

CISCO *Live!*

ALL IN

#CiscoLive



The bridge to possible

# A Network Engineer's Blueprint for ACI Forwarding

## Part 2 – Debugging ACI Forwarding

Joseph Young, ACI Technical Leader – Customer Experience  
BRKDCN-3900b

CISCO *Live!*

#CiscoLive

# Cisco Webex App


## Questions?

Use Cisco Webex App to chat with the speaker after the session

## How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 17, 2022.



<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-3900b>



# Agenda

- Understanding the Tools
  - UI Tools
  - Elam
  - Friage
  - Span / ERSPAN
  - Flow Telemetry / netflow
- Debugging and Walking Through ACI Flows
  - (Routed, Bridged, BUM, Proxied)


# Glossary of Acronyms

## VxLAN packet acronyms

Acronyms	Definitions
ACI	Application Centric Infrastructure
APIC	Application Policy Infrastructure Controller
EP	Endpoint
EPG	Endpoint Group
BD	Bridge Domain
VRF	Virtual Routing and Forwarding
COOP	Council of Oracle Protocol
VxLAN	Virtual eXtensible LAN

Acronyms	Definitions
dXXXo	Outer Destination XXX (dIPo = Outer Destination IP)
sXXXo	Outer Source XXX (sIPo = Outer Source IP)
dXXXi	Inner Destination XXX (dIPi = Inner Destination IP)
sXXXi	Inner Source XXX (sIPi = Inner Source IP)
GIPO	Outer Multicast Group IP
VNID	Virtual Network Identifier

# Understanding the Tools



# Start with High-level Tools

## Use Endpoint Tracker for Building a Topology

### EP Tracker

End Point Search

EP Locally Learned on pod 2, nodes 401-402

172.16.31.100

Search

Learned At	Tenant	Application	EPG	IP
2/401-2/402, vPC: vpc-esxi-10.2.10.19 (learned,vmm)	CiscoLive	Database	DB	172.16.31.100

End Point Search

No EP Learn, is this an L3out?

10.255.255.100

Search

Learned At	Tenant	IP
No items have been found.		

# Start with High-level Tools

## Use Atomic Counters to Check for Overlay Drops and Latency (PTP)

### Add EP to EP Policy



Name:

Description:

Administrative State:  Disabled  Enabled

Features:  Atomic Counter  
 Latency Statistics

Source Type:  EP  IP

Source IP:

Application Profile EPG/ESG Client Endpoint Internet Protocol

Destination IP:

Application Profile EPG/ESG Client Endpoint Internet Protocol

Filters:

Name	Protocol	Source port	Destination port	Description
ip	Unspecified	Unspecified	Unspecified	



# Start with High-level Tools

Use Atomic Counters to Check for Overlay Drops and Latency (PTP)

The screenshot displays the CiscoLive interface with a navigation menu on the left and two data tables on the right. The top table, 'EP-to-EP Atomic Counter - CL-AC', shows a 'Dropped' count of 0. The bottom table, 'EP-to-EP Latency Average - CL-AC', shows an 'Average(μs)' of 104.8575. Two callout boxes highlight these values: '104 Microseconds of delay in overlay' points to the latency average, and 'No overlay drops!' points to the zero dropped count.

**EP to EP CL-AC**

EP-to-EP Atomic Counter - CL-AC

Source	Destination	Last Collection (30 seconds) Pkt			
		Transmit	Admitted	Dropped	Excess
uni/tn-CiscoLive/ap-Databas...	uni/tn-CiscoLive/ap-APP/epg...	29	29	0	0

**EP-to-EP Latency Average - CL-AC**

Last 30 Seconds Collection 04/25/2022 16:06:05			Cumulative (04/25/2022 15:04:45 - 04/25/2022 16:06:05)		
Average(μs)	Standard Deviation(μs)	Packet Count	Average(μs)	Max(μs)	Packet Count
104.8575	0.0000	29	104.8575	104.8575	3768

# Start with High-level Tools

Use Tenant Visibility tools to check for Contract Drops

The screenshot shows the CiscoLive interface for a tenant. The left sidebar contains navigation options: Quick Start, CiscoLive, Application Profiles, Networking, Contracts, Policies, Services, and Security. The main content area is titled 'Tenant - CiscoLive' and has several tabs: Summary, Dashboard, Policy, Operational (selected), Stats, Health, Faults, and History. Under the 'Operational' tab, there are sub-tabs: Endpoints, Flows, Packets (selected), Policy Tags, and Resource IDs. Further down, there are more sub-tabs: L2 Permit, L3 Permit, L2 Drop, and L3 Drop (selected). A table displays a list of dropped packets. The first row is highlighted with an orange box, and a callout box with the text 'This flow is being contract dropped' points to it. The table columns are: Timestamp, VRF, Src IP, Dest IP, Protocol, Src Port, Dest Port, and Node.

Timestamp	VRF	Src IP	Dest IP	Protocol	Src Port	Dest Port	Node
2022-04-25T17:19:44.070+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402
2022-04-25T17:19:39.430+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402
2022-04-25T17:18:53.350+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402
2022-04-25T17:11:12.545+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402
2022-04-25T17:18:52.870+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402
2022-04-25T17:18:52.326+00:00	CustA	172.16.31.100	10.255.255.100	icmp	unspecified	unspecified	node-402

```
apic4# show aclog deny l3 pkt tenant common vrf CORE
srcIp dstIp protocol srcPort dstPort node srcIntf vrfEncap
-----
<EMPTY>
```

# Start with High-level Tools

## Port Counters are as Useful as Ever

```
leaf1# show interface eth1/8
Ethernet1/8 is up
admin state is up, Dedicated Interface
Last link flapped 03:07:41
RX
 3527922 unicast packets !omitted
 4041582 input packets 609518993 bytes
 12 jumbo packets 0 storm suppression bytes
 0 runts 0 giants 0 CRC 0 Stomped CRC 0 no buffer
 0 input error 0 short frame 0 overrun !omitted
 0 watchdog 0 bad etype drop 0 bad proto drop !omitted
 0 input with dribble 0 input discard
 0 input buffer drop 0 input total drop
TX
 32262479565 unicast packets !omitted
 32395063346 output packets 49034781261
 32249687943 jumbo packets
 0 output error 0 collision 0 deferred
 0 lost carrier 0 no carrier 0 babble 0 output discard
 0 output buffer drops 0 output total drops
```

Frames received with bad FCS

Indicates a previously stomped frame was received

What is a Stomp?

- When a frame is received with a bad FCS and/or is malformed

AND

- The frame is cut-through switched

The switch will invert the new CRC to tell the first store-and-forward device to drop it

Frame transmitted with stomped CRC

Buffer drops, sign of congestion

# Start with High-level Tools

Using moquery to check port counters fabric-wide

## #Check Fabric-wide for FCS Errors

```
moquery -c rmonDot3Stats -f 'rmon.Dot3Stats.fCSErrors>="1"' | egrep "dn|fCSErrors"
```

## #Check Fabric-wide for total CRC Stomp + FCS Errors

```
moquery -c rmonEtherStats -f 'rmon.EtherStats.cRCAlignErrors>="1"' | egrep "dn|cRCAlignErrors"
```

## #Check Fabric-wide for Output Buffer Drops

```
moquery -c rmonEgrCounters -f 'rmon.EgrCounters.bufferdropPkts>="1"' | egrep "dn|bufferdropPkts"
```

## #Check Fabric-wide Output Errors

```
moquery -c rmonIfOut -f 'rmon.IfOut.errors>="1"' | egrep "dn|errors"
```

# ELAM - Embedded Logic Analyzer Module

- It is a tripwire in hardware
- The first frame to match a specified condition 'trips' it
- Report is created with vast amount of data regarding asic decisions

Dst - TCP 10.0.0.1:3000

Dst - TCP 10.0.0.1:3001

Dst - TCP 10.0.0.1:3002



```
vsh_lc
debug platform internal tah elam asic 0
trigger reset
trigger init in-select 6 out-select 1
set outer ipv4 dst_ip 10.0.0.1
set outer 14 dst-port 3001
start
```

Frame was not dropped in lookups!

```
module-1(DBG-elam-insel6)# stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Armed
Asic 0 Slice 1 Status Triggered

module-1(DBG-elam-insel6)# ereport | grep "drop reason"
RW drop reason           : no drop
LU drop reason           : no drop
```

Matching frame was caught!

# What ASIC should be set in the ELAM?

```
vsh_lc  
debug platform internal <asic> elam asic 0
```

Model	Role	Asic for Elam
N9K-C*C	Fixed Spine	roc
N9K-C*GX	Fixed Spine	app
N9K-C*-EX	Leaf	tah
N9K-C*-FX/FXP/FX2	Leaf	roc
N9K-C*-GX	Leaf	app
N9K-C*-GX2	Leaf	cho
N9K-X97*-EX	Spine LC	tah
N9K-X97*-FX	Spine LC	roc
N9K-X97*-GX	Spine LC	app
N9K-C95*-FM-E	Spine FM	tah
N9K-C950*-FM-E2	Spine FM	roc
N9K-C95*-FM-G	Spine FM	app

# Steps to Using Elam on Gen2+ Leaf or Fixed Spine

Elams are run from the line card shell

Refer to "What ASIC should be set in the ELAM" slide

Leafs and fixed spines are single ASIC switches. Always use ASIC 0

```
vsh_lc
```

```
debug platform internal tah elam ASIC 0
```

```
trigger reset
```

```
trigger init in-select 6 out-select 0
```

```
set outer ipv4 dst_ip 10.0.0.1
```

```
set outer 14 dst-port 3001
```

```
start
```

Failing to reset the trigger can cause past elam configurations to take effect. Always reset the trigger!

Use 0 or 1

```
module-1 (DBG-elam) # trigger init in-select ?  
!omitted  
14 Outer(12(vntag)|13|14)-inner(12|13|14)-ieth  
6 Outer12-outer13-outer14  
7 Inner12-inner13-inner14  
!omitted
```

Determines which headers conditions can be matched in. Use 14 or 7 when matching vxlan encapsulated headers.

# Steps to Using Elam on Gen2+ Leaf or Fixed Spine

Use "set outer" or "set inner" depending on in-select and if matching outer or inner headers in vxlan packet

Which headers to match conditions for?

```
vsh_lc
debug platform internal tah elam asic 0
trigger reset
trigger init in-select 6 out-select 0
set outer ipv4 dst_ip 10.0.0.1
set outer 14 dst-port 3001
start
```

What to match in the header?

Finally enable the elam!

When running `stat` if `Triggered` is seen, this means a matching packet was received



# Reading an Elam

ereport available since 4.2

At a high-level...

```
module-1 (DBG-elam-inse16) # ereport
!omitted
-----
Outer L3 Header
-----
L3 Type           : IPv4
IP Version        : 4
DSCP              : 0
IP Packet Length  : 84 ( = IP header(28 bytes) + IP payload )
Don't Fragment Bit : set
TTL               : 64
IP Protocol Number : ICMP
Destination IP    : 192.168.200.11
Source IP         : 192.168.100.10
!omitted
Contract Result
Contract Drop     : no
Contract Logging  : no
Contract Applied  : yes
Contract Hit      : yes
```

- ereport provides a simple, human-readable report output
- ereport requires  $\geq 5.2$  code for modular spines
- Groups data into outer/inner, headers, and lookup results

# Reading an Elam

ereport available since 4.2

At a low-level...

```
report detail | grep -F "-----" | grep -v VECTOR | grep -v end
LU BEGIN -----
LUA -----
LUB -----
LUC -----
LUD -----
LU END -----
*** FP latch results -----
*** LBX latch results -----
*** ACX latch results -----
RW BEGIN -----
RW END -----
```

- An elam report provides a walkthrough of each ASIC block
- Each decision in each block is recorded
- Refer to “Inside an ACI Switch ASIC” from part 1 for more details
- All output is in HEX

# What if Elam Shows a Drop?

ereport available since 4.2

## ereport

Lookup Drop

-----  
LU drop reason : SECURITY\_GROUP\_DENY

## Common Drop Reasons

Drop Code	What Does it Mean?	What to Do?
ACL_DROP	For traffic destined to the CPU on an FX switch it is expected and cosmetic. Also seen when traffic was received from a fabric port and the leaf has a remote EP learn with no bounce flag.	Ignore if its an FX switch and destined to local switch IP/process. Otherwise, check for incorrect EP learn.
DCI_*_XLATE_MISS	For multisite / remote-leaf, there was no matching vnid or ptag translation found.	Check contracts between local and remote resources.
INFRA_ENCAP_SRC_TEP_MISS	No route and/or tunnel found back to the outer source IP	Check for a tunnel pointing back to the outer source IP. Also, check for a route in overlay.
SECURITY_GROUP_DENY	<b>Frame was contract dropped</b>	<b>Make sure a contract is configured to allow the flow.</b>
SRC_VLAN_MBR	Received vlan not programmed on ingress port.	Check if the frame was correct tagged/untagged. Make sure no invalid-path faults exist for the epg.
UC_PC_CFG_TABLE_DROP	No route was found for the destination.	Check the routing table for the destination.
VLAN_XLATE_MISS	Received vlan doesn't exist on the switch.	Check if the frame is tagged with correct vlan. Check for invalid-path faults on the epg.

# Steps to Using Elam on Gen2+ Modular Spine

## Challenges of Modular Spines

- Line cards (and potentially FM's) have multiple asics
- Elam must specify asic number
- Ingress/Egress ports may be internal LC – FM connections
- ereport only available in 5.2 and later

Fortunately, spine elams aren't needed as commonly as leaf elams!

# Steps to Using Elam on Gen2+ Modular Spine

Ingress LC



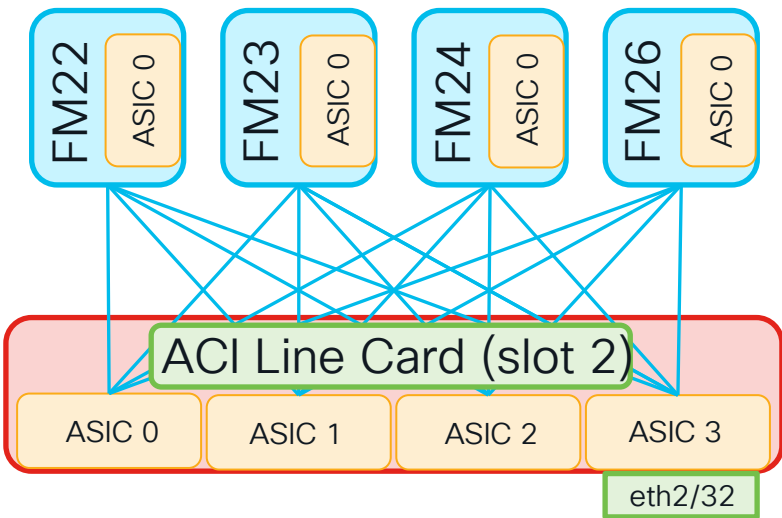
Determine the Asic, Slice, and Srcid of the ingress port

1

```
sp# vsh
sp# attach mod 2
module-2# show plat internal hal 12 port gpd
```

IfId	Ifname	P	Cfg	MbrID	As	AP	S1	Sp	Ss	Ovec
!omitted										
1a09f000	Eth2/32	0	b9	38	3	31	1	8	10	90

Eth2/32 is on Asic 3, Slice 1, with srcid 0x10. Use for Elam!

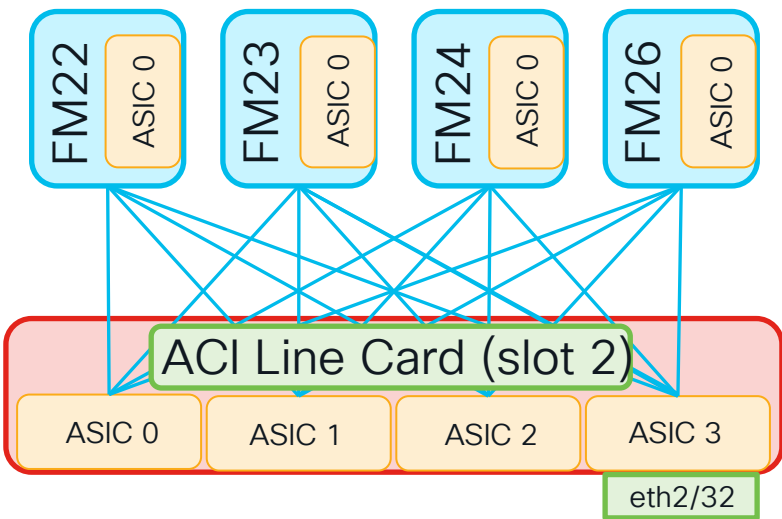


Ingress Traffic:

Inner Headers -  
Src - 10.10.10.10  
Dst - 10.10.11.11

# Steps to Using Elam on Gen2+ Modular Spine

Ingress LC



Ingress Traffic:  
Inner Headers -  
Src - 10.10.10.10  
Dst - 10.10.11.11

2

```
sp# vsh
sp# attach mod 2
debug plat internal tah elam asic 3 slice 1
trigger reset
trigger init in-select 14 out-select 1
set srcid 0x10
set inner ipv4 src_ip 10.10.10.10 dst_ip 10.10.11.11
start
```

Asic and slice of eth1/32  
(see last slide)

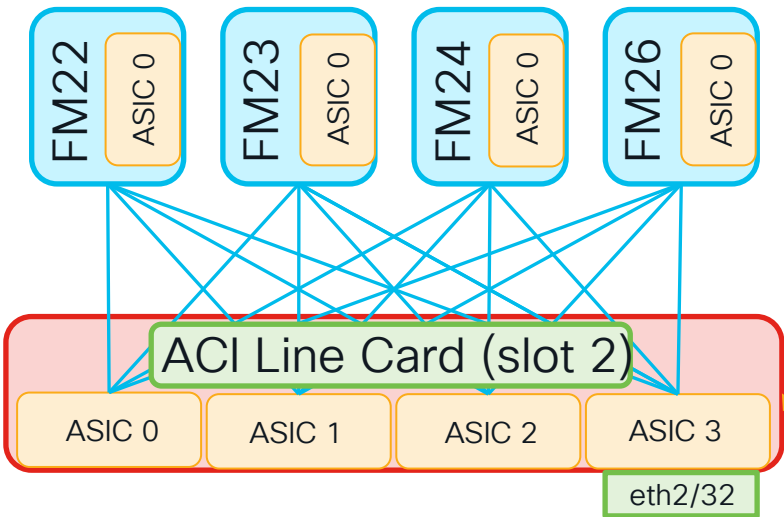
Source ID value of  
eth1/32 (see last slide)

3

```
module-2 (DBG-elam-insel14) # stat
ELAM STATUS
=====
Asic 3 Slice 1 Status Triggered
```

Packet was matched!

