Troubleshoot ACI Intra-Fabric Forwarding Tools

Contents

Introduction

Background Information

What Can These Tools Help With?

SPAN and ERSPAN

ELAM

- Overview
- ASICs
- ASICs Table
- ELAM Trigger In-select
- ELAM Trigger Out-select
- ELAM set Conditions
- Viewing the ELAM Report
- Full ELAM Example
- ELAM Assistant Native/Embedded Application
  - ELAMExample

ping/iping

- Examples
- Examples

fTriage

- Examples

tcpdump

On-demand Atomic Counters

Introduction

This document describes tools natively included in ACI that can be used to debug forwarding problems.

Background Information

The material from this document was extracted from the Troubleshooting Cisco Application Centric Infrastructure, Second Edition book, specifically the***Intra-Fabric fowarding - Tools***chapter.

Additionally, deeper explanations of ELAM and Ftriage can be found in the CiscoLive On-Demand library in session BRKDCN-3900b.

What Can These Tools Help With?

In order to troubleshoot a forwarding problem from an ACI perspective, understand:

1. Which switch is receiving a flow?
2. What forwarding decision is that switch making?
3. Is the switch dropping it?

ACI includes several tools which allow the user to gain in-depth insights into what is happening to a specific flow. The next several sections demonstrate these tools in detail so only a high-level introduction is provided here.

SPAN and ERSPAN

SPAN and ERSPAN are both tools that allow all or some traffic received at a specific location to be replicated to another location. The end device that the replicated traffic is sent to is expected to be running some type of packet sniffer/analyzer application. Traditional SPAN involves replicating traffic that is being received on one port and passing out through another port. ACI supports doing this in addition to ERSPAN.

ERSPAN replicates the traffic out a local port; the replicated traffic is encapsulated in GRE and sent to a remote destination. In ACI, this ERSPAN destination must only be learned as a Layer 3 endpoint and it can be any EPG in any VRF.

It is a good idea to always have SPAN destinations connected to the fabric to minimize preparation time during troubleshooting and allow for rapid ERSPAN session config and capture.

ELAM

Overview

Embedded Logic Analyzer Module (ELAM) is a tool that allows a user to set conditions in hardware and capture the first packet or frame that matches the set conditions. A successful capture causes the ELAM status to show as "triggered." Once triggered, the ELAM is disabled and a dump can be collected to analyze the vast number of forwarding decisions that the switch ASIC is making with that packet/frame. ELAM is implemented at the ASIC level and does not impact CPU or other resources on the switch.

The forwarding examples in this book uses ELAM as a means of verifying what is happening with the flow. Examples shows both the leaf CLI version and the ELAM Assistant App.

This guide does not cover usage of ELAM on first generation leaf switches (switches without EX, FX, or FX2 suffix).

Before using the tool, it is important to understand the structure of the command syntax.

Example on leaf CLI:

```
vsh_lc                                          [This command enters the line card shell where ELAMs are r
debug platform internal <asic> elam asic 0      [refer to the ASICs table]
```

Set Conditions to **trigger**:

```
trigger reset                                                   [ensures no existing triggers are running]
trigger init in-select <number> out-select <number>             [determines what information about a packet is
set outer/inner                                                 [sets conditions]
start                                                           [starts the trigger]
status                                                          [checks if a packet is captured]
```

Generate the dump containing the packet analysis

```
ereport             [display detailed forwarding decision for the packet]
```

Continue to enter the **status** command to view the state of the trigger. Once a packet matching the defined conditions is detected on the ASIC, the output of **status** shows **triggered**. Once the ELAM has been triggered, the details of the switch forwarding decisions can be shown with 'ereport'. Prior to ACI version 4.2, **report** must be used.

