac IfId IfName IfId IfName
=====
===== 18010002 Tunnel2 I 3005 2 overlay-1 192.168.120.670 0 0 0 0 0 0 0 1 0 E 2 2 2 overlay-1
0.0.0.0 0d:0d:0d:0d:0d:00 1a030001 Eth1/49.1 1a030000 Eth1/4 9 2 overlay-1 0.0.0.0
0d:0d:0d:0d:0d:00 1a031002 Eth1/50.2 1a031000 Eth1/5 0
```

此输出给我们我们关心的一些个值：

Ifld -接口ID分配到通道

IP -目的地的IP。这应该匹配ELTMC。

L3 Ifld -交换机能使用转发到适当的目的地第3层接口。

一旦我们认识Ifld，我们能验证我们在elam获得匹配隧道目的地的encap：

```
module-1(DBG-TAH-elam-insel9)# show platform internal hal tunnel rtep apd Non-Sandbox Mode
LEGEND: ----- ifId: Interface Id IP: IP address HwVrfId: Hardware Vrf Id SrcTepIdx: Source Tep
Index BDxlate: Egress BDxlate DstInfoIdx: Destination info index RwEncapIdx: Rw Encap Index
```

```


ECMPIdx: ECMP Index Num: Number of hops ECMPMbrIdx: ECMP member Index L2 Index: L2 Index
RwDmacIdx: Rw Dmax Index Num. of Sandboxes: 1 Sandbox_ID: 0, BMP: 0x0 Remote Tep Count: 15
=====
===== ifId IP HwVrfId BDxlate SrcTepIdx DstInfoIdx RwEncapIdx ECMPIdx
ECMPMbrIdx Num L2Index RwDmacIdx
=====
===== 18010002 192.168.120.67 2 1 3a9a 3005 6 0 0 2 1a030000 0 <----
RwEncapIdx is 6! Same as the "encap_idx" in the ELAM Report. 1a031000 1

```

此通道有RwEncapIdx (重写Encap索引) 6，是什么在elam显示。

## 1个EP --> L3外的已路由流

### 拓扑



在本例中，我们将跟踪一数据包的数据包流从发送ICMP的EP1的对在运行OSPF的N5K的一环回。N5K通过在一个对的L3Out连接EX交换机。

因为我们验证编程在本文初的本地EP，假设EP在硬件方面正确地学习和继续到路由验证。

首先，请检查OSPF状态和路由表：

```

leaf6# show ip ospf neighbors vrf jr:sb OSPF Process ID default VRF jr:sb Total number of
neighbors: 2 Neighbor ID Pri State Up Time Address Interface 27.27.27.1 1 FULL/BDR 00:22:39
10.10.27.1 Vlan28 <---- Leaf5 27.27.27.3 1 FULL/DROTHER 00:22:37 10.10.27.3 Vlan28 <---- N5K
leaf6# show ip route vrf jr:sb 100.100.100.100 IP Route Table for VRF "jr:sb" '*' denotes best
ucast next-hop '**' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>'
in via output denotes VRF <string> 100.100.100.100/32, ubest/mbest: 1/0 *via 10.10.27.3, vlan28,
[110/5], 00:16:58, ospf-default, intra

```


因此我们知道routing表表示下一跳作为5K在10.10.27.3。好开始，但是如何能验证什么硬件有？





## 1个EP -->远程EP或SVI -脊椎验证

### 拓扑



### 逻辑

在本例中，我们将跟踪一数据包的数据包流从EP1的被注定对远程BD交换的虚拟接口(SVI)。此示例目的将验证脊椎转发保证数据包发送对正确分支。假设数据包发送给入口分支的脊椎代理。

在脊椎，请首先验证委员会Oracles协议(小屋)目的地IP的，因为数据包发送给查找的脊椎代理：

```
cal01-spine1# show coop internal info ip-db | grep -A 10 192.168.20.1 <----- IP address :  
192.168.20.1 Vrf : 2129921 Flags : 0 EP vrf vnid : 2129921 EP IP : 192.168.20.1 Publisher Id :  
10.0.224.88 Record timestamp : 11 04 2016 16:41:16 422062712 Publish timestamp : 11 04 2016  
16:41:16 424633605 Seq No: 0 Remote publish timestamp: 01 01 1970 00:00:00 0 URIB Tunnel Info  
Num tunnels : 1 Tunnel address : 10.0.224.88 <---- REMOTE LEAF Tunnel ref count : 1
```

请验证什么分支有该TEP地址：

```
spine1# acidiag fmvread | grep 10.0.224.88 105 1 cal01-leaf5 FDO20160TPS 10.0.224.88/32 leaf  
active 0
```


因为我们知道数据包进入在模块2的脊椎，波尔特6，我们能附加到模块2并且查看波尔特布局。











## 逻辑

有我们捉住一数据包不安排在“显示平台内部hal I2内部端口pi的” Ovector表的一些方案。在下面的方案中，我们实际上捉住回来从FM的数据包，因此我们需要查看一个不同的表发现哪个前面板端口数据包选择。

注意以上的拓扑是了解中转流量的一个完全不同的环境(没有代理路由)。模块是N9K-X9732C-EX。

```
@module-1# debug platform internal tah elam asic 3
@module-1(DBG-elam)# trigger reset
@module-1(DBG-elam)# trigg init in-select 13 out-select 0
@module-1(DBG-elam-insel13)# set inner ipv4 src_ip 192.85.1.2 dst_ip 192.85.2.67
@module-1(DBG-elam-insel13)# star
@module-1(DBG-elam-insel13)# stat
ELAM STATUS
=====
Asic 3 Slice 0 Status Armed
Asic 3 Slice 1 Status Triggered

@module-1(DBG-elam-insel13)# report | grep ovector
```



```

L3 Tid      Tid      Lbl  Lbl  | S V | ID  I
=====
=====
lf5         SpInBndMgmt 0 9de  1a   0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 D-2d4     D-3e1  0   0   0 0   1   0
1a000000   Eth1/1   0 1b  1c   0 11 0 10 20 20 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1
1 D-13b     D-33b  500  0   1 0   3   0
1a01c000   Eth1/29  0 37  1e   3 3d 1 14 28 a8 1 0 0 0 0 0 0 0 0 0 0 1 8 8 1
1 D-3f2     D-7a   100  0   0 0   2   0
1a01d000   Eth1/30  0 38  20   3 39 1 10 20 a0 1 0 0 0 0 0 0 0 0 0 0 1 5 5 1
1 D-36e     D-362  100  0   0 0   2   0 <<<<<<< 1/30 is the phys interface that connects to
leaf 102, verified by topology, ASIC 3, Slice 1
1a01e000   Eth1/31  0 39  22   3 35 1 c 18 98 1 0 0 0 0 0 0 0 0 0 0 1 9 9 1
1 D-273     D-8    100  0   0 0   2   0
1a01f000   Eth1/32  0 3a  24   3 31 1 8 10 90 1 0 0 0 0 0 0 0 0 0 0 1 a a 1
1 D-154     D-5d   100  0   0 0   2   0

```