# Steps to Using Elam on Gen2+ Modular Spine



Ingress Traffic:  
 Inner Headers -  
 Src - 10.10.10.10  
 Dst - 10.10.11.11

5

```
module-2# show plat internal hal 12 internal-port pi
=====
IfId      IfName                                     As  Ovec
=====
96        lc(0)-fc(0):22:pc2:p1                     0  b8
98        lc(1)-fc(0):22:pc2:p1                     1  b8
9a        lc(2)-fc(0):22:pc2:p1                     2  b8
9c        lc(3)-fc(0):22:pc2:p1                   3  b8
```

Packet forwarded to FM 23! (output is zero-based)

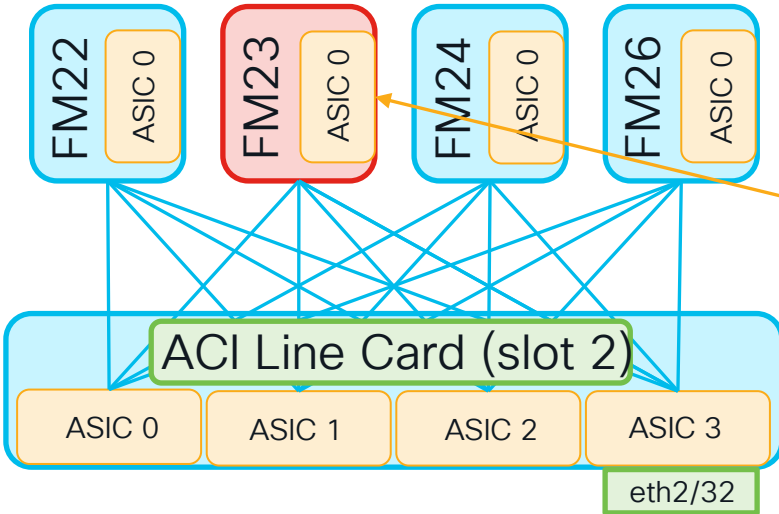
Ovector indicates the egress port to FM

4

```
report | egrep "drop\_vec|ovec|asic"
Dumping report for ASIC inst 3 slice 1 insel 14 outsel 1
*_sidebnd_no_spare_vec. ovector_idx: 0xB8
*_vec.pb_x_header_sidebnd_drop_vec. lux_drop_vec: 0x0000000
```

Packet wasn't dropped in lookups!

# Steps to Using Elam on Gen2+ Modular Spine



Ingress Traffic:  
Inner Headers -  
Src - 10.10.10.10  
Dst - 10.10.11.11

9508 and 9516 FM's have 2 asics; if no trigger on 0, try 1.

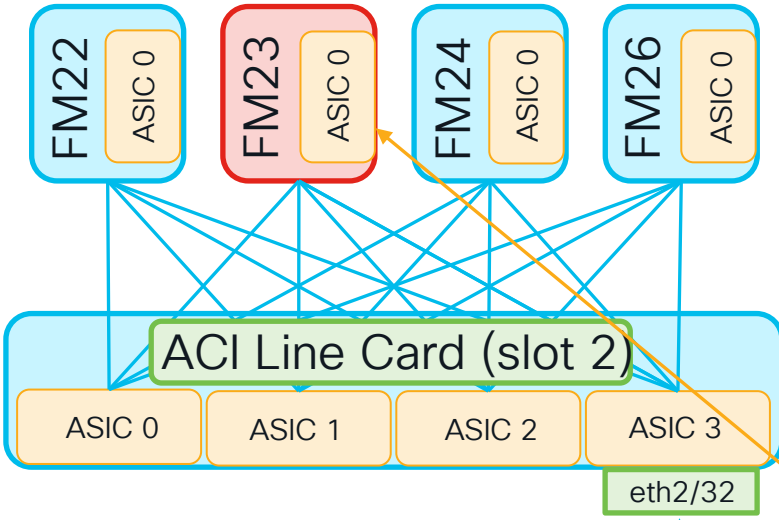
```
6 sp# vsh
sp# attach mod 23
debug plat internal tah elam asic 0
trigger reset
trigger init in-select 14 out-select 1
set inner ipv4 src_ip 10.10.10.10 dst_ip 10.10.11.11
start
```

```
7 module-23(DBG-elam-insel14)#
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Armed
Asic 0 Slice 1 Status Triggered
Asic 0 Slice 2 Status Armed
Asic 0 Slice 3 Status Armed
Asic 0 Slice 4 Status Armed
Asic 0 Slice 5 Status Armed
```

Packet was matched!



# Steps to Using Elam on Gen2+ Modular Spine



Ingress Traffic:  
 Inner Headers -  
 Src - 10.10.10.10  
 Dst - 10.10.11.11

```

9 module-23# show plat internal hal 12 port gpd
-----
IfId      Ifname      As  Ovec
-----
f5        fc0-1c1:3-1 0  58
    
```

Packet forwarded to LC 2  
 (zero based - Asic 3, Slice 1)

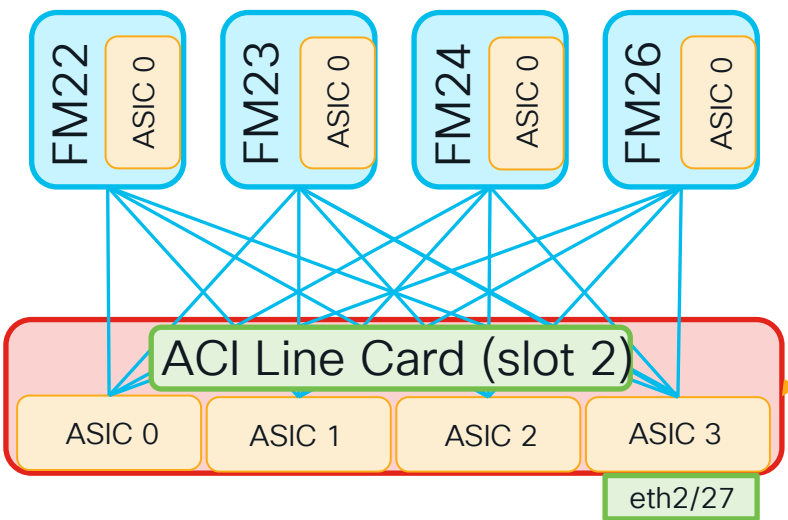
Ovector indicates the egress port to LC

```

8 report | egrep "drop\_vec|ovec|asic"
Dumping report for asic inst 0 slice 1 inseq 14 outseq 1
*_sidebnd_no_spare_vec ovector_idx: 0x58
*_vec.pbx_header_sidebnd_drop_vec.lux_drop_vec: 0x000000000000
    
```

Packet wasn't dropped in lookups!

# Steps to Using Elam on Gen2+ Modular Spine



Inner Headers -  
Egress Traffic: Src - 10.10.10.10  
Dst - 10.10.11.11

```
10 sp# vsh
sp# attach mod 2
debug plat internal tah elam asic 3 slice 1
trigger reset
trigger init in-select 14 out-select 1
set outer 12 vntag_vld 1
set inner ipv4 src_ip 10.10.10.10 dst_ip 10.10.11.11
start
```

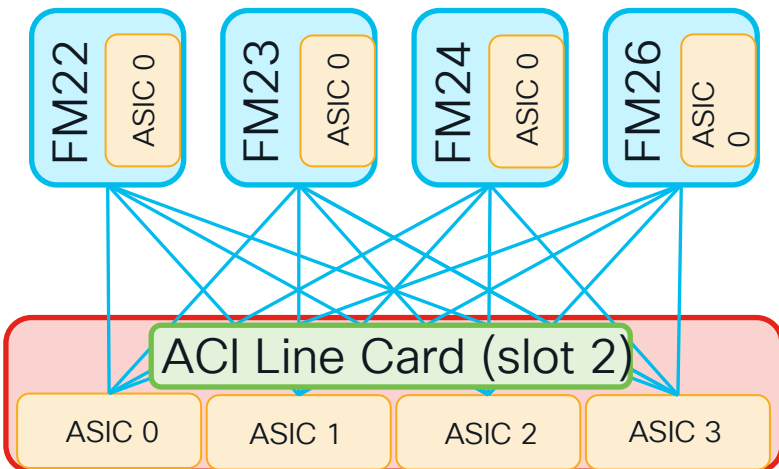
Asic 3 / slice 1 as seen on last slide

Vntag present only coming from FM

```
11 module-2 (DBG-elam-insel14) # stat
ELAM STATUS
=====
Asic 3 Slice 1 Status Triggered
```

Packet was matched!

# Steps to Using Elam on Gen2+ Modular Spine



eth2/27

Egress Traffic: Inner Headers -  
Src - 10.10.10.10  
Dst - 10.10.11.11

```
13 module-2# show plat internal hal 12 port gpd
=====
IfId      Ifname      As AP Sl Sp Ss Ovec
=====
1a08a000 Eth2/11     1  5  0  4  8  8
1a09a000 Eth2/27  3 5  0  4  8  8
```

Spine forwards out front-panel Eth2/27!

```
12 report | egrep "drop\ _vec|ovec|asic"
Dumping report for asic inst 3 slice 1 insel 14 outsel 1
*_sidebnd_no_spare_vec. ovector_idx: 0x8
*_vec.pbh_header_sidebnd_drop_vec. lux_drop_vec: 0x00000000
```

Ovector indicates the egress port to Leaf

Packet wasn't dropped in lookups!

# Automating Modular Spine ELAMs

CLI-based Modular Spine Elam tool available at – [EasySpineElam](https://github.com/CiscoSpine/easy-spine-elam)

Easily Set Conditions on  
All or Some Modules

```
spine1#./easy-spine-elam.sh -m all -d ingress
Final module list is:
2 23 26 3
2022-06-08T14:55:57 In-select - 14 and out-select - 0 are being used.
!ommitted
70. inner ipv4 destination ip > Format : d.d.d.d
71. inner ipv4 protocol > Format : 0-255
73. inner ipv4 source ip > Format : d.d.d.d
91. inner l4 dest port > Format : 0-65535

Select corresponding numbers of conditions to set. Separate numbers with commas.
Ex: 1,2,3,4,5
```

```
Enter selections: 70,73,71,91
```

Which conditions to match?

```
Enter inner ipv4 destination ip > Format : d.d.d.d: 80.0.0.1
Enter inner ipv4 source ip > Format : d.d.d.d: 150.0.0.100
Enter inner ipv4 protocol > Format : 0-255: 6
Enter inner l4 dest port > Format : 0-65535: 8989
```

Set conditions

# Automating Modular Spine ELAMs

CLI-based Modular Spine Elam tool available at – [EasySpineElam](#)

```
2022-06-08T14:56:28 Checking elam status for module 2
2022-06-08T14:56:28 Checking elam status for module 23
2022-06-08T14:56:28 Checking elam status for module 26
2022-06-08T14:56:28 Checking elam status for module 3
```

Generate and view ereport  
from all Triggered Modules!

```
ELAM TRIGGERED on module 26:  
ASIC: 0 SLICE: 1
```

ELAM triggered on  
LC and FM!

```
ELAM TRIGGERED on module 2:  
ASIC: 3 SLICE: 1
```

```
Type "status" to check elam status again. Type "ereport", "report" or "report detail"  
to collect all reports: ereport
```

```
2022-06-08T14:57:36 Collecting report for module 26 asic 0...
2022-06-08T14:57:36 Collecting report for module 2 asic 3...
2022-06-08T14:57:46 Converting reports to ereport format!
```

Locally view or copy  
off the final ereports

```
The following decoded elams are available -
```

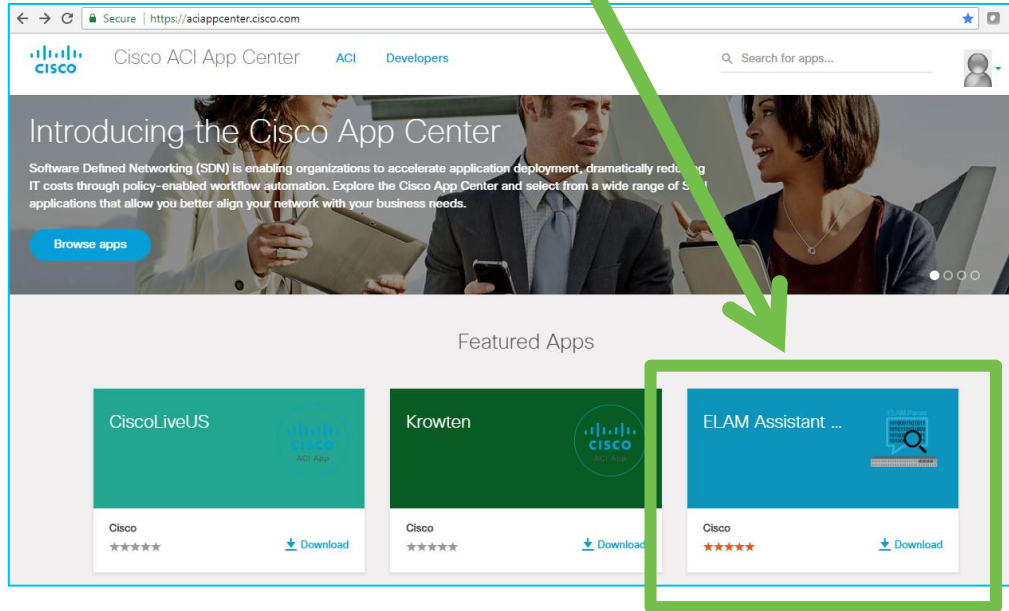
```
/data/techsupport/mod26-asic0-elamreport-2022-06-08T14-57-36-EREPORT  
/data/techsupport/mod2-asic3-elamreport-2022-06-08T14-57-36-EREPORT
```

```
2022-06-08T14:57:49 FINISHED!
```

# Shouldn't ELAM be More Simple?

## Elam Assistant in DCAppCenter

<https://dcappcenter.cisco.com>



### ELAM (Embedded Logic Analyzer Module)

- Perform an ASIC level packet capture

### ELAM Assistant

- You can perform ELAM like a TAC engineer!
- With a nicely formatted result report

### Detail Explanations:

- <https://dcappcenter.cisco.com/elam-assistant.html>
- How to use video, pictures
  - A download link for ELAM Assistant

# ELAM Assistant in ACI AppCenter (example)

## 1. Perform an Elam

The screenshot shows the ACI AppCenter interface with the 'ELAM Assistant' application selected. The main content area is titled 'ELAM Parameters' and contains a table of configurations. A callout box points to the 'Set Parameters' button in the table.

Status	Node	Direction	Source I/F	Parameters	VxLAN (outer) header
Set	node-401	from downlink	any	dst ip 10.255.255.100	
Report Ready	node-402	from downlink	any	dst ip 10.255.255.100	
Report Ready	node-303	from LEAF/IPN	any	dst ip 10.255.255.100	

Buttons: Quick Add, Add Node, Set ELAM(s), Check Trigger

ELAM Report Parse Result ( report name: )

Express Detail Raw

Select a report.

Triggered!!  
and  
Report is Ready

Set Parameters

# ELAM Assistant in ACI AppCenter (example)

## 2. Read a Report

The screenshot displays the ACI AppCenter ELAM Assistant interface. On the left, a sidebar lists capture nodes: node-101, node-102, node-303, node-402, and Unsupported Nodes. The main area is titled 'ELAM Parameters' and includes a 'Name your capture' field, a table of configurations, and a 'Set ELAM(s)' button. Below this is the 'ELAM Report Parse Result' section with tabs for 'Express', 'Detail', and 'Raw'. The 'Express' tab is active, showing 'Captured Packet Information' and 'Basic Information'. A callout box highlights the 'Packet Forwarding Information' section, which contains details about the forward result, contract, and drop status. Another callout box points to the 'Set' button in the configuration table, and a third points to the 'Express' tab. A 'Scroll Down' callout is also present.

**ELAM Parameters**

Name your capture

Status	Node	Direction	Source I/F	Parameters
Set	node-401	from downlink	any	
Report Ready	node-402	from downlink	any	
Report Ready	node-303	from LEAF/IPN	any	

**ELAM Report Parse Result ( report name: node-402\_slot1\_asic...)**

Express Detail Raw

**Captured Packet Information**

**Basic Information**

Device Type	LEAF
Packet Direction	ingress (from downlink)
Incoming I/F	eth1/4

**L2 Header**

Destination MAC	0022.BDF8.19FF
Source MAC	0050.569A.65DB
Access Encap VLAN	844

**Packet Forwarding Information**

**Forward Result**

Destination Type	To another ACI node (LEAF, AVS/AVE etc.)
Destination TEP	10.1.240.33 (MAC Spine-Proxy)
Destination Physical Port	eth1/49

**Contract**

Destination EPG pcTag (dclass)	0x4002 / 16386 (L3OUT CiscoLive:L3out-CUST:EPEG2)
Source EPG pcTag (sclass)	0x8005 / 32773 (CiscoLive:Database:DB)
Contract was applied	1 (Contract was applied on this node)

**Drop**

Drop Code	no drop
-----------	---------

Click to see report

Report shows up here

Scroll Down



# FTRIAGE – Automating Elams

Orchestrate End-to-End  
ELAMs from the APIC!

```
apic1# ftriage route -ii LEAF:101,102 -dip 10.99.99.100 -sip 192.168.100.10
20:19:54 INFO main:1295 L3 packet Seen on leaf102 Ingress: Eth1/34 (Po5) Egress: Eth1/54 Vnid: 2523136
20:19:55 INFO main:1364 leaf102: Packet's egress outer [SIP:10.0.176.67, DIP:10.0.64.70]
20:19:55 INFO main:1371 leaf102: Outgoing packet's Vnid: 2523136
20:19:56 INFO main:353 Computed ingress encap string vlan-3501
20:20:03 INFO main:464 Ingress BD(s) CL2022:bd1
20:20:03 INFO main:476 Ingress Ctx: CL2022:vrfl Vnid: 2523136
!
20:21:46 INFO main:1295 L3 packet Seen on spine1005 Ingress: Eth1/1 Egress: Eth1/3 Vnid: 2523136
20:22:38 INFO fib:737 spine1005: Transit in spine
20:23:32 INFO main:1295 L3 packet Seen on leaf103 Ingress: Eth1/29 Egress: Eth1/27/4 Vnid: NULL
!
20:24:02 INFO fib:219 leaf103: L3 out interface Ethernet1/27/4
20:24:10 INFO main:781 Computed egress encap string vlan-1055
20:24:17 INFO main:1796 Packet is Exiting fabric with peer-device: N3K-1 and peer-port: Ethernet1/31
```

# SPAN / ERSPAN

Don't neglect old friends!