ASICs

Within the ELAM syntax, note that the ASIC must be specified. Since the ASIC is dependent on the switch model, refer to this table to determine which ASIC to specify:

ASICs Table

Switch/Line card Family	Asic for Elam
-EX switches/LCs	TAH
-FX(P) switches/LCs	ROC
-FX2 switches/LCs	ROC

Switch/Line card Family	Asic for Elam
C switches (9364C,9332C)	ROC
-GX switches	APP
-GX2 switches	CHO
-FX3 switches	ROC

ELAM Trigger In-select

The other component of the ELAM that must be understood when running from the CLI is the **in-select**. The **in-select** defines which headers the packet/frame is expected to have, and which to match on.

For example, a packet coming from a downlink port that is not VXLAN encapsulated would only have outer Layer 2, Layer 3, and Layer 4 headers.

A packet coming from a front-panel (downlink) port that is VXLAN encapsulated (such as Cisco ACI Virtual Edge in VXLAN mode) or coming from an upstream spine would have VXLAN encapsulation. This means it would have potentially both outer and inner Layer 2, Layer 3, and Layer 4 headers.

All trigger options are:

<#root>

```
leaf1#
```

```
vsh_lc
```

```
module-1#
```

```
debug platform internal tah elam asic 0
```

```
module-1(DBG-elam)#
```

```
trigger reset
```

```
module-1(DBG-elam)#
```

```
trigger init in-select ?
```

```
 10  Outerl4-innerl4-ieth
 13  Outer(l2|l3|l4)-inner(l2|l3|l4)-noieth
 14  Outer(l2(vntag)|l3|l4)-inner(l2|l3|l4)-ieth
 15  Outer(l2|l3|l4)-inner(l2|l3|l4)-ieth
 6   Outerl2-outerl3-outerl4
 7   Innerl2-innerl3-innerl4
 8   Outerl2-innerl2-ieth
 9   Outerl3-innerl3
```

If **in-select 6** is selected the only option is to set conditions and display headers from the outer Layer 2, 3, or 4 headers. If **in-select 14** is selected the only option is to set conditions for and see the details of the outer and inner Layer 2, 3, and 4 headers.

Best practices note:

- To capture a packet coming with VLAN encapsulation on a downlink port, use **in-select 6**.
- To capture a packet with VXLAN encapsulation (either from a spine or from a vleaf with VXLAN encapsulation) use **in-select 14**.

ELAM Trigger Out-select

The **out-select** command allows some ability to control which lookup results are displayed in the ELAM report. For most practical purposes **out-select 0** can be used as it contains most information including the **drop vector**, which tells if the result of the lookup is to drop the packet/frame.

Note that when **report** instead of **ereport** or **report detail** is used to get ELAM results, **drop vector** only shows up in **out-select 1**. However, one can always perform **ereport** or **report detail** with **out-select 0**.

ELAM set Conditions

ELAM supports a large amount of Layer 2, 3, and 4 conditions to look for in a packet. Specifying **inner** vs. **outer** determines if the condition can be checked in the inner header (VXLAN encapsulated packet) or outer header.

ARP example:

```
set outer arp source-ip-address 10.0.0.1 target-ip-address 10.0.0.2
```

MAC address example:

```
set outer l2 src_mac aaaa.bbbb.cccc dst_mac cccc.bbbb.aaaa
```

IP address in inner header example:

```
set inner ipv4 src_ip 10.0.0.1 dst_ip 10.0.0.2
```

Viewing the ELAM Report

Verify that the ELAM has triggered with **status**:

<#root>

```
module-1(DBG-elam-insel6)#
```

status

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

ereport can be used to display the result of the ELAM in an easy to understand format. Note that the ELAM report is saved in the /var/log/dme/log/ folder on the switch. There are two files for the ELAM under the folder:

- elam_<timestamp>.txt
- pretty_elam_<timestamp>.txt

Full ELAM Example

This example would capture a non-VXLAN encapsulated traffic (matching on outer header) coming from a downlink port on an -EX switch:

