

Troubleshoot Link Aggregation Control Protocol (LACP) on Nexus

Contents

Introduction

This document describes how to troubleshoot Link Aggregation Control Protocol (LACP) on Nexus 9000 cloudscale family.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:


- LACP protocol
- NXOS Platform
- ELAM understanding
- Ethalyzer understanding

Components Used

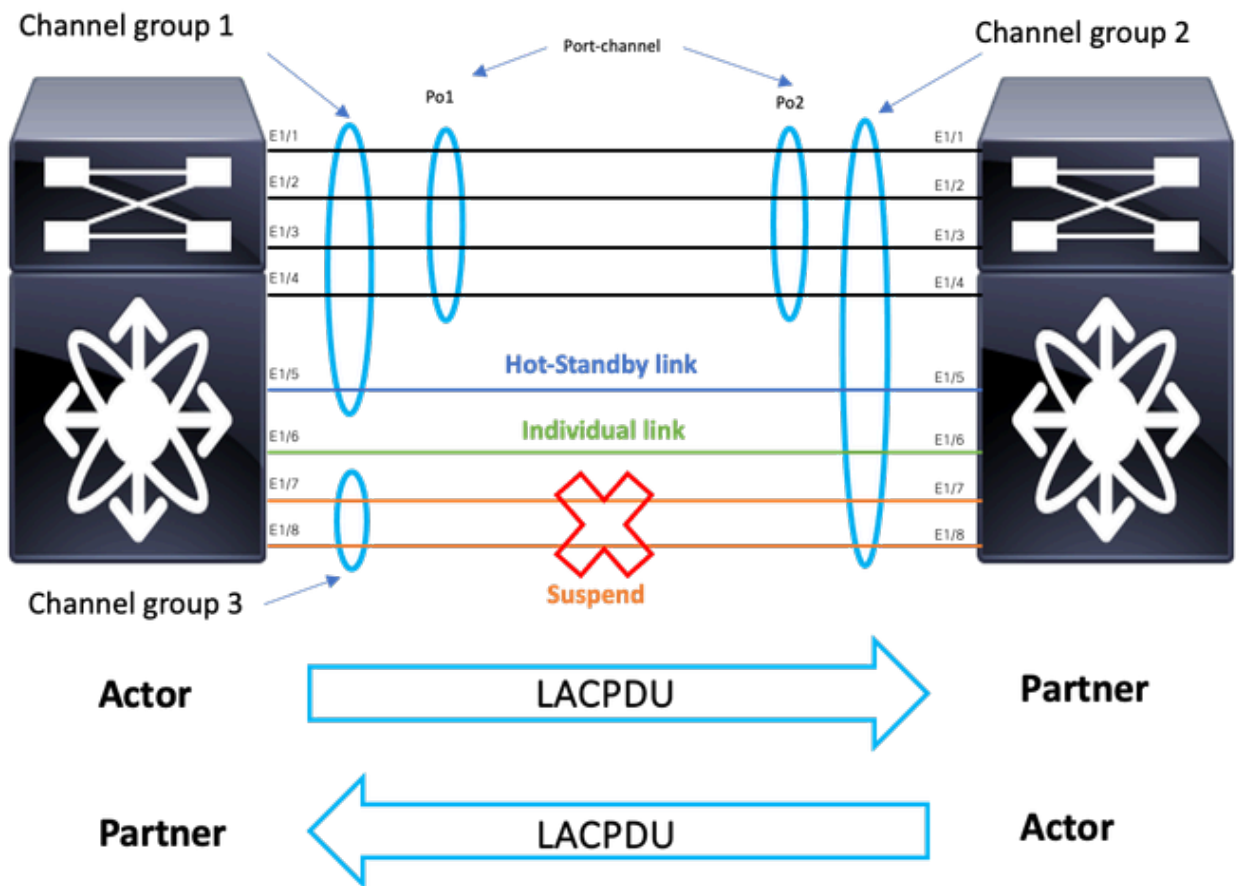
Name	Platform	Version
N9K1	N9K-C93108TC-EX	9.3(10)
N9K2	N9K-C93108TC-EX	9.3(10)


The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Topology



Verify LACP Link Status





Note: Image 1.1 LACP link status.