- Both local span and erspan supported
- ERSPAN requires an I3 endpoint learned anywhere in the fabric
- Still the best tool for checking –
  - Packet contents
  - Frame format
  - Retransmissions
  - ...and anything else that can be seen in a pcap

# Other Tools Requiring External Resources

## Netflow

- Captures flow information based on specified criteria
- Useful for troubleshooting packet loss and latency

## Flow Telemetry

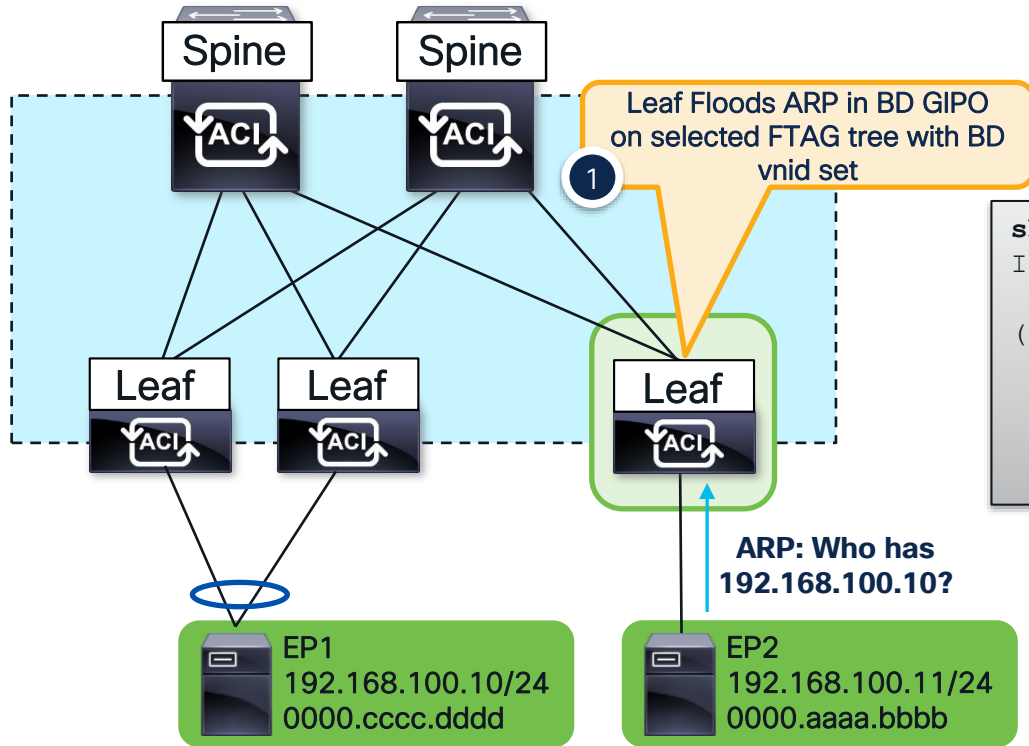
- Hardware directly streams flow data to Nexus Dashboard Insights
- Useful for troubleshooting packet loss and latency
- Latency measurements leverage PTP for additional accuracy
- NDI can perform additional flow analytics

# Debugging ACI BUM Flows



# ARP - Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled



Check GIPO Route

```
show ip mroute 225.0.2.128 vrf overlay-1
IP Multicast Routing Table for VRF "overlay-1"

(*, 225.0.2.128/32), uptime: 22w2d, isis
Incoming interface: Null, RPF nbr: 0.0.0.0
Outgoing interface list: (count: 2)
  Ethernet1/29.9, uptime: 8w2d
  Ethernet1/30.10, uptime: 22w2d
```

# ARP – How to Find the GiPo

## From the GUI...

The screenshot shows the Cisco APIC GUI with the 'Tenants' tab selected. The left sidebar shows a tree view with 'CL2022' expanded to 'Networking' and then 'Bridge Domains'. The main content area displays a table titled 'Networking - Bridge Domains' with the following data:

Name	Segment	VRF	Multicast Address
bd1	14811121	vrf1	225.0.2.128
bd2	16613259	vrf1	225.0.8.48
bd3	16187328	vrf2	225.0.159.112

## From the APIC CLI...

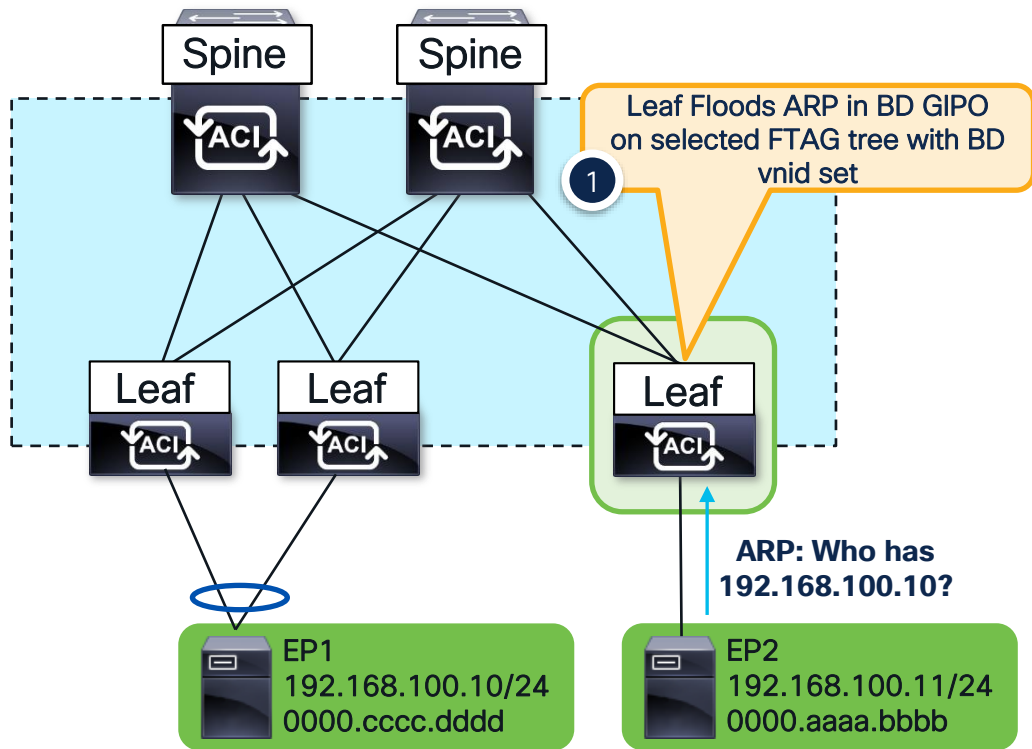
```
moquery -c fvBD -f 'fv.BD.dn*"tn-CL2022/Bd-bd1"'  
  
# fv.BD  
arpFlood : yes  
bcastP : 225.0.2.128  
dn : uni/tn-CL2022/Bd-bd1
```

## From the Switch CLI...

```
moquery -c l2BD -f 'l2.BD.name=="CL2022:bd1"' -x rsp-subtree=full rsp-subtree-class=fmcastGrp  
# fmcast.Grp  
addr : 225.0.2.128  
dn : sys/ctx-[vxlan-2523136]/bd-[vxlan-14811121]/fmgrp-[225.0.2.128]  
rn : fmgrp-[225.0.2.128]
```

# ARP - Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled



ELAM the ARP request!

```
vsh_lc
debug plat internal app elam asic 0
trigger reset
trigger init in-select 6 out-select 0
set outer arp source-ip 192.168.100.11
set outer arp target-ip 192.168.100.10
start
!
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Armed
Asic 0 Slice 1 Status Armed
Asic 0 Slice 2 Status Triggered
Asic 0 Slice 3 Status Armed
```

# ARP - Ingress Leaf Elam Results (ereport)

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled

Outer L2 Header

-----  
Access Encap VLAN : 3502 ( 0xDAE )

Make sure this matches  
what is expected

Outer L3 Header

-----  
ARP Opcode : Request( 0x1 )  
ARP Sender IP : 192.168.100.11  
ARP Target IP : 192.168.100.10

Contract Result

-----  
Contract Drop : no  
Contract Applied : no

**FINAL FORWARDING LOOKUP**

-----  
Bits set in Final Forwarding Block: : IFABRIC\_IG MC TENANT MYTEP **BRIDGE** MISS **FLOOD**

Frame is flooded in the Bridge Domain!

Lookup Drop

-----  
LU drop reason : no drop

Not Dropped in lookups!



# ARP – How to Find the FTAG

No other way than Elam...

```
module-1(DBG-elam-insel6)# ereport | grep "nopad.ftag"  
wol_lu2ba_sb_info.mc_info.mc_info_nopad.ftag: 0x8
```

Selected ftag is 0x8

- Leaf forwards to root port and OIF's for ftag 8
- Since GIPO is 225.0.2.128, Dest multicast address is 225.0.2.136 (gipo + ftag)
- Check ftag topology with **show isis internal mcast routes ftag**

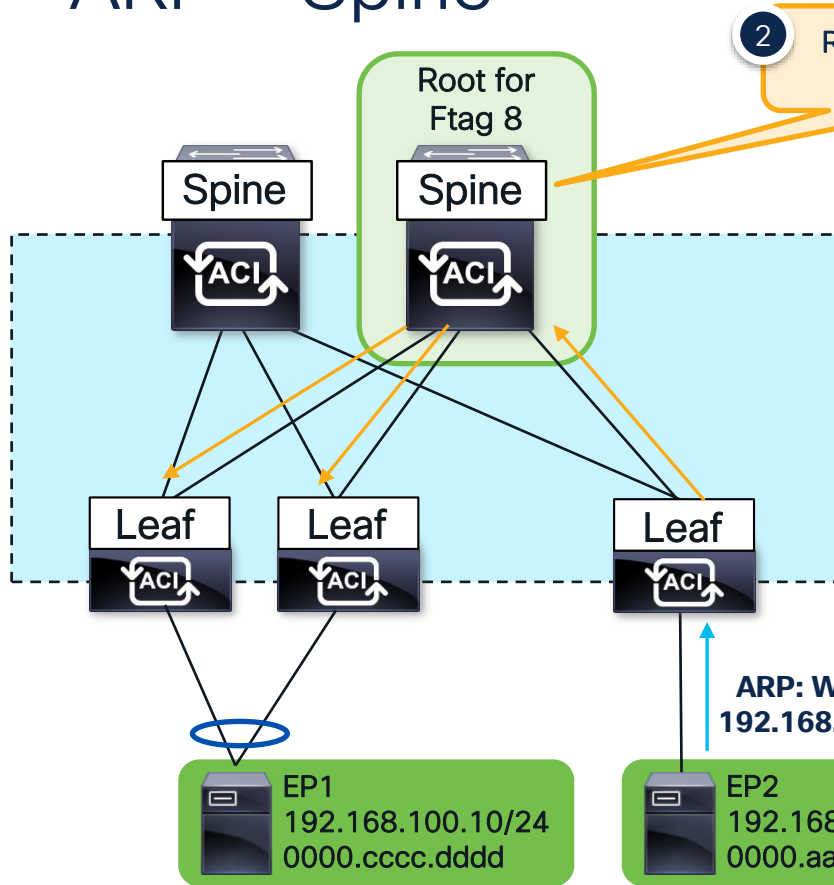
```
leaf103# show isis internal mcast routes ftag  
IS-IS process: isis_infra  
  VRF : default  
FTAG Routes  
=====
```

FTAG ID:	8	[Enabled]	Cost:(	1/	6/	0)
----------	---	-----------	--------	----	----	----

```
-----  
  Root port: Ethernet1/29.9  
  OIF List:
```

Leaf appends ftag to gipo and forwards out Eth1/29 to spine

# ARP - Spine



2 Root spine for ftag 8 forwards out OIFs

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled

This spine is the root!

```
spine1005# show isis internal mcast routes ftag
IS-IS process: isis_infra
VRF : default
FTAG Routes
=====
FTAG ID: 8 [Root] [Enabled] Cost:( 0/ 0/ 0)
-----
Root port: -
OIF List:
Ethernet1/1.20
Ethernet1/2.21
Ethernet1/3.19
```

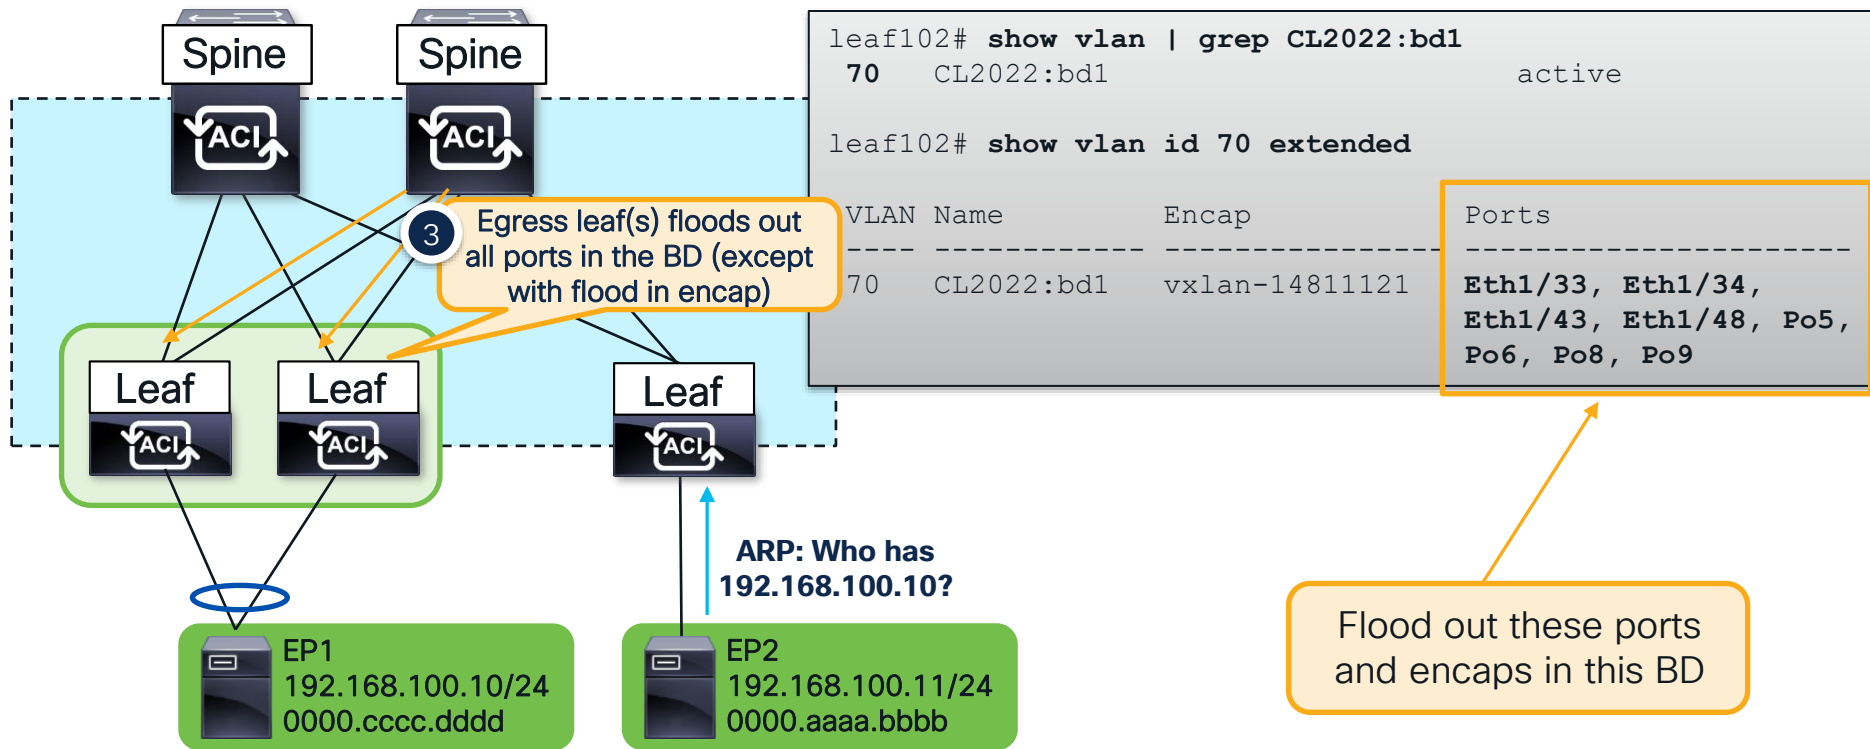
ARP: Who has 192.168.100.10?