<#root>

```
module-1#
```

debug platform internal tah elam asic 0

```
module-1(DBG-elam)#
```

trigger reset

```
module-1(DBG-elam)#
```

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


module-1(DBG-elam-insel6)#

set outer ipv4 src_ip 10.0.0.1 dst_ip 10.0.0.2


module-1(DBG-elam-insel6)#

start


module-1(DBG-elam-insel6)#

status

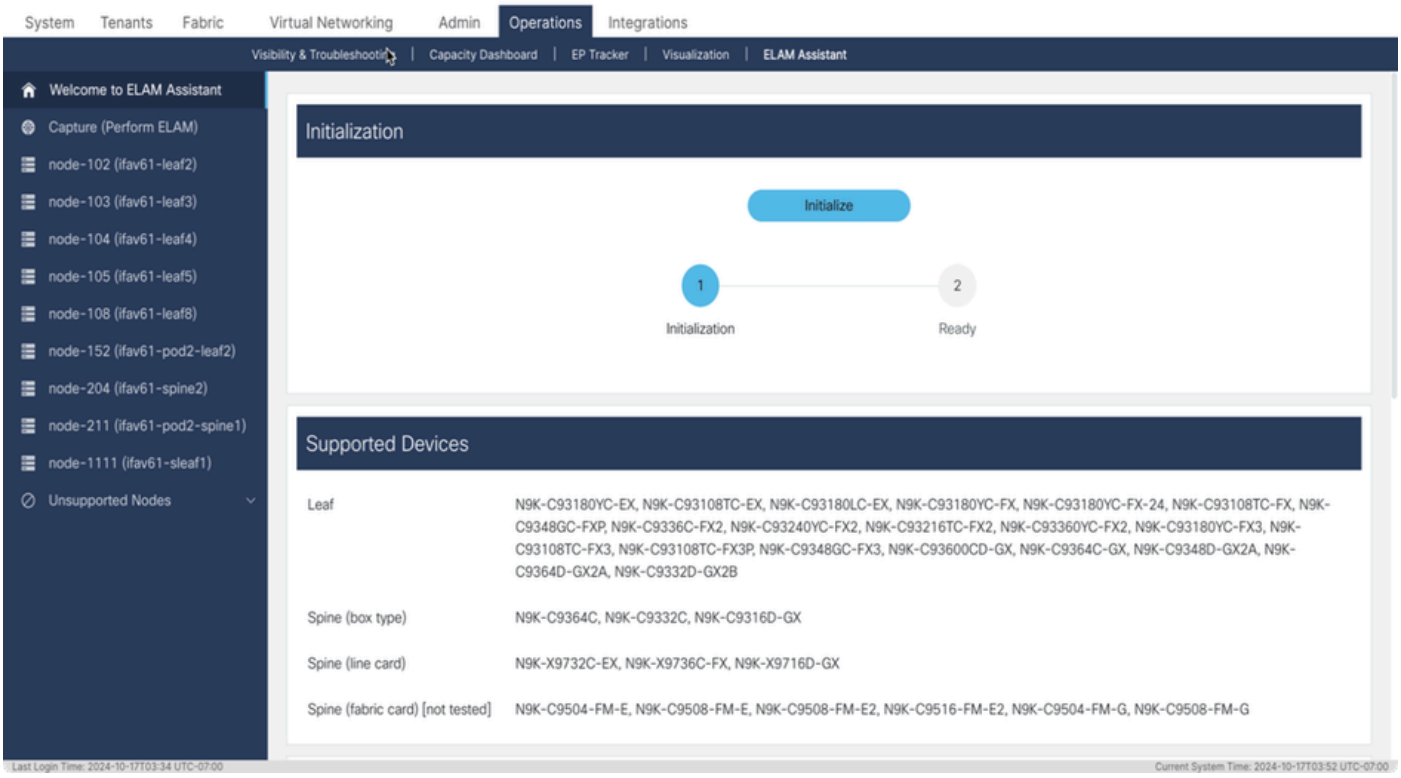

module-1(DBG-elam-insel6)#

ereport
```

ELAM Assistant Native/Embedded Application

The troubleshooting examples in this book also show the usage of the ELAM Assistant.

Formerly, it was an app which can be downloaded through the Cisco DC App Center (https://dcappcenter.cisco.com). but starting in ACI version 6.1(2)+, the ELAM Assistant is integrated directly into APIC and accessed via the Operations tab in the APIC UI.



For ELAM Assistant embedded, only APIC users with read-write privilege are able to use Elam assistant, and menu option is not be visible for users without read-write privilege.

This tool automates the deployment and interpretation of ELAMs through the GUI on the APIC.

ELAM Example