Configure LACP:

N9K1	N9K2
<pre>show run interface port-channel 1 membership interface port-channel1 switchport switchport mode trunk interface Ethernet1/1 switchport switchport mode trunk channel-group 1 mode active no shutdown interface Ethernet1/2 switchport switchport mode trunk channel-group 1 mode active</pre>	<pre>show run interface port-channel 1 membership interface port-channel1 switchport switchport mode trunk interface Ethernet1/1 switchport switchport mode trunk channel-group 1 mode active no shutdown interface Ethernet1/2 switchport switchport mode trunk channel-group 1 mode active</pre>

no shutdown	no shutdown
-------------	-------------

Verify port-channel status

N9K1							
<pre>sh port-channel summary interface port-channel 1 Flags: D - Down P - Up in port-channel (members) I - Individual H - Hot-standby (LACP only) s - Suspended r - Module-removed b - BFD Session Wait S - Switched R - Routed U - Up (port-channel) p - Up in delay-lacp mode (member) M - Not in use. Min-links not met</pre>					<pre>sh port-channel summary Flags: D - Down I - Individual s - Suspended b - BFD Session S - Switched U - Up (port-ch p - Up in delay M - Not in use.</pre>		
Group	Port-Channel	Type	Protocol	Member Ports	Group	Port-Channel	Type
1	Po1(SU)	Eth	LACP	Eth1/1(P)	1	Po1(SU)	Eth



Note: Most common failure scenarios is Nexus suspending the interface, which is covered in **LACP Suspended interface** section.

Verify LACP Suspended interface

```
sh port-channel summary interface port-channel 1
```

Flags: D - Down P - Up in port-channel (members)
I - Individual H - Hot-standby (LACP only)
s - Suspended r - Module-removed
b - BFD Session Wait
S - Switched R - Routed
U - Up (port-channel)
p - Up in delay-lacp mode (member)
M - Not in use. Min-links not met

Group	Port-Channel	Type	Protocol	Member Ports
1	Po1(SD)	Eth	LACP	Eth1/1(s)

```
sh int e1/1 status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	--	suspended	trunk	auto	auto	10Gbase-SR

```
sh int e1/1
```

```
Ethernet1/1 is down (suspended(no LACP PDUs))
```

```
admin state is up, Dedicated Interface
```


```
Belongs to Po1
```

```
Hardware: 100/1000/10000/25000 Ethernet, address: 003a.9c08.68ab (bia 003a.9c08.68ab)
```