EP2  
192.168.100.11/24  
0000.aaaa.bbbb

Spine forwards out OIFs

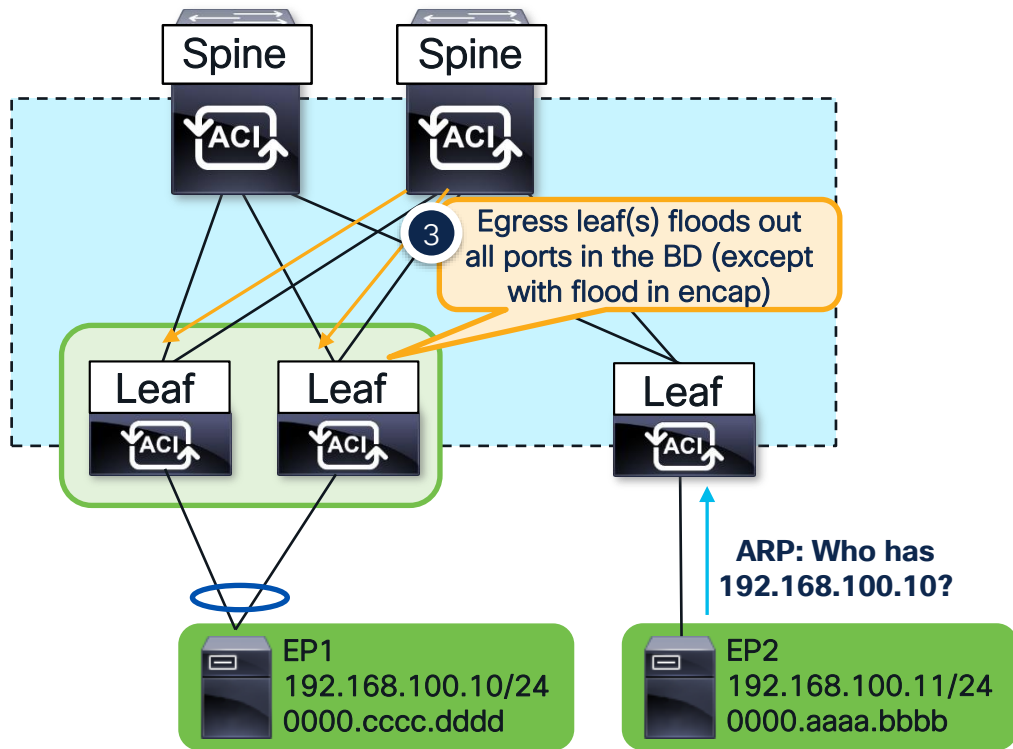
# ARP - Egress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled



# ARP - Egress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled



ELAM the ARP request!

```
vsh_lc
debug plat internal tah elam asic 0
trigger reset
trigger init in-select 14 out-select 1
set inner arp source-ip 192.168.100.11
set inner arp target-ip 192.168.100.10
set inner 12 dst_mac ffff.ffff.ffff
start
```

```
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# ARP - Egress Leaf Elam Results (ereport)

Bridge Domain Settings:  
Unicast Routing Disable  
ARP Flooding Enabled

Outer L3 Header

Destination IP : 225.0.2.136

Destination is GIPO  
(225.0.2.128) + FTAG (0x8)

Inner L3 Header

ARP Sender IP : 192.168.100.11

ARP Target IP : 192.168.100.10

Outer L4 Header

VRF or BD VNID : 14811121( 0xE1FFF1 )

Contract Result

Contract Drop : no

FINAL FORWARDING LOOKUP

Bits set in Final Forwarding Block: : IFABRIC\_EG MC INFRA ENCAP MYTEP **BRIDGE** MISS **FLOOD**

Frame is flooded in the Bridge Domain!

Lookup Drop

Not Dropped in lookups!

**LU drop reason : no drop**

# ARP – Egress Leaf Port is VPC

- Both VPC members receive a flooded copy
- One VPC member is the Designated Forwarder (DF) for the flow
- DF is hashed per flow
- Only DF floods out VPC interfaces

## Non-DF Leaf

```
module-1(DBG-elam-insell14)# ereport | grep df | grep vpc  
  sug_lub_latch_results_vec.lub4_1.vpc_df: 0x0  
    sug_fpx_lookup_vec.lkup.dciptvec.pt.vpc_df: 0x0  
      sug_fpc_lookup_vec.fplu_vec.lkup.dciptvec.pt.vpc_df: 0x0  
        sug_fpc_lookup_vec.fplu_vec.lkup.dciptvec.pt.vpc_df: 0x0
```

## DF Leaf

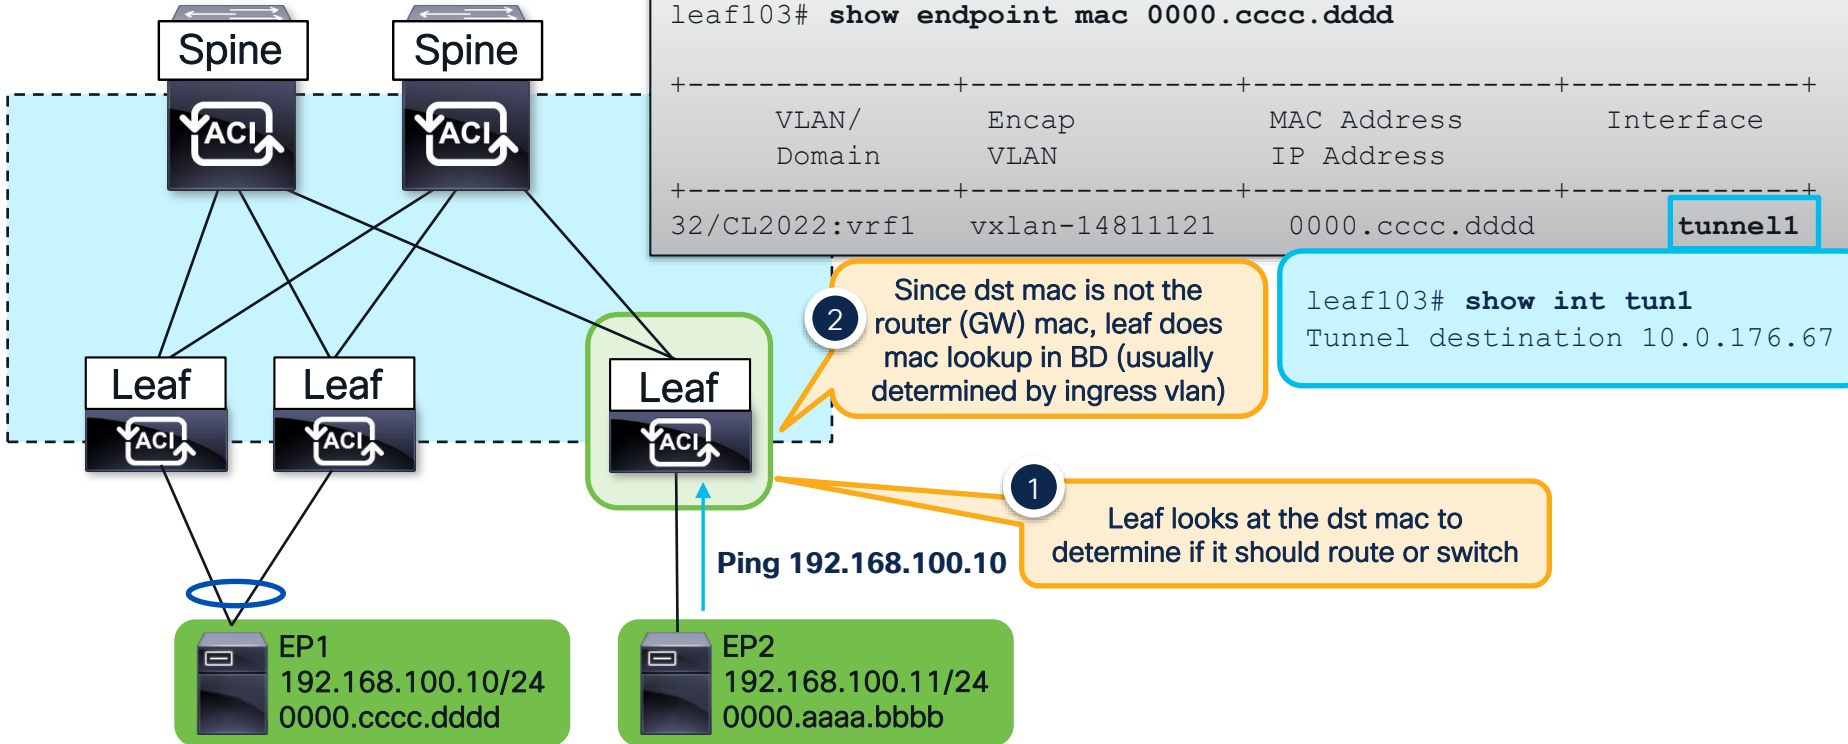
```
module-1(DBG-elam-insell14)# ereport | grep df | grep vpc  
  sug_lub_latch_results_vec.lub4_1.vpc_df: 0x1  
    sug_fpx_lookup_vec.lkup.dciptvec.pt.vpc_df: 0x1  
      sug_fpc_lookup_vec.fplu_vec.lkup.dciptvec.pt.vpc_df: 0x1  
        sug_fpc_lookup_vec.fplu_vec.lkup.dciptvec.pt.vpc_df: 0x1
```

# Debugging ACI Bridged Flows

# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood

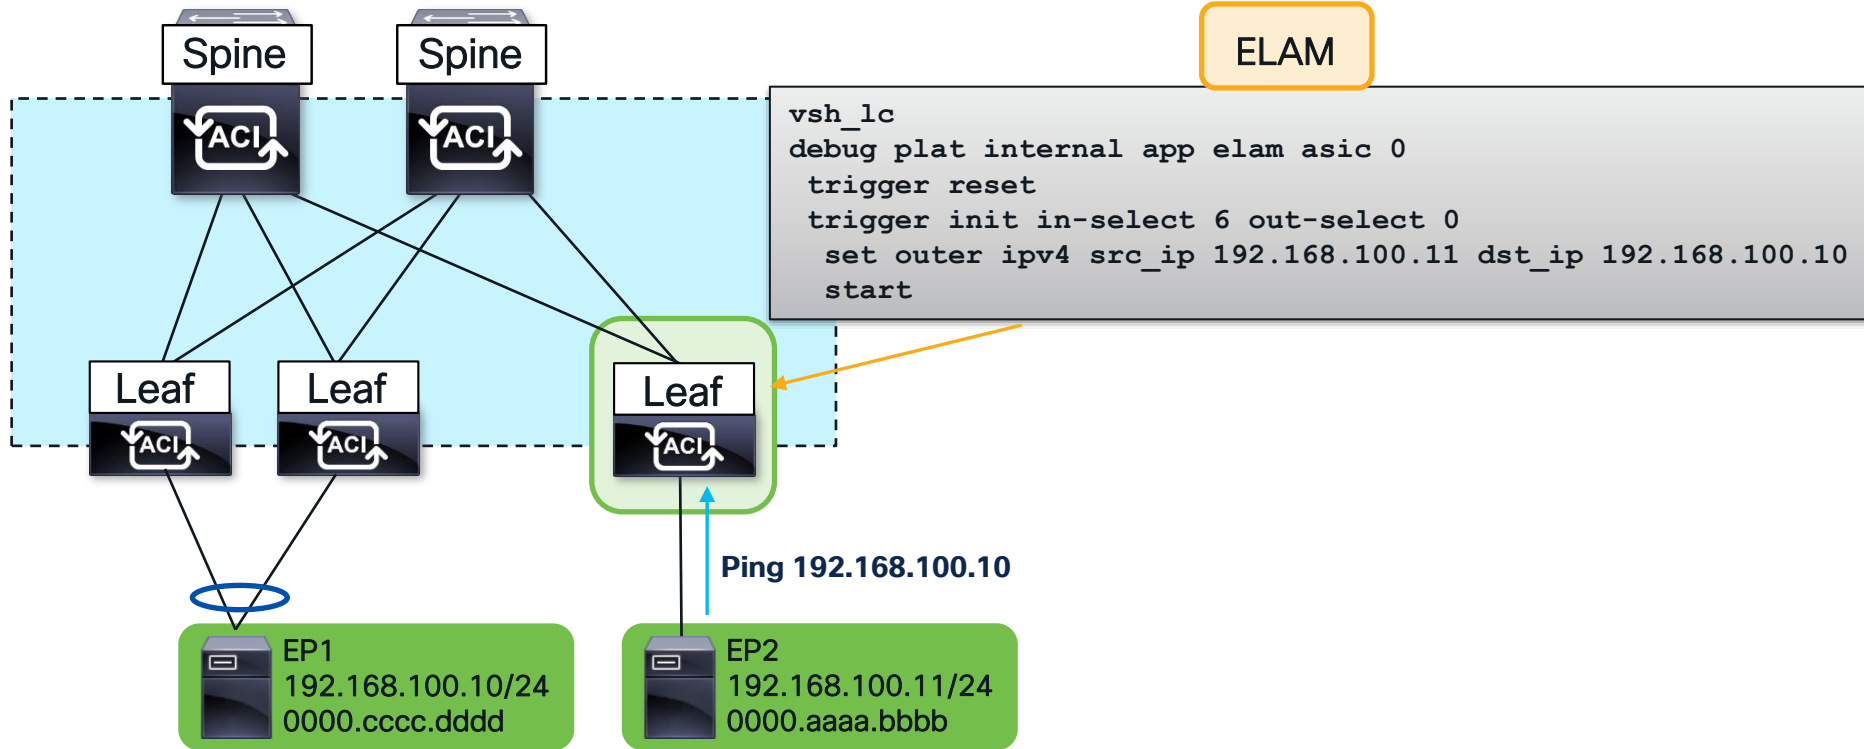
Lookup dst mac in ingress BD





# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood



# Known Unicast – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood

### Outer L2 Header

```
-----  
Destination MAC      : 0000.cccc.dddd  
Source MAC           : 0000.aaaa.bbbb  
Access Encap VLAN    : 3502 ( 0xDAE )
```

Dest mac that is looked up within BD

Make sure this is the expected vlan

### Outer L3 Header

```
-----  
IP Protocol Number  : ICMP  
Destination IP       : 192.168.100.10  
Source IP            : 192.168.100.11
```

Dest is tunnel

### Other Forwarding Information

```
-----  
Encap Index is valid : yes  
Encap Index          : 34 ( 0x22 )
```

```
show plat internal hal tunnel rtep apd  
=====
```

ifId	IP	RwEncapIdx
18010001	10.0.176.67	22

```
=====
```

Forward to this overlay TEP

### FINAL FORWARDING LOOKUP

```
-----  
Bits set in Final Forwarding Block: IFABRIC_IG UC TENANT MYTEP BRIDGE HIT
```

### Lookup Drop

```
-----  
LU drop reason      : no drop
```

Not Dropped in lookups!

Unicast + Bridge (L2 lookup) +  
Destination Known

# Known Unicast – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood

```
ereport | grep "ovector "  
ovector : 152 ( 0x98 )
```

```
show platform internal hal 12 port gpd
```

```
=====
```

IfId	Ifname	As	AP	Sl	Sp	Ss	Ovec
1a01c000	<b>Eth1/29</b>	0	59	2	18	18	<b>98</b>

```
=====
```

Traffic is forwarded out Eth1/29!

# Known Unicast – Ingress Leaf

## Contract Verification

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood

```
Contract Lookup Key
-----
IP Protocol           : ICMP( 0x1 )
L4 Src Port          : 2048( 0x800 )
L4 Dst Port          : 35914( 0x8C4A )
sclass (src pCtag)   : 49154( 0xC002 )
dclass (dst pCtag)   : 49154( 0xC002 )
src pCtag is from local table : yes
Unknown Unicast / Flood Packet : no