This example shows the deployment of an ELAM matching a specific source and destination IP on node-101 downlink port:

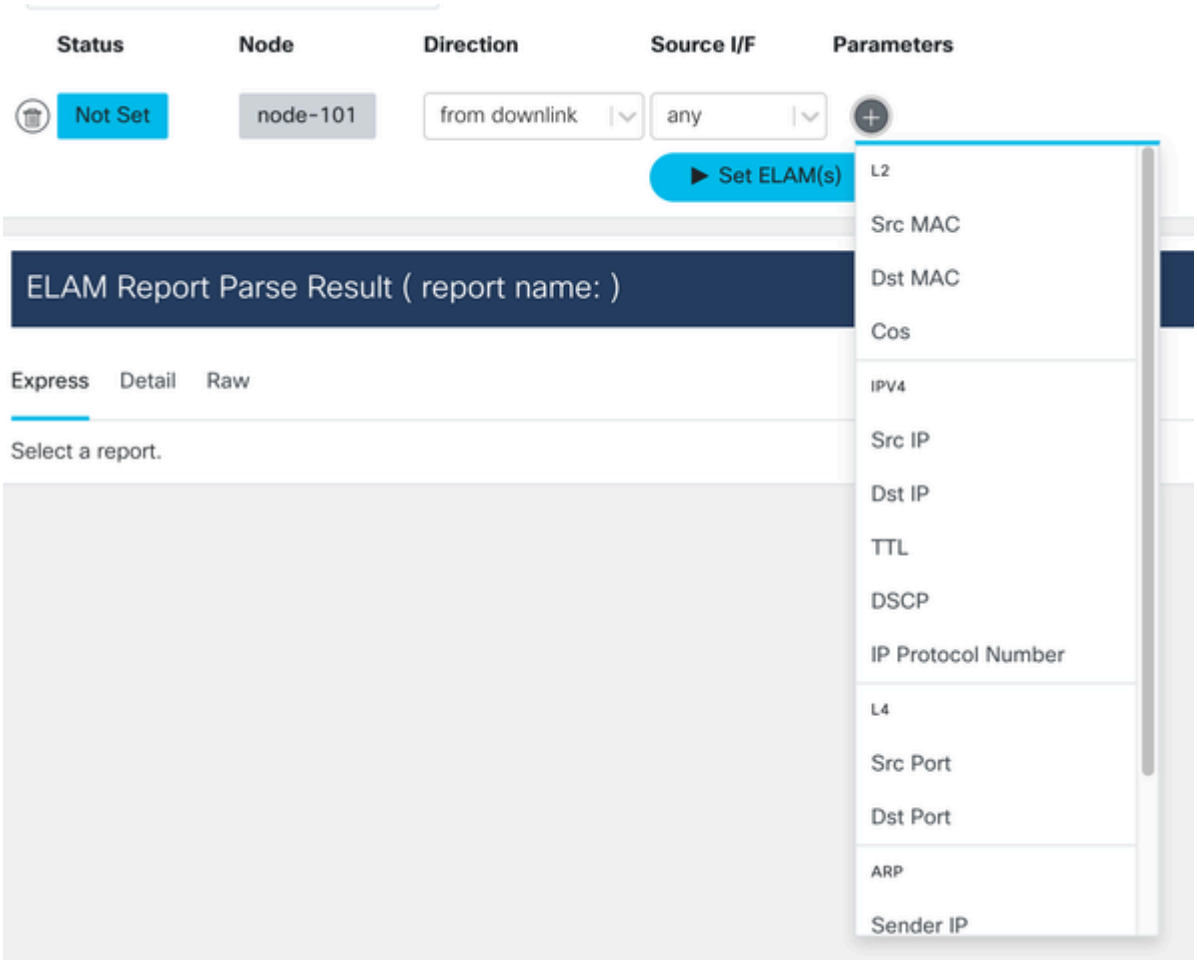
To start using it, it is needed to click on **Capture (Perform ELAM)** in the left side panel, and then in the main screen, click on **Add node** to start adding nodes where teh ELAM needs to be setup. In this example, node-101 is added:

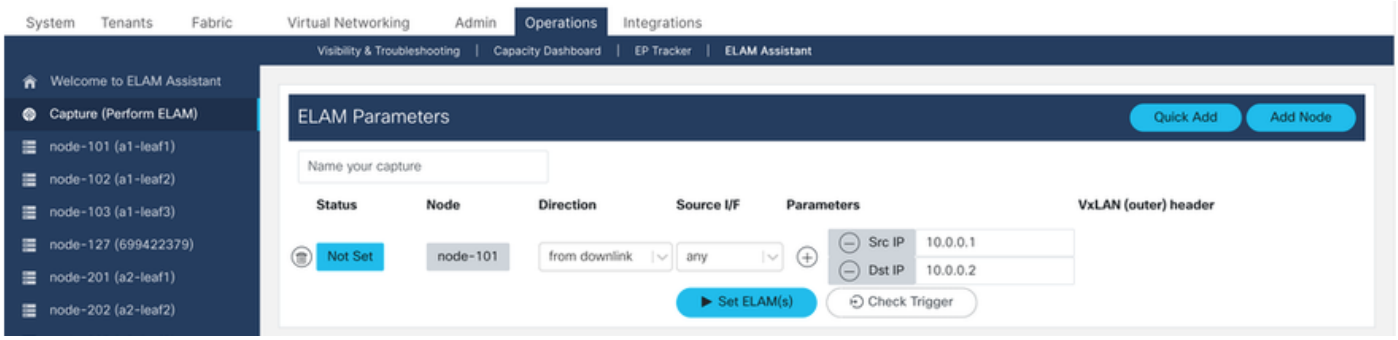


Once the nodes of interest are added, it is needed to configure the Direction (downlink/downlink VXLAN/uplink).

It is also possible to specify if any particular interface needs to be observed, the default is **any**.

Finally, and a series of parameters can be setup (similarly to the CLI version):



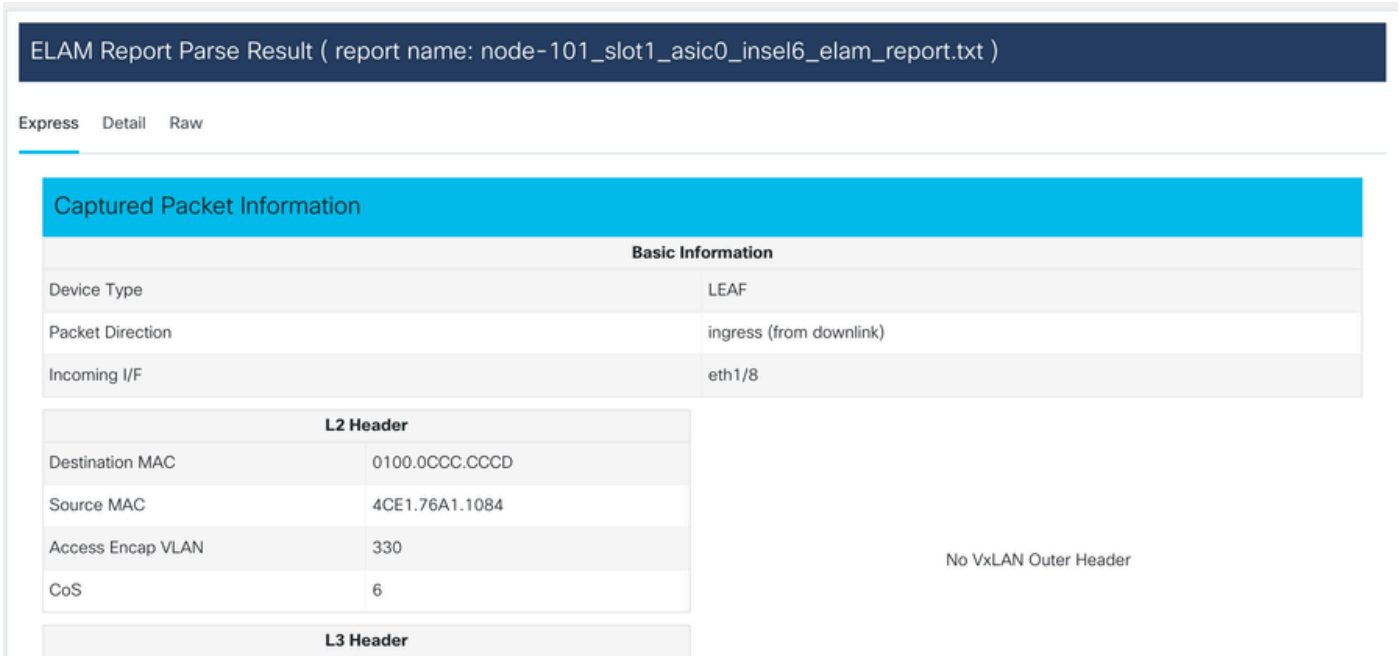


Finally, click on **Set ELAM(s)** to arm the ELAMs on the specified nodes.

To confirm if the ELAM triggered, the button **Check Trigger** can be clicked and if the ELAM gets triggered in a node, the Status changes to **Report Ready** next to the node name:



This **Report Ready** button displays in a pretty format the ELAM report below:



The output is displayed in three flavours: Express, Detail and Raw.

The ELAM Assistant also allows for easy usage of more complex matching parameters such as the source interface or VXLAN values.

ping/iping

The **ping** command is a standard network utility used in Linux to test the reachability of a host on an Internet Protocol (IP) network. It sends ICMP Echo Request packets to a specified destination (an IP address or domain name) and listens for an Echo Reply.
The main function of this command is to perform:
• Connectivity Testing: Verifies if a remote device is active and reachable.
• Latency Measurement: Calculates the round-trip time (RTT) for packets to travel to the destination and back, measured in milliseconds (ms).
• Packet Loss Detection: Reports the percentage of packets that failed to return, which helps identify network congestion or stability issues.

This utility can be used when reachability issues are identified in the APIC.

Examples

Confirm reachability of a SNMP external server (example 10.31.125.144) from the APIC OOB address:

```
ping 10.31.125.144 -I oobmgmt
```

As can be seen, the **-I** option is useful to specify the source interface, in this example the OOB management address.

It can be used the option -h to see all available options.

The **iping** command on the other hand, is a specialized diagnostic tool within the Cisco Application Centric Infrastructure (ACI) environment.

iping is a powerful network diagnostic utility used within the Cisco ACI CLI to verify connectivity from a specific node or context within the fabric. Because ACI is a multi-tenant environment, **iping** allows administrators to test connectivity while accounting for the logical isolation provided by VRFs and tenants. *How it works:* It initiates an ICMP echo request from a specified source node or context to a destination IP. It ensures that the packet traverses the fabric correctly according to the configured EPGs and contracts. To do this, the leaf switch must encapsulate the ICMP packet in the same way a production traffic would be encapsulated if it were entering the fabric. The iVXLAN header contains the VNID (Virtual Network Identifier) that maps the traffic to the correct VRF or EPG. This ensures the packet is subject to the same forwarding rules and security policies as actual production traffic.

Primary Functions:

- Context-Aware Testing: Allows you to specify the Tenant and VRF, ensuring the test mimics the actual traffic flow of a specific application.
- Fabric Visibility: Helps determine if connectivity issues are due to path misconfiguration, contract restrictions, or physical link failures within the leaf/spine architecture.
- Source Selection: Allows the user to specify which leaf switch (node) and interface originating the **iping,** which is critical for identifying localized issues.

Examples

Confirm endpoint reachability, pinging from the SVI deployed from the Bridge Domain subject (SVI)