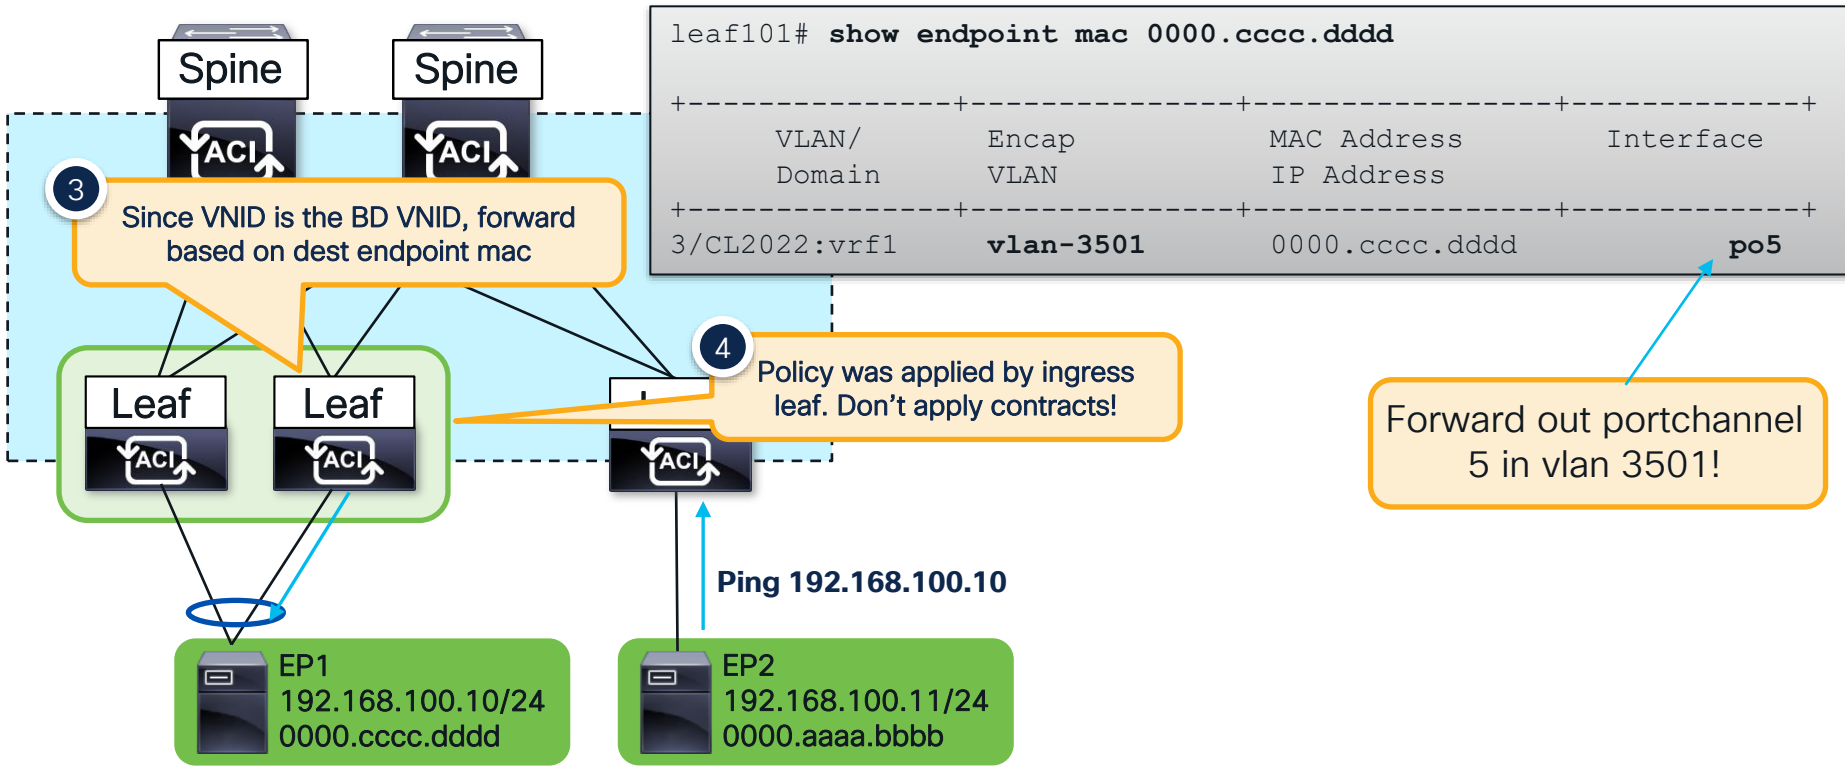
Contract Result
-----
Contract Drop        : no
Contract Applied     : yes
Contract Hit         : yes
Contract Aclqos Stats Index : 131025
( show sys int aclqos zoning-rules | grep -B 9 "Idx: 131025" )
```

Source and Dest EPG is the same. Implicitly permit!  
(unless isolation enabled)

Contract Applied and  
no Drop!

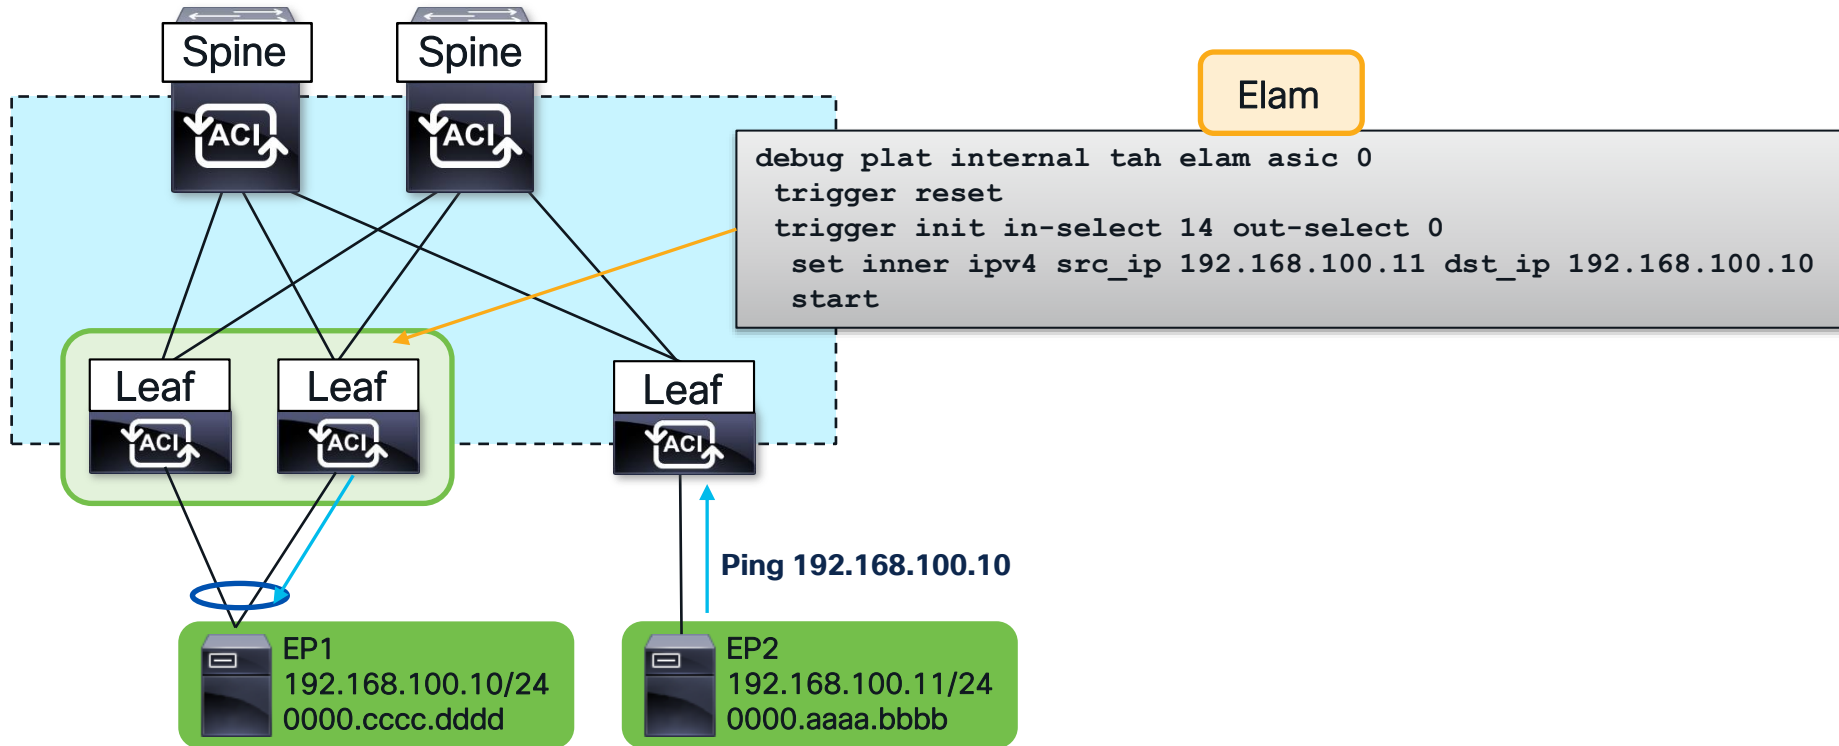
# Known Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood



# Known Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood



# Known Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Disable  
Unknown Unicast Flood

Inner L2 Header

-----  
Inner Destination MAC : 0000.cccc.dddd

Inner L3 Header

-----  
Destination IP : 192.168.100

Outer L4 Header

-----  
L4 Type : iVxLAN  
Src Policy Applied Bit : **1**  
Dst Policy Applied Bit : **1**  
VRF or BD VNID : **14811121 ( 0xE1FFF1 )**

Contracts have already been applied. No need to check.

Mac lookup done in bridge domain with this VNID

Sideband Information

-----  
ovector : **146 ( 0x92 )**

```
show platform internal hal 12 port gpd
```

```
=====
IfId      Ifname      As AP  S1 Sp Ss  Ovec
=====
1a021000  Eth1/34     0  32  1  9  12  92
=====
```

Forward out Eth1/34!

FINAL FORWARDING LOOKUP

-----  
Bits set in Final Forwarding Block: IFABRIC\_EG **UC** INFRA ENCAP MYTEP **BRIDGE HIT**

Lookup Drop

-----  
LU drop reason : **no drop**

Unicast + Bridge (L2 lookup) + Destination Known

# Debugging ACI Routed Flows





# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled

Lookup dst IP in ingress VRF

```
leaf103# show endpoint ip 192.168.100.10
+-----+-----+-----+
| VLAN/ | MAC Address | Interface |
| Domain | IP Address  |           |
+-----+-----+-----+
| CL2022:vrf1 | 192.168.100.10 | tunnel1 |
```

```
leaf103# show int tun1
Tunnel destination 10.0.176.67
```

2 Since dst mac is the router (GW) mac, leaf does IP lookup in VRF of source IP

1 Leaf looks at the dst mac to determine if it should route or switch

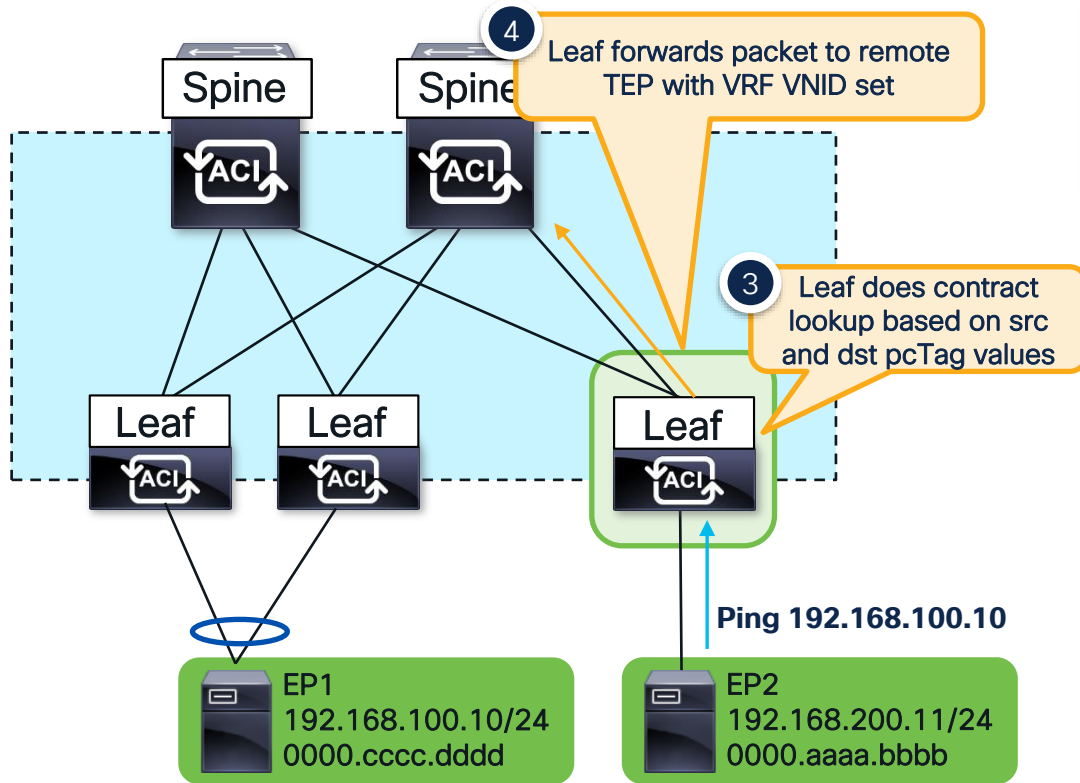
Ping 192.168.100.10

EP1  
192.168.100.10/24  
0000.cccc.dddd

EP2  
192.168.200.11/24  
0000.aaaa.bbbb

# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



Get Sclass

```
103# show sys internal epm endpoint ip
192.168.200.11
!omitted
BD vnid : 16613259 ::: VRF vnid : 2523136
sclass : 32771
```

Get Dclass

```
103# show sys internal epm endpoint ip
192.168.100.10
!omitted
BD vnid : 0 ::: VRF vnid : 2523136
sclass : 49154
```


Check Contract

```
103# show zoning-rule src-epg 32771
dst-epg 49154 scope 2523136
```

RuleID	Name	Action
4209	CL2022:allow-all	permit

# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



ELAM

```
vsh_lc
debug plat internal app elam asic 0
trigger reset
trigger init in-select 6 out-select 0
set outer ipv4 src_ip 192.168.200.11
set outer ipv4 dst_ip 192.168.100.10
start
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# Known Unicast – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Enabled



### Outer L2 Header

```
-----  
Destination MAC      : 0022.BDF8.19FF  
Access Encap VLAN   : 3769 ( 0xEB9 )
```

ACI Router Mac. Route this packet!

Make sure this is the expected vlan

### Outer L3 Header

```
-----  
Destination IP       : 192.168.100.10  
Source IP            : 192.168.200.11
```

Dest is tunnel

### Other Forwarding Information

```
-----  
Encap Index is valid : yes  
Encap Index          : 34 ( 0x22 )
```

```
show plat internal hal tunnel rtep apd
```

```
=====
```

ifId	IP	RwEncapIdx
18010001	10.0.176.67	22

```
=====
```

Forward to this overlay TEP

### FINAL FORWARDING LOOKUP

```
-----  
Bits set in Final Forwarding Block: IFABRIC_IG UC TENANT MYTEP ROUTE HIT
```

### Lookup Drop

```
-----  
LU drop reason       : no drop
```

Not Dropped in lookups!

Unicast + Route (L3 lookup) +  
L3 Route Found

# Known Unicast – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Enabled



```
ereport | grep "ovector "  
ovector : 152 ( 0x98 )
```

```
show platform internal hal 12 port gpd
```

```
=====
```

IfId	Ifname	As	AP	Sl	Sp	Ss	Ovec
1a01c000	Eth1/29	0	59	2	18	18	98

```
=====
```

Traffic is forwarded out Eth1/29!

# Known Unicast – Ingress Leaf

## Contract Verification

Bridge Domain Settings:  
Unicast Routing Enabled



```
Contract Lookup Key
-----
IP Protocol           : ICMP( 0x1 )
L4 Src Port          : 2048( 0x800 )
L4 Dst Port          : 31219( 0x79F3 )
sclass (src pCtag)   : 32771( 0x8003 )
dclass (dst pCtag)   : 49154( 0xC002 )
src pCtag is from local table : yes
Unknown Unicast / Flood Packet : no
```

Source and Dest EPG  
used for contract lookup

```
Contract Result
-----
Contract Drop        : no
Contract Applied     : yes
Contract Hit         : yes
Contract Aclqos Stats Index : 131025
```

Contract Applied and  
no Drop!

But how do I know which  
contract this is actually hitting?

# Known Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



## Contract Verification

```
Contract Result
-----
Contract Drop      : no
Contract Applied   : yes
Contract Hit       : yes
Contract Aclqos Stats Index : 131025
```

Hardware Index of matching contract

Run this from vsh\_lc

```
show sys int aclqos zoning-rules | grep -B 9 "Idx: 130974"
=====
Rule ID: 4163 Scope 8 Src EPG: 32771 Dst EPG: 49154 Filter 532
Curr TCAM resource:
=====
=== SDK Info ===
Result/Stats Idx: 130974
```

Zoning-rule ID

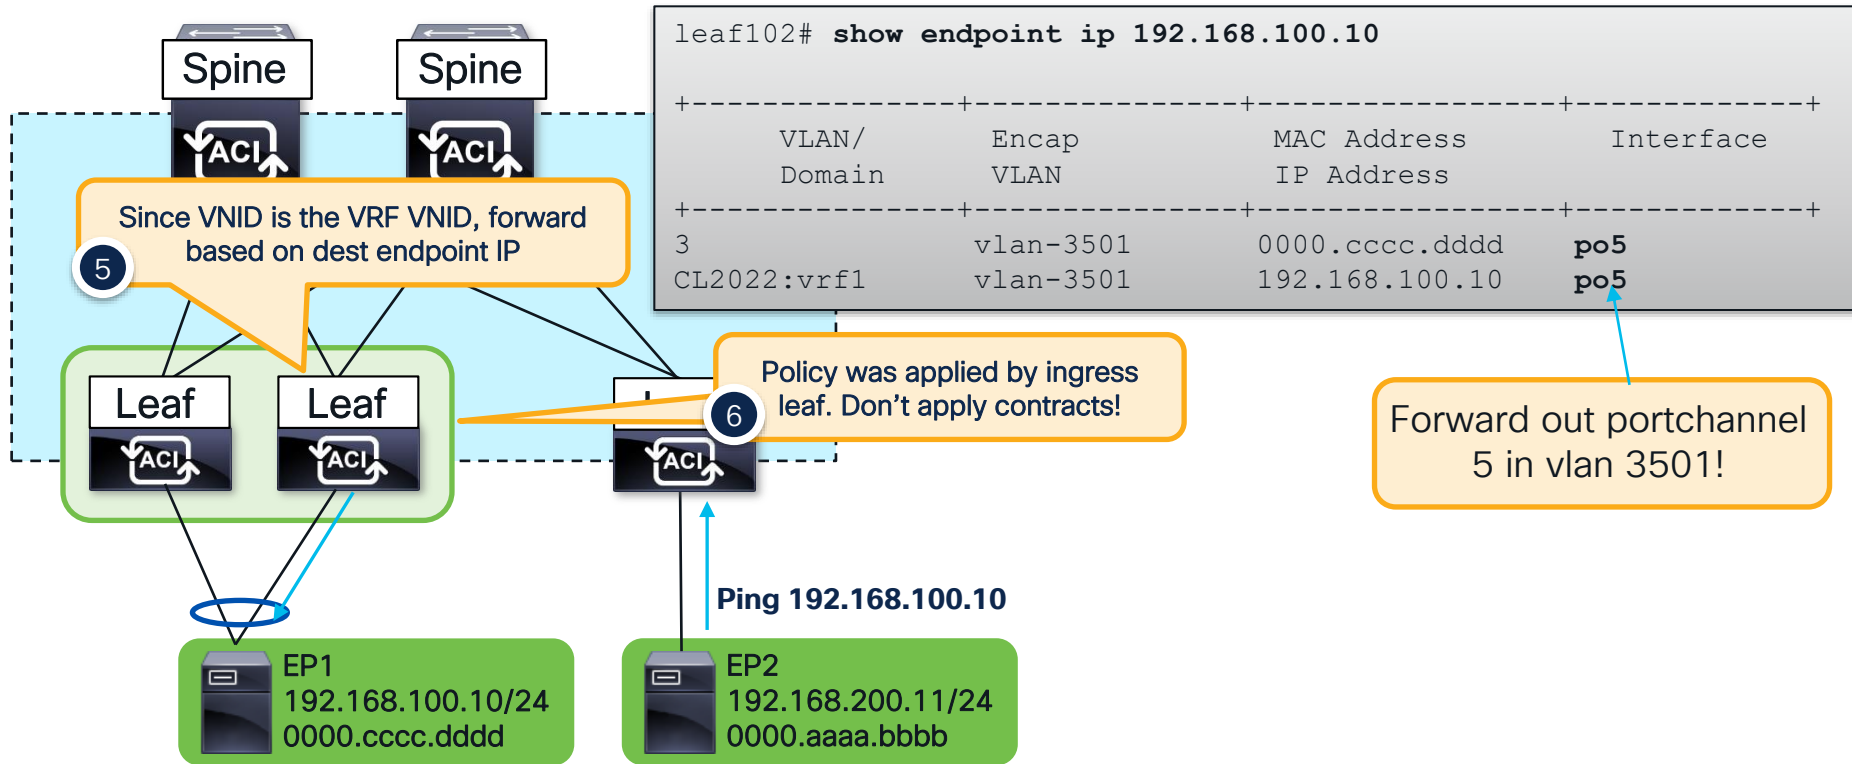
Run this from normal shell

```
show zoning-rule rule-id 4163
+-----+-----+-----+-----+-----+-----+-----+
| Rule ID | SrcEPG | DstEPG | FilterID | Scope | Name | Action |
+-----+-----+-----+-----+-----+-----+-----+
| 4163 | 32771 | 49154 | 532 | 2523136 | CL2022:allow-all | permit |
+-----+-----+-----+-----+-----+-----+-----+
```

Traffic hit this contract!