```
iping 192.168.104.2 -V tenant_prod:VRF1
```

It is needed to select at least the VRF source, with the parameter **-V**.

When it is provided the **-V <tenant>:<vrf>** flag, the system identifies the specific VNID (VXLAN Network Identifier) associated with that VRF. This VNID is critical for ensuring the packet is routed within the correct logical routing table.

It can be used the option -h to see all available options.

fTriage

fTriage is an APIC CLI-based tool that is intended to provide end-to-end automation of ELAM configuration and interpretation. The premise of the tool is that a user can define a specific flow as well as the leaf where the flow is expected to start. The tool then executes ELAMs on each node, one by one, to examine the forwarding flow. It is particularly useful in large topologies where it is unclear which path a packet takes.

fTriage generates a large log file containing the output of each command executed. The name of this file is visible on the first few lines of the **fTriage** output.

fTriage completion can take up to 15 minutes.

Examples

Map out the flow for routed communication between 10.0.1.1 and 10.0.2.1 starting on leaf 104:

```
ftriage route -ii LEAF:104 -dip 10.0.2.1 -sip 10.0.1.1
```

Map out a Layer 2 flow starting on leaf 104:

```
ftriage bridge -ii LEAF:104 -dmac 02:02:02:02:02:02
```

Full fTriage help can be seen by running **ftriage --help** on the APIC.

tcpdump

tcpdump can be leveraged on ACI switches to capture traffic to and from the control-plane. Note that **only control plane traffic** sent to the switch CPU can be observed in a **tcpdump** capture. Some examples are: routing protocols, LLDP/CDP, LACP, ARP, and so on. To capture dataplane (and control plane) traffic please make use of SPAN and/or ELAM.

To capture on the CPU, specify the kpm_inb interface**.** Most traditional **tcpdump** options and filters are available.

Example to capture ICMP destined to an SVI on the leaf switch:

<#root>

```
leaf205#
```

```
tcpdump -ni kpm_inb icmp
```

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on kpm_inb, link-type EN10MB (Ethernet), capture size 65535 bytes
20:24:12.921981 IP 10.0.2.100 > 10.0.2.1: ICMP echo request, id 62762, seq 4096, length 64
20:24:12.922059 IP 10.0.2.1 > 10.0.2.100: ICMP echo reply, id 62762, seq 4096, length 64
20:24:13.922064 IP 10.0.2.100 > 10.0.2.1: ICMP echo request, id 62762, seq 4352, length 64
20:24:13.922157 IP 10.0.2.1 > 10.0.2.100: ICMP echo reply, id 62762, seq 4352, length 64
20:24:14.922231 IP 10.0.2.100 > 10.0.2.1: ICMP echo request, id 62762, seq 4608, length 64
20:24:14.922303 IP 10.0.2.1 > 10.0.2.100: ICMP echo reply, id 62762, seq 4608, length 64
```

In addition, the **-w** option allows the tcpdump to write the packet capture to a PCAP file so that it can be opened in tools such as Wireshark.

To use **tcpdump** on the eth0 interface, which is the out-of-band interface on the switch. This is useful to troubleshoot connectivity of any traffic going through the out-of-band physical port of the switch. This would mainly be control plane-based traffic such as SSH, SNMP, and so on.

On-demand Atomic Counters

On-demand atomic counters are intended to count packets within a specific flow as they leave on a leaf uplink and are received on another leaf fabric port. They allow some granularity into whether packets were missed or received in excess.