# Known Unicast – Egress Leaf

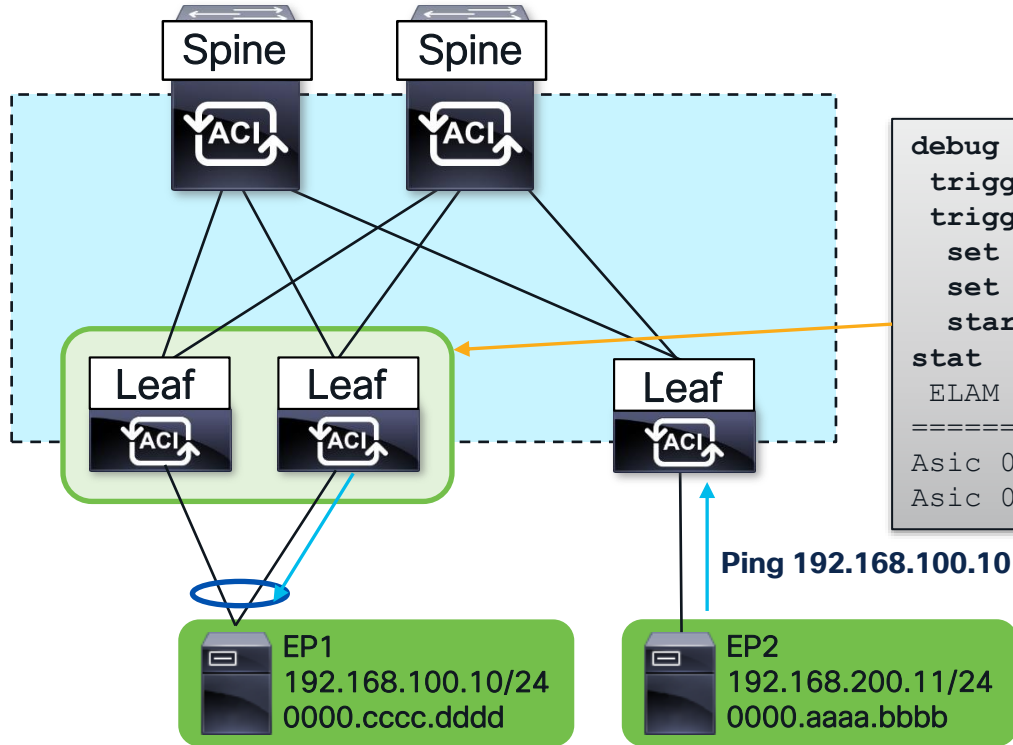
Bridge Domain Settings:  
Unicast Routing Enabled





# Known Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



Elam

```
debug plat internal tah elam asic 0
trigger reset
trigger init in-select 14 out-select 0
set inner ipv4 src_ip 192.168.200.11
set inner ipv4 dst_ip 192.168.100.10
start
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# Known Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



Inner L2 Header

Inner Destination MAC : 000C.0C0C.0C0C

Inner L3 Header

Destination IP : 192.168.100

Outer L4 Header

L4 Type : iVxLAN  
Src Policy Applied Bit : **1**  
Dst Policy Applied Bit : **1**  
VRF or BD VNID : **2523136 ( 0x268000 )**

Contracts have already been applied. No need to check.

IP lookup done in VRF with this VNID

```
show platform internal hal 12 port gpd
```

IfId	Ifname	As	AP	S1	Sp	Ss	Ovec
1a021000	Eth1/34	0	32	1	9	12	92

Forward out Eth1/34!

Sideband Information

ovector : 146 ( 0x92 )

FINAL FORWARDING LOOKUP

Bits set in Final Forwarding Block: IFABRIC\_EG UC INFRA ENCAP MYTEP ROUTE HIT

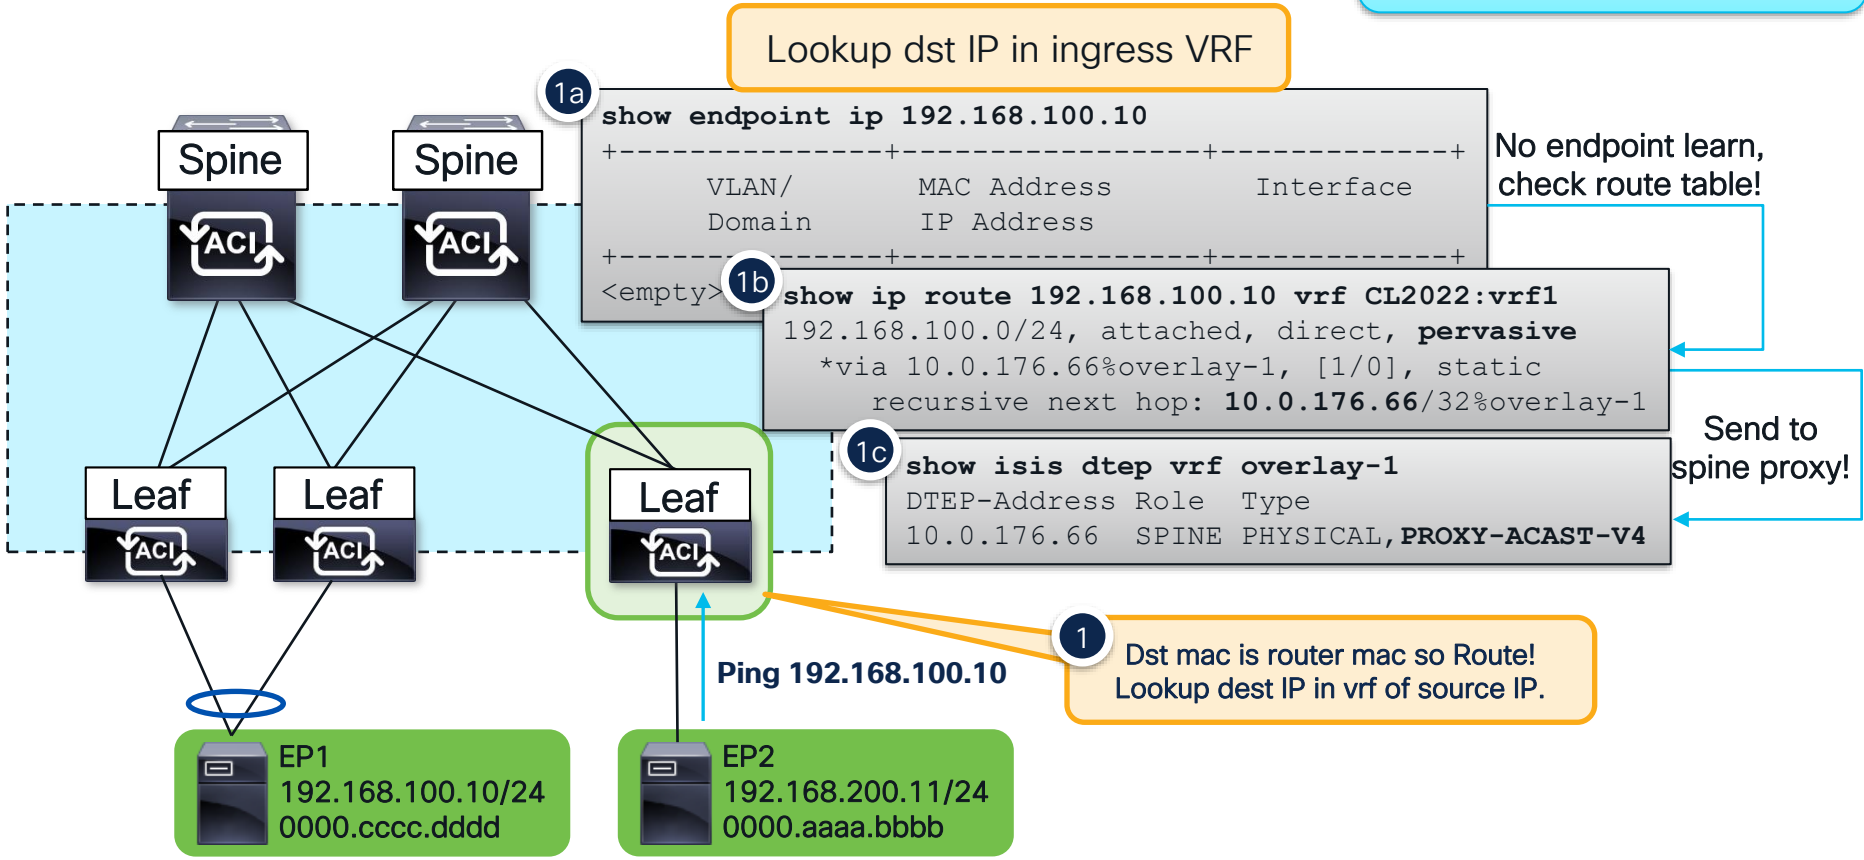
Lookup Drop

LU drop reason : no drop

Unicast + Route (L3 lookup) + L3 Route Found

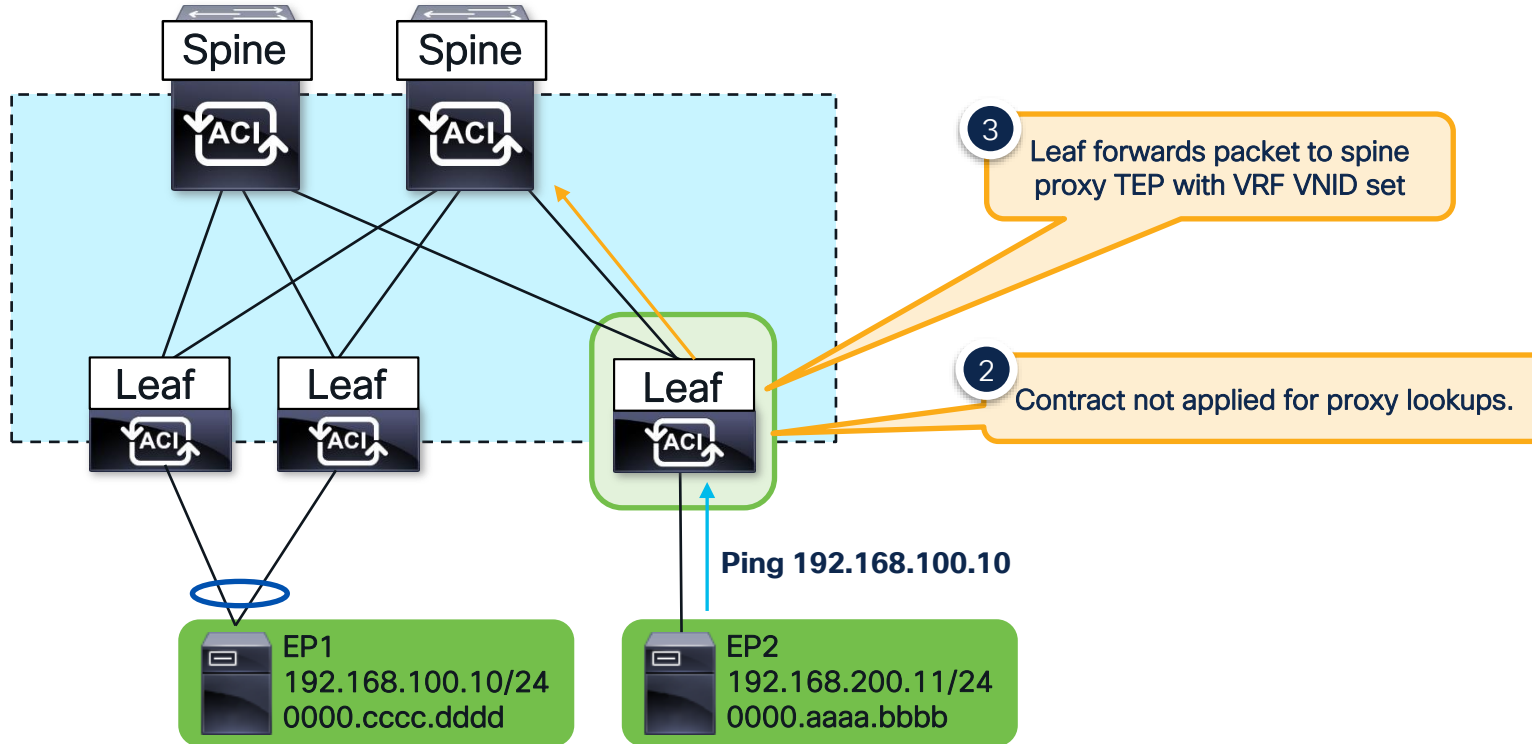
# Proxied Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



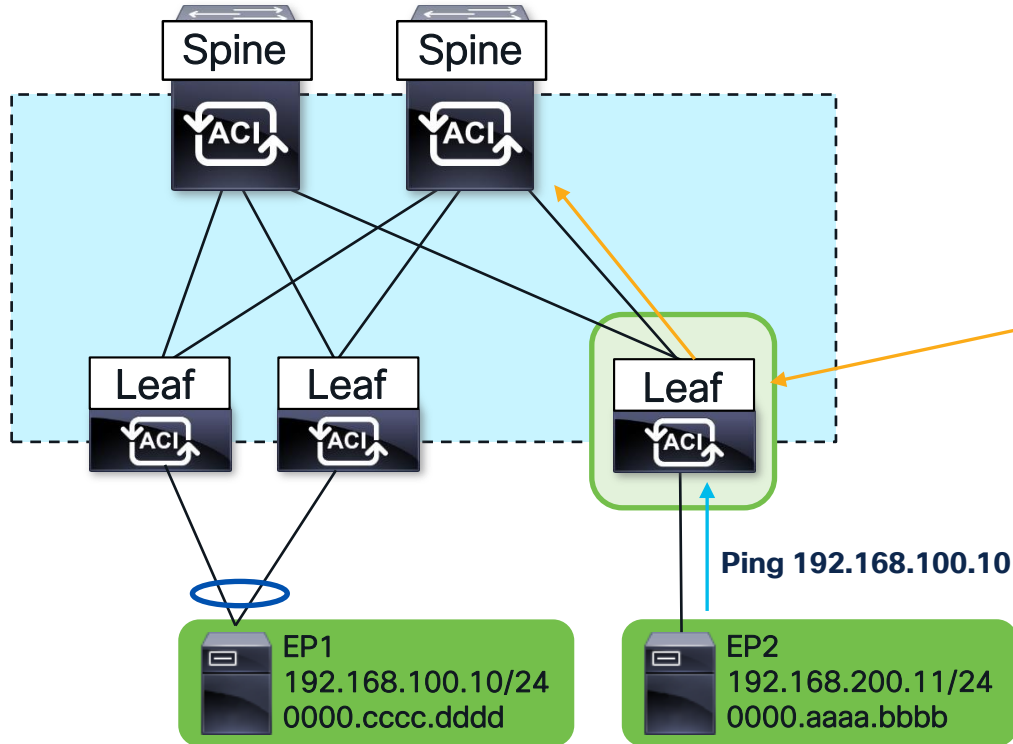
# Proxied Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



# Proxied Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



ELAM

```
vsh_lc
debug plat internal app elam asic 0
trigger reset
trigger init in-select 6 out-select 0
set outer ipv4 src_ip 192.168.200.11
set outer ipv4 dst_ip 192.168.100.10
start
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# Proxied Unicast – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled

## Forwarding Verifications

```
Outer L2 Header
-----
Destination MAC      : 0022.BDF8.19FF
Access Encap VLAN   : 3769 ( 0xEB9 )

Outer L3 Header
-----
Destination IP      : 192.168.100.10
Source IP           : 192.168.200.11

Other Forwarding Information
-----
Encap Index is valid : yes
Encap Index          : 1 ( 0x1 )

FINAL FORWARDING LOOKUP
-----
Bits set in Final Forwarding Block: IFABRIC_IG UC_TENANT MYTEP ROUTE_HIT

Lookup Drop
-----
LU drop reason      : no drop
```

ACI Router Mac. Route this packet!

Make sure this is the expected vlan

Dest is tunnel

```
show plat internal hal tunnel rtep apd
=====
ifId      IP          RwEncapIdx
=====
18010007 10.0.176.66 1
```

Forward to this overlay TEP

Not Dropped in lookups!

Unicast + Route (L3 lookup) +  
L3 Route Found

# Proxied Unicast – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Enabled

```
ereport | grep "ovector "  
ovector : 152 ( 0x98 )
```

```
show platform internal hal 12 port gpd
```

```
=====
```

IfId	Ifname	As	AP	Sl	Sp	Ss	Ovec
1a01c000	<b>Eth1/29</b>	0	59	2	18	18	<b>98</b>

```
=====
```

Traffic is forwarded out Eth1/29!

# Proxied Unicast – Ingress Leaf

## Contract Verification

Bridge Domain Settings:  
Unicast Routing Enabled

```
Contract Lookup Key
-----
IP Protocol           : ICMP( 0x1 )
L4 Src Port          : 2048( 0x800 )
L4 Dst Port          : 31219( 0x79F3 )
sclass (src pcTag)   : 32771( 0x8003 )
dclass (dst pcTag)   : 1( 0x1 )
src pcTag is from local table : yes
Unknown Unicast / Flood Packet : no
```

Dest EPG is 1 for fabric owned subnets

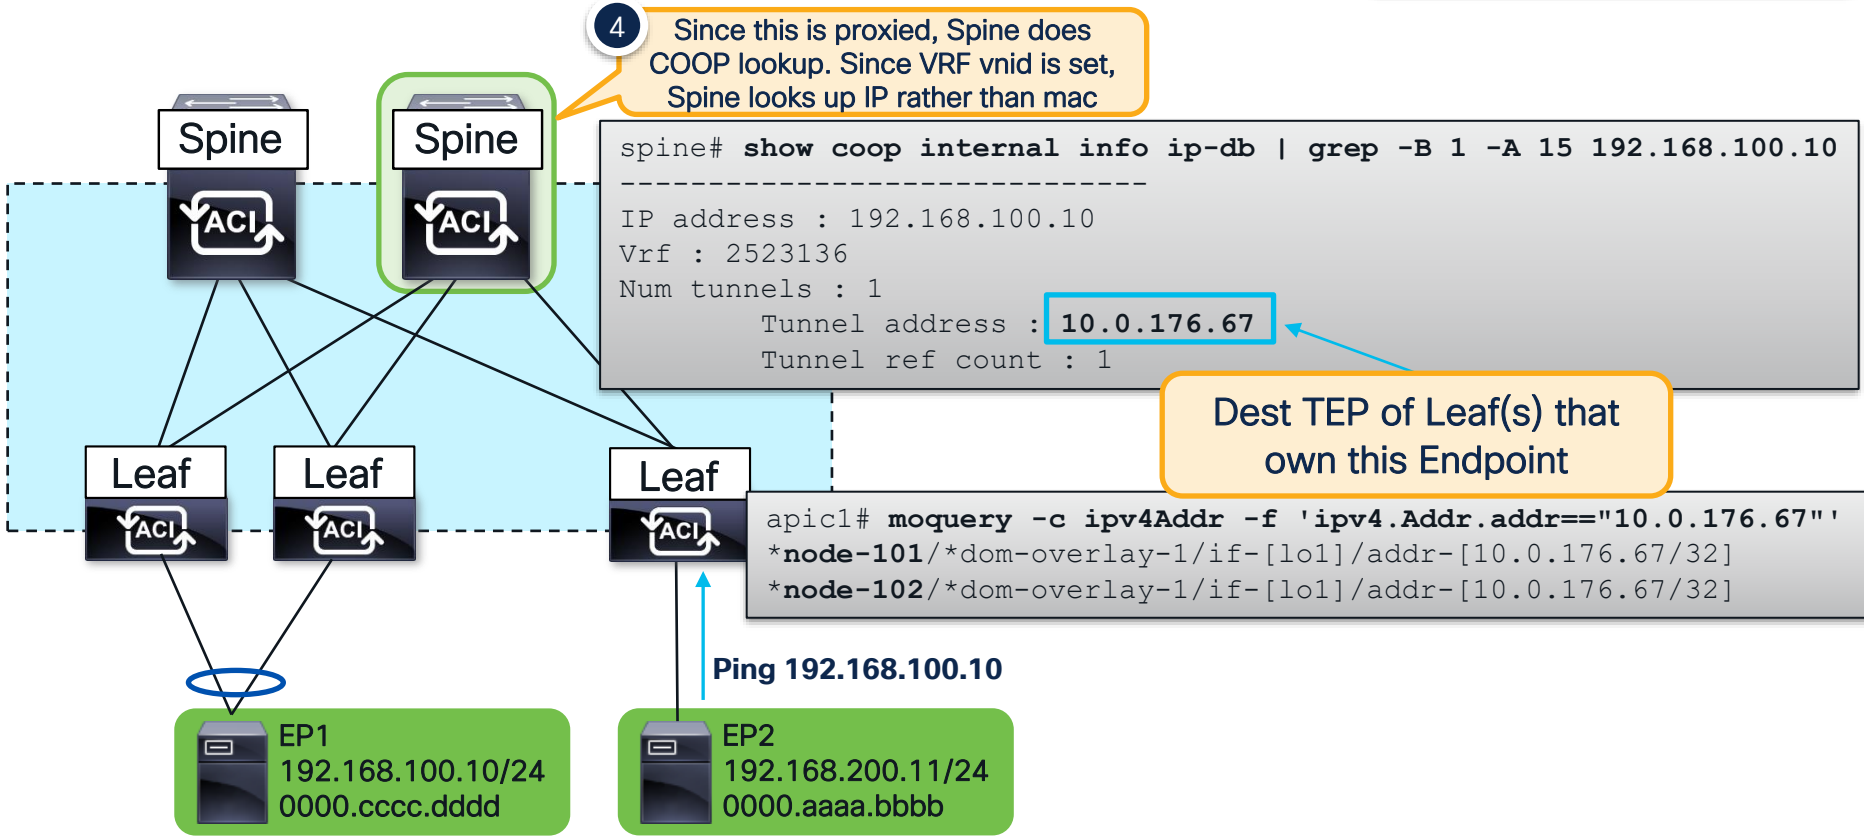
```
Contract Result
-----
Contract Drop        : no
Contract Applied     : no
Contract Hit         : yes
Contract Aclqos Stats Index : 131025
```

Contract not applied since this is proxied!



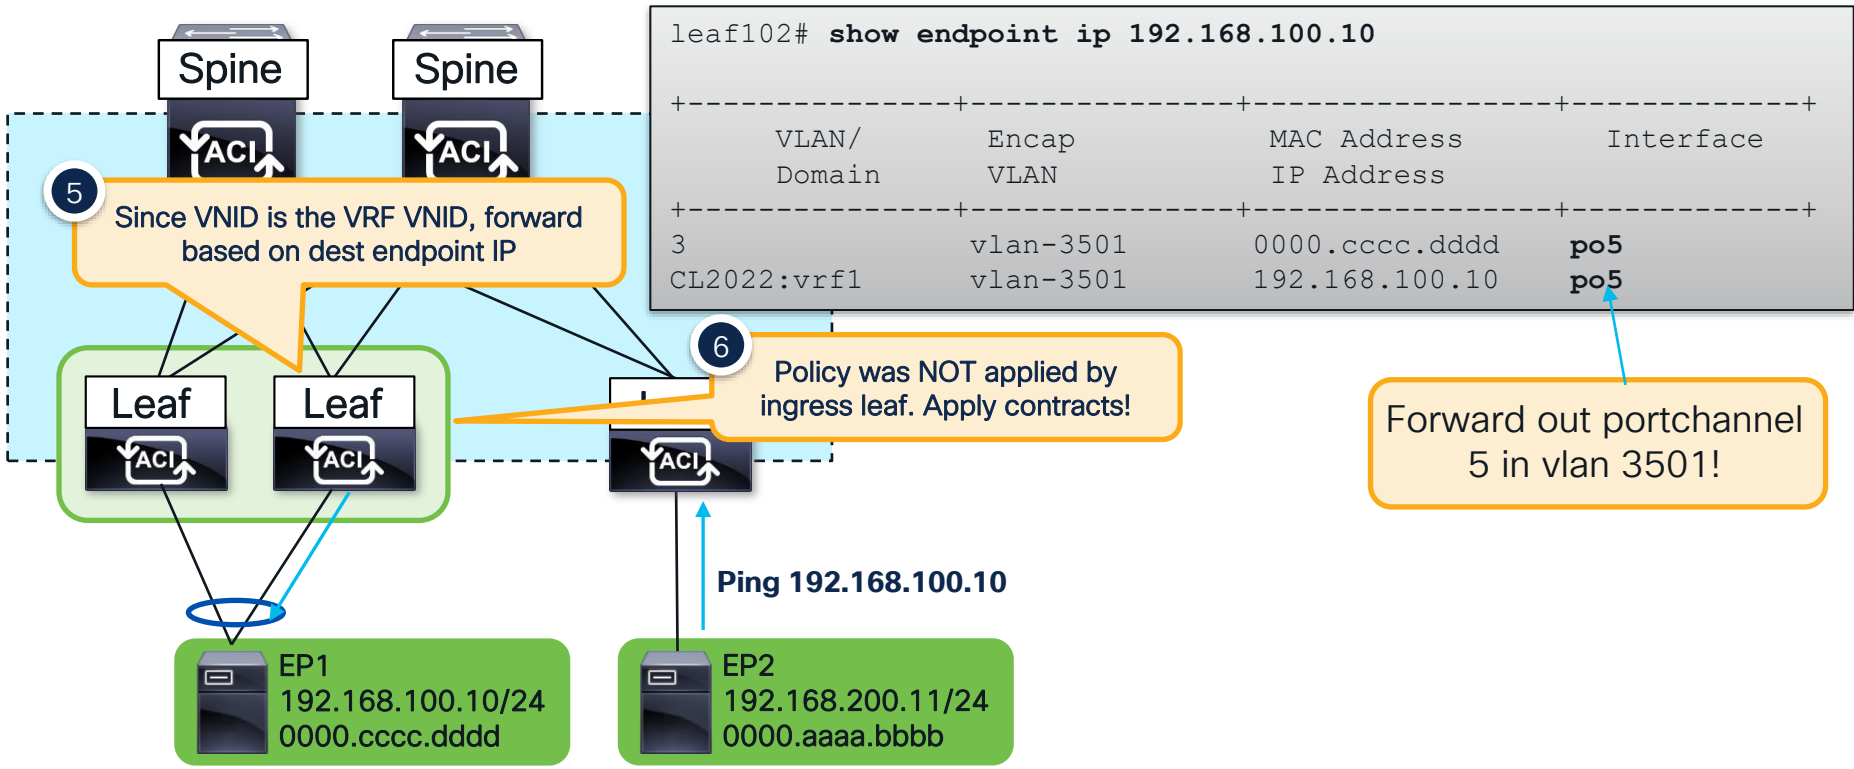
# Proxied Unicast – Spine

Bridge Domain Settings:  
Unicast Routing Enabled



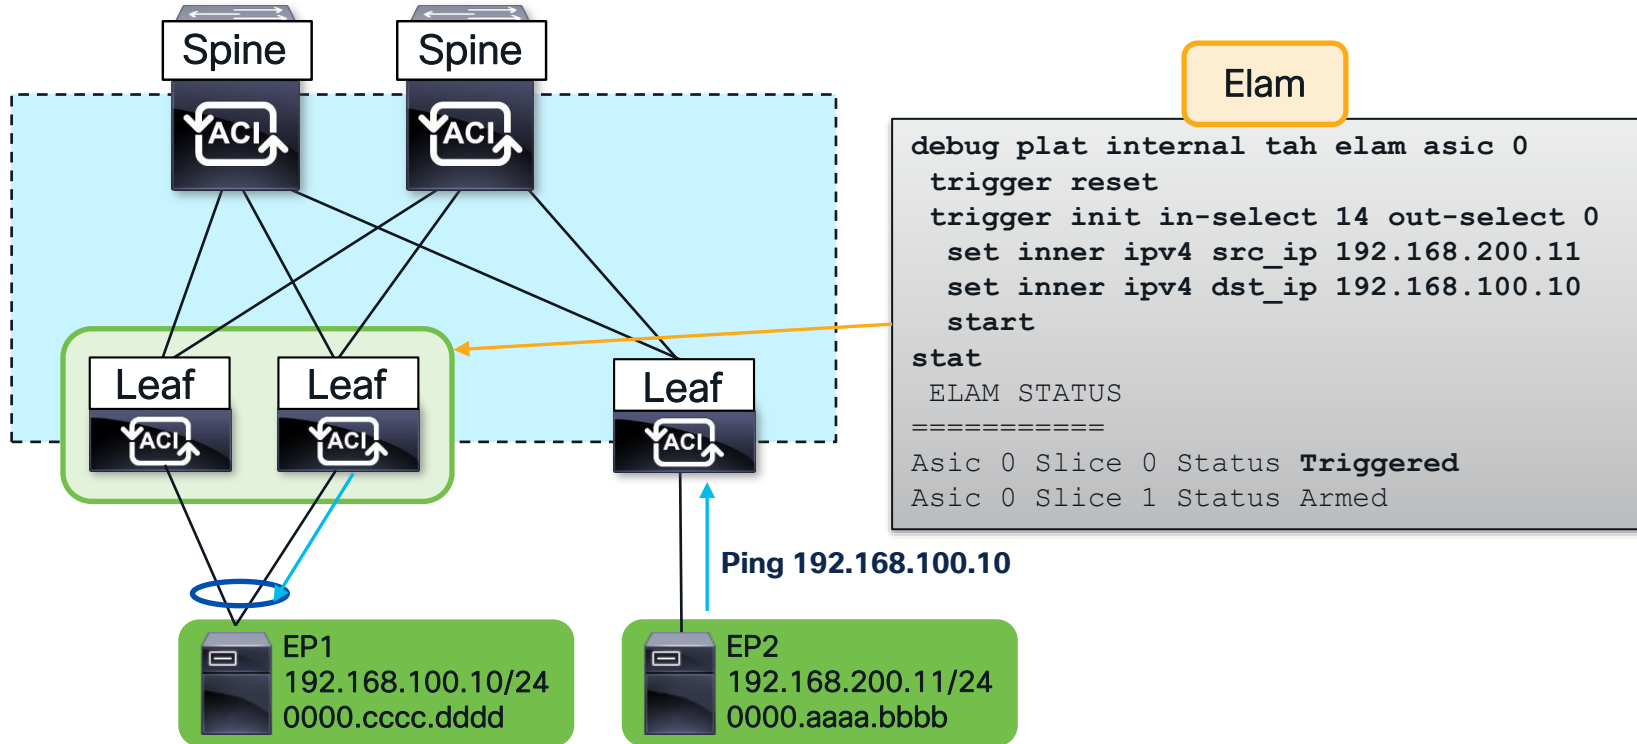
# Proxied Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



# Proxied Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



# Proxied Unicast – Egress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Enabled

```
Inner L3 Header
-----
Destination IP      : 192.168.100
Outer L4 Header
-----
L4 Type             : IPv4
Src Policy Applied Bit : 0
Dst Policy Applied Bit : 0
VRF or BD VNID      : 2523136 ( 0x268000 )
Sideband Information
-----
ovector             : 146 ( 0x92 )
FINAL FORWARDING LOOKUP
-----
Bits set in Final Forwarding Block: IFABRIC_EG UC INFRA ENCAP MYTEP ROUTE HIT
Lookup Drop
-----
LU drop reason      : no drop
```

Contracts have not been applied yet!

IP lookup done in VRF with this VNID

```
show platform internal hal 12 port gpd
=====
IfId      Ifname      As AP   S1 Sp Ss  Ovec
=====
1a021000  Eth1/34      0  32  1  9  12  92
```

Forward out Eth1/34!

Not Dropped in lookups!

Unicast + Route (L3 lookup) +  
L3 Route Found

# Proxied Unicast – Egress Leaf

## Contract Verification

Bridge Domain Settings:  
Unicast Routing Enabled

```
Contract Lookup Key
-----
IP Protocol           : ICMP( 0x1 )
L4 Src Port          : 2048( 0x800 )
L4 Dst Port          : 33226( 0x81CA )
sclass (src pCtag)   : 32771 ( 0x8003 )
dclass (dst pCtag)   : 49154 ( 0xC002 )
src pCtag is from local table : no
Unknown Unicast / Flood Packet : no

Contract Result
-----
Contract Drop        : no
Contract Applied     : yes
Contract Hit         : yes
Contract Aclqos Stats Index : 131025
```

Source and Dest EPG used  
for contract lookup.

Contract Applied and  
no Drop!

But how do I know which  
contract this is actually hitting?

# Proxied Unicast – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled

## Contract Verification

```
Contract Result
-----
Contract Drop           : no
Contract Applied        : yes
Contract Hit            : yes
Contract Aclqos Stats Index : 81836
```

Hardware Index of matching contract

Run this from vsh\_lc

```
show sys int aclqos zoning-rules | grep -B 9 "Idx: 81836"
=====
Rule ID: 4234 Scope 16 Src EPG: 32771 Dst EPG: 49154 Filter
532
=====
=== SDK Info ===
Result/Stats Idx: 81836
```

Zoning-rule ID

Run this from normal shell

```
show zoning-rule rule-id 4234
+-----+-----+-----+-----+-----+-----+-----+
| Rule ID | SrcEPG | DstEPG | FilterID | Scope | Name | Action |
+-----+-----+-----+-----+-----+-----+-----+
| 4163 | 32771 | 49154 | 532 | 2523136 | CL2022:allow-all | permit |
+-----+-----+-----+-----+-----+-----+-----+
```

Traffic hit this contract!

# L3Out Destination – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled

Lookup dst IP in ingress VRF

2  
Since dst mac is the router (GW) mac,  
leaf does IP lookup in VRF of source IP

2a

```
show endpoint ip 10.99.99.100
+-----+-----+-----+
| VLAN/  | MAC Address | Interface |
| Domain | IP Address  |           |
+-----+-----+-----+
| <empty> |             |           |
+-----+-----+-----+
```

No endpoint learn,  
check route table!

2b

```
show ip route 10.99.99.100 vrf CL2022:vrf1
10.99.99.0/24, ubest/mbest: 1/0
 *via 10.0.64.70%overlay-1, [200/20], bgp-65100
  recursive next hop: 10.0.64.70/32%overlay-1
```

Send to BL  
PTEP!

2c

```
acidiag fmvread | grep 10.0.64.70
Name      IP Address      Role
-----
leaf103   10.0.64.70/32   leaf
```

1  
Leaf looks at the dst mac to  
determine if it should route or switch

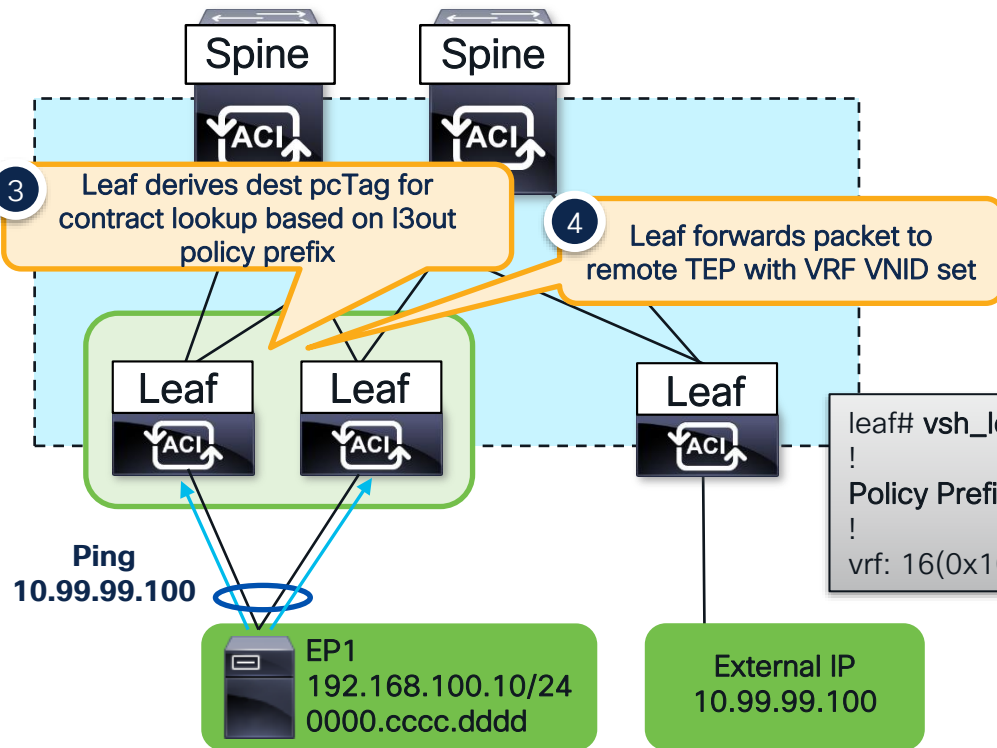
Ping  
10.99.99.100

EP1  
192.168.100.10/24  
0000.cccc.dddd

External IP  
10.99.99.100

# L3Out Destination – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



## External EPGs

### External EPGs

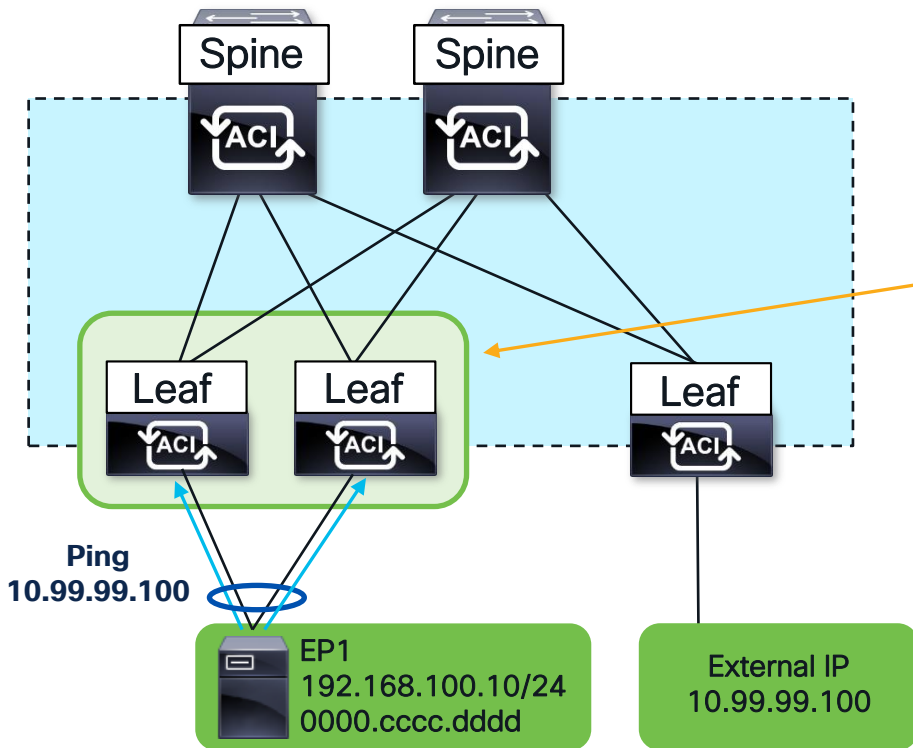
Name	Description	pcTag
all	10.99.99.0/24 Network	32772

```
leaf# vsh_lc -c "show forwarding route 10.99.99.100 platf vrf CL2022:vrf1"
!
Policy Prefix 10.99.99.0/24
!
vrf: 16(0x10), routed_if: 0x0 epc_class: 32772(0x8004)
```



# L3Out Destination – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



ELAM

```
vsh_lc
debug plat internal tah elam asic 0
trigger reset
trigger init in-select 6 out-select 0
set outer ipv4 src_ip 192.168.100.10
set outer ipv4 dst_ip 10.99.99.100
start
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# L3Out Destination – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



## Forwarding Verifications

Outer L2 Header

```
-----  
Destination MAC      : 0022.BDF8.19FF  
Access Encap VLAN   : 3501 ( 0xDAD )
```

ACI Router Mac. Route this packet!

Make sure this is the expected vlan

Outer L3 Header

```
-----  
Destination IP       : 10.99.99.100  
Source IP            : 192.168.100.10
```

Dest is tunnel

Other Forwarding Information

```
-----  
Encap Index is valid : yes  
Encap Index          : 37 ( 0x25 )
```

```
show plat internal hal tunnel rtep apd
```

```
=====
```

ifId	IP	RwEncapIdx
18010004	10.0.64.70	25

```
=====
```

Forward to this overlay TEP

FINAL FORWARDING LOOKUP

```
-----  
Bits set in Final Forwarding Block: IFABRIC_IG UC_TENANT MYTEP ROUTE HIT
```

Lookup Drop

```
-----  
LU drop reason       : no drop
```

Not Dropped in lookups!

Unicast + Route (L3 lookup) +  
L3 Route Found

# L3Out Destination – Ingress Leaf

## Forwarding Verifications

Bridge Domain Settings:  
Unicast Routing Enabled



```
ereport | grep "ovector "  
ovector : 48 ( 0x30 )
```

```
show platform internal hal 12 port gpd
```

```
=====
```

IfId	Ifname	As	AP	Sl	Sp	Ss	Ovec
1a035000	<b>Eth1/54</b>	0	19	0	18	30	<b>30</b>

```
=====
```

Traffic is forwarded out Eth1/54!

# L3Out Destination – Ingress Leaf

## Contract Verification

Bridge Domain Settings:  
Unicast Routing Enabled



```
Contract Lookup Key
-----
IP Protocol           : ICMP( 0x1 )
L4 Src Port          : 2048( 0x800 )
L4 Dst Port          : 12063( 0x2F1F )
sclass (src pcTag)   : 49154( 0xC002 )
dclass (dst pcTag)   : 32772( 0x8004 )
src pcTag is from local table : yes
Unknown Unicast / Flood Packet : no
```

Source and Dest EPG  
used for contract lookup

```
Contract Result
-----
Contract Drop        : no
Contract Applied     : yes
Contract Hit         : yes
Contract Aclqos Stats Index : 81765
```

Contract Applied and  
no Drop!

But how do I know which  
contract this is actually hitting?

# L3Out Destination – Ingress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



## Contract Verification

```
Contract Result
-----
Contract Drop           : no
Contract Applied        : yes
Contract Hit            : yes
Contract Aclqos Stats Index : 81765
```

Hardware Index of matching contract

Run this from vsh\_lc

```
show sys int aclqos zoning-rules | grep -B 9 "Idx: 81765"
=====
Rule ID: 4248 Scope 16 Src EPG: 0 Dst EPG: 32772 Filter 532
Curr TCAM resource:
=====
=== SDK Info ===
Result/Stats Idx: 81765
```

Zoning-rule ID

Run this from normal shell

```
show zoning-rule rule-id 4248
+-----+-----+-----+-----+-----+-----+-----+
| Rule ID | SrcEPG | DstEPG | FilterID | Scope | Name | Action |
+-----+-----+-----+-----+-----+-----+-----+
| 4248 | 0 | 32772 | 532 | 2523136 | CL2022:l3out-allow-all | permit |
+-----+-----+-----+-----+-----+-----+-----+
```

Traffic hit this contract!

# L3Out Destination – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled

Lookup dst IP in received VRF

5 Since received VNID is the VRF VNID, forward based on dest endpoint IP

```
5a show endpoint ip 10.99.99.100
-----+-----+-----+
VLAN/      MAC Address      Interface
Domain     IP Address
-----+-----+-----+
<empty>
```

No endpoint learn, check route table!

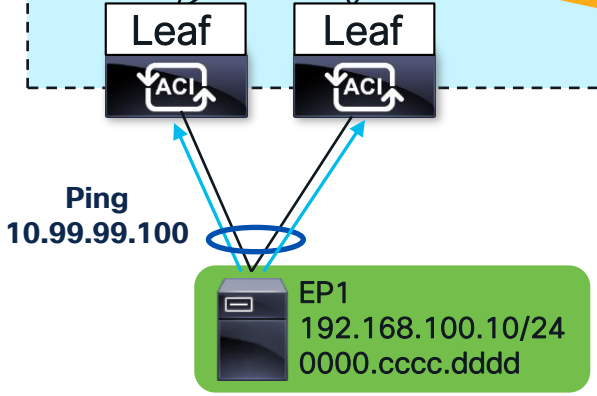
```
5b show ip route 10.99.99.100 vrf CL2022:vrf1
10.99.99.0/24, ubest/mbest: 1/0
 *via 10.55.0.100, vlan25, [110/20], ospf, type-2
```

```
5c show ip arp 10.55.0.100 vrf CL2022:vrf1
Address      MAC Address      Interface
10.55.0.100  0005.73ff.593c  vlan25
```

```
5d show mac address addr 0005.73ff.593c vl 25
VLAN      MAC Address      Ports
-----+-----+-----+
* 25      0005.73ff.593c  eth1/27/4
```

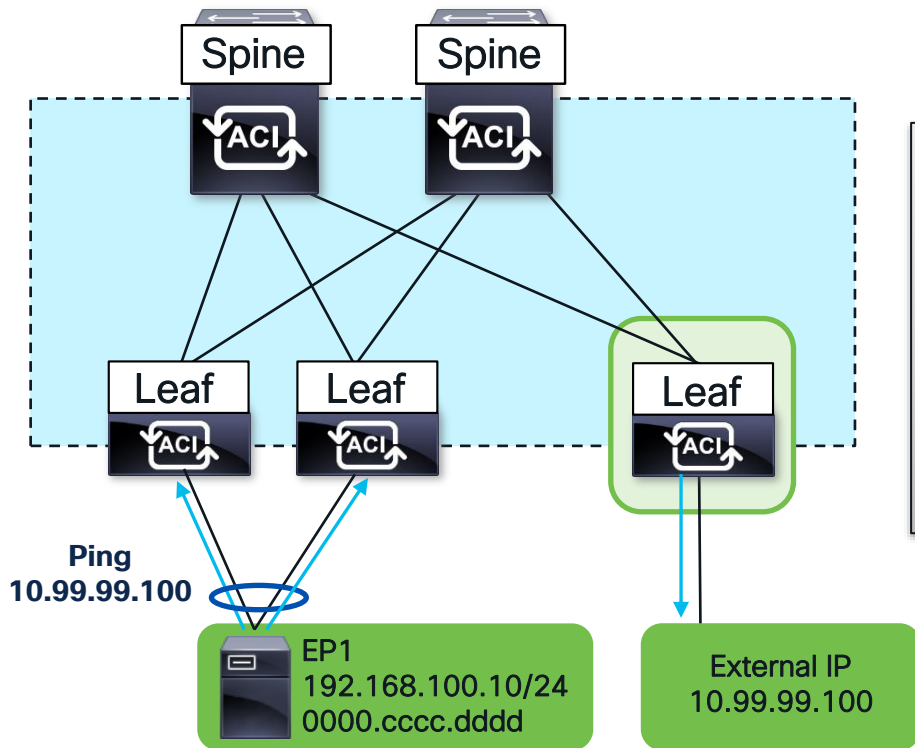
Forward based on ARP and MAC Adjacencies

6 Policy was applied by ingress leaf. No need to apply contracts



# L3Out Destination – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



Elam

```
debug plat internal app elam asic 0
trigger reset
trigger init in-select 14 out-select 0
set inner ipv4 src_ip 192.168.100.10
set inner ipv4 dst_ip 10.99.99.100
start
stat
ELAM STATUS
=====
Asic 0 Slice 0 Status Triggered
Asic 0 Slice 1 Status Armed
```

# L3Out Destination – Egress Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



Inner L2 Header

Inner Destination MAC : 000C.0C0C.0C0C

Inner L3 Header

Destination IP : 10.99.99.10

Outer L4 Header

L4 Type : iVxLAN

Src Policy Applied Bit : 1

Dst Policy Applied Bit : 1

VRF or BD VNID : 2523136 ( 0x268000 )

Sideband Information

ovector : 147 ( 0x93 )

FINAL FORWARDING LOOKUP

Bits set in Final Forwarding Block: IFABRIC\_EG UC INFRA ENCAP MYTEP ROUTE HIT

Lookup Drop

LU drop reason : no drop

Contracts have already been applied. No need to check.

IP lookup done in VRF with this VNID

```
show platform internal hal 12 port gpd
```

IfId	Ifname	As	AP	S1	Sp	Ss	Ovec
4301a000	Eth1/27/4	0	54	2	13	13	93

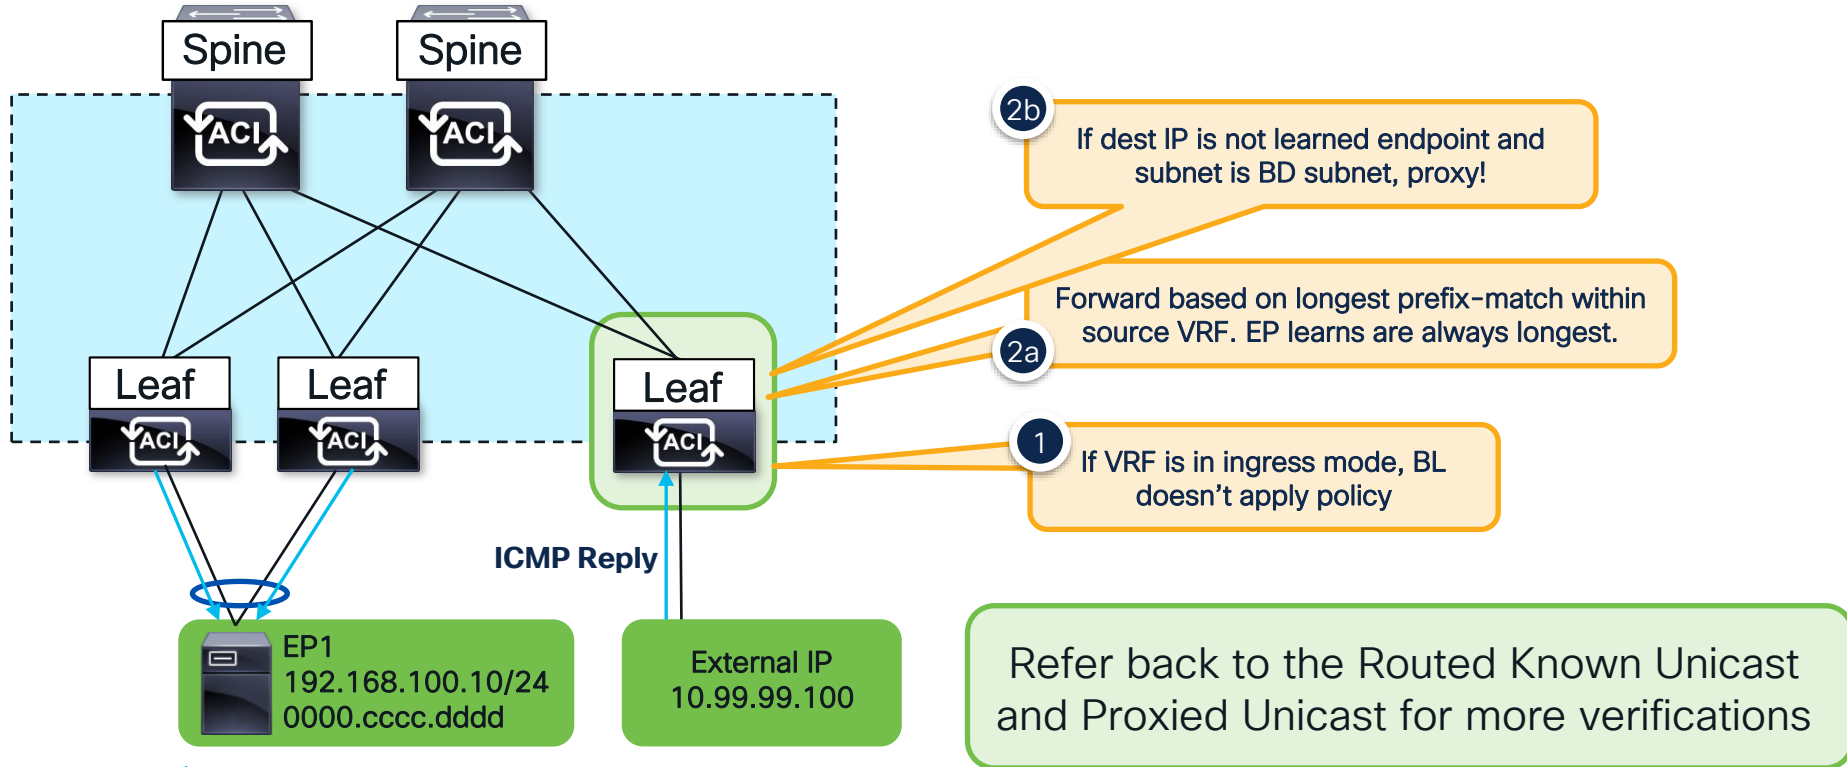
Forward out Eth1/27/4!

Unicast + Route (L3 lookup) +  
L3 Route Found



# L3Out Source – Ingress Border Leaf

Bridge Domain Settings:  
Unicast Routing Enabled



# Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.



# Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. [www.cisco.com/go/certs](http://www.cisco.com/go/certs)

## Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.



## Learn



### Cisco U.

IT learning hub that guides teams and learners toward their goals

### Cisco Digital Learning

Subscription-based product, technology, and certification training

### Cisco Modeling Labs

Network simulation platform for design, testing, and troubleshooting

### Cisco Learning Network

Resource community portal for certifications and learning



## Train



### Cisco Training Bootcamps

Intensive team & individual automation and technology training programs

### Cisco Learning Partner Program

Authorized training partners supporting Cisco technology and career certifications

### Cisco Instructor-led and Virtual Instructor-led training

Accelerated curriculum of product, technology, and certification courses



## Certify



### Cisco Certifications and Specialist Certifications

Award-winning certification program empowers students and IT Professionals to advance their technical careers

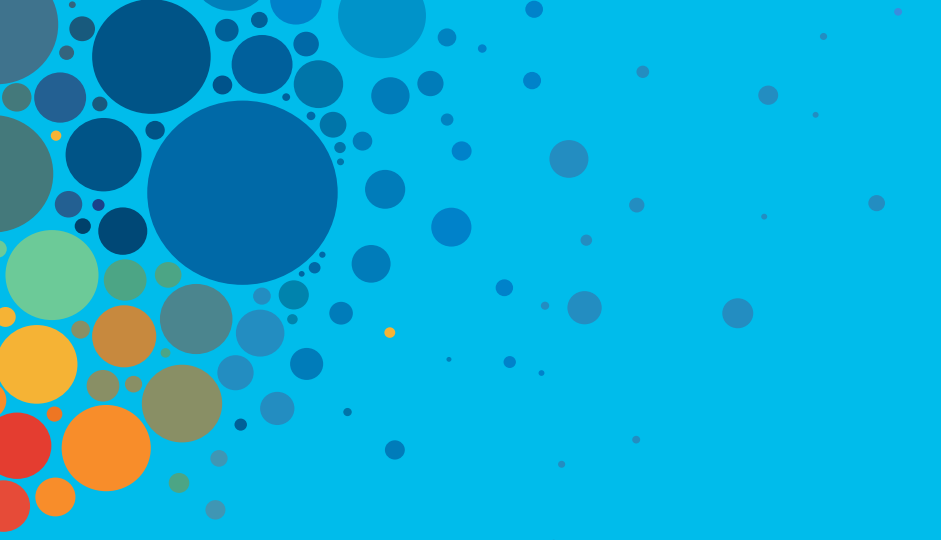
### Cisco Guided Study Groups

180-day certification prep program with learning and support

### Cisco Continuing Education Program

Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**



# Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at [www.CiscoLive.com/on-demand](http://www.CiscoLive.com/on-demand)



The bridge to possible

# Thank you

CISCO *Live!*

#CiscoLive

CISCO *Live!*

ALL IN

#CiscoLive