```
MTU 9216 bytes, BW 10000000 Kbit , DLY 10 usec
```

```
reliability 255/255, txload 1/255, rxload 1/255
```

```
<Snipped>
```



Note: When facing this scenario Nexus is not receiving LACP PDUS from partner, LACP interfaces counters can be verified, as written section: **Verify LACP interface counters** or packet captures such as SPAN or ELAM (Described on section **LACP ELAM**) can be taken.

Verify LACP interface counters


Each device must send and receive LACPDU at the same rate in order for the port-channel to come up.

N9K1								
sh lacp counters interface port-channel 1 NOTE: Clear lacp counters to get accurate statistics							sh lacp counters interface NOTE: Clear lacp counter	
-----							-----	
Port	Sent	LACPDUs		Markers/Resp LACPDUs			Port	Sent
		Recv		Recv	Sent	Pkts Err		
-----							-----	
port-channel1							port-channel1	
Ethernet1/1	445		445	0	0	0	Ethernet1/1	445
Ethernet1/2	445		445	0	0	0	Ethernet1/2	445

Verify LACP Actor State Bits

In each LACP PDU Actor state information is exchanged between Partner and Actor.

Activity	1: Active mode	0: Passive mode
Timeout	1: Short timeout	0: Long Timeout
Aggregation	1: Aggregatable	0: Individual
Sync	1: In Sync	0: Out of Sync
Collecting	1: Collecting Enabled	0: Collecting Disabled
Distributing	1: Distributing Enabled	0: Distributing Disabled
Defaulted	1: Use default for partner	0: Use rx LACPDU for partner
Expired	1: Partner PDU Expired	0: Not Expired




Note: Table 2.0 Actor State Bits meaning

Verify LACP State Hexadecimal Value:

state: **0x3d** (Ac-1 To-0 Ag-1 Sy-1 Co-1 Di-1 De-0 Ex-0)

		State							
		Ex	De	Di	Co	Sy	Ag	To	Ac
0x3d=		0	0	1	1	1	1	0	1



Note: Image 3.0 LACP state conversion from binary to Hexadecimal

Verify LAG ID

Link Agregation Identifier is the information that each physical interface member of the same port-channel are sharing, to appear as a single "virtual interface" . It can be verified using commands.

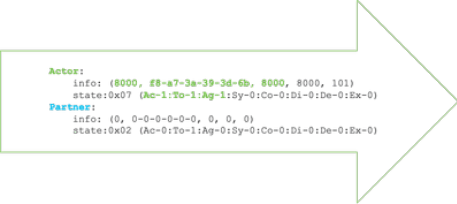
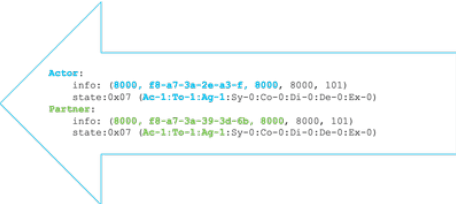
N9K1 LAG ID	N9K2 LAG ID
<pre>sh lacp interface e1/1 include ignore local lag Lag Id: [[(1770, a8-c-d-96-43-7f, 5, 8000, 1c9), (1770, a8-c-d-96-c9-bf, 5, 8000, 1c9)]] Local Port: Eth1/1 MAC Address=a8-c-d-96-c9-bf sh lacp interface e1/2 include ignore local lag Lag Id: [[(1770, a8-c-d-96-43-7f, 5, 8000, 1ca), (1770, a8-c-d-96-c9-bf, 5, 8000, 1ca)]] Local Port: Eth1/2 MAC Address=a8-c-d-96-c9-bf</pre>	<pre>sh lacp inter Lag Id: [[(1 Local Port: E sh lacp inter Lag Id: [[(1 Local Port: E</pre>

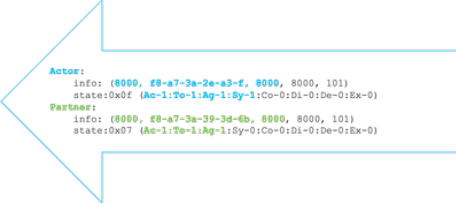
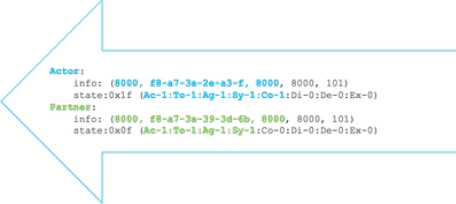
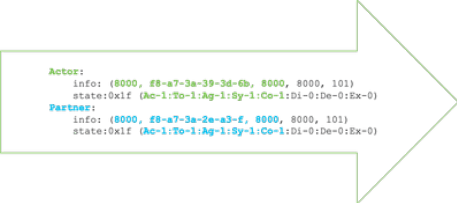
Verify LACP PDU Exchange

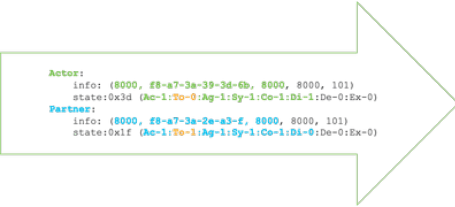
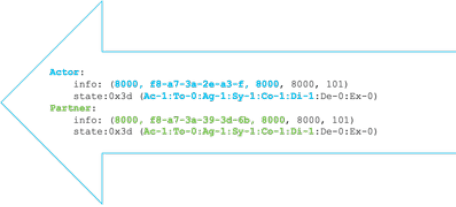
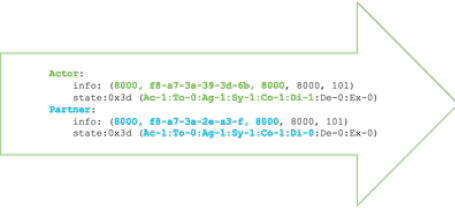
There are certain scenarios where even if Nexus is exchanging LACP PDUs at correct rate, port-channel does not come up.

This could be due a failure on LACP negotiation.

An example of a correct LACP transaction for a port-channel to come up is showed in this table.

1	N9K1 Actor	N9K2 Partner
<ul style="list-style-type: none"> N9K1 send LACP with actor information with the state bits. Partner information is at 0s since N9K1 has not received any LACP PDU from partner. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) Partner: info: (0, 0-0-0-0-0-0, 0, 0, 0) state:0x02 (Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0) </pre>	
2	N9K1 Partner	N9K2 Actor
<ul style="list-style-type: none"> N9K2 receive information from N9K1. N9K2 send LACP PDU with its information and acknowledge of N9K1 information . 		 <pre> Actor: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) </pre>
3	N9K1 Actor	N9K2 Partner
<ul style="list-style-type: none"> N9K1 Acknowledge N9K2 information. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) </pre>	
4	N9K1 Partner	N9K2 Actor

<ul style="list-style-type: none"> N9K2 receives acknowledge from N9K1. N9K2 sends LACP PDU adding SYNC bit at 1. 		 <pre> Actor: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0) </pre>
5	N9K1 Actor	N9K2 Partner
<ul style="list-style-type: none"> N9K1 acknowledge SYNC bit from N9K2. N9K1 adds SYNC bit at 1 to LACP PDU. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0) </pre>	
6	N9K1 Partner	N9K2 Actor
<ul style="list-style-type: none"> N9K2 Acknowledge SYNC bit from N9K1. N9K2 adds Collect bit at 1 to LACP PDU. 		 <pre> Actor: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x1f (Ac-1:To-1:Ag-1:Sy-1:Co-1:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0) </pre>
7	N9K1 Actor	N9K2 Partner
<ul style="list-style-type: none"> N9K1 Acknowledge Collect bit from N9K2. N9K1 Adds Collect bit to its LACP PDU. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x1f (Ac-1:To-1:Ag-1:Sy-1:Co-1:Di-0:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x1f (Ac-1:To-1:Ag-1:Sy-1:Co-1:Di-0:De-0:Ex-0) </pre>	
8	N9K1 Actor	N9K2 Partner

<ul style="list-style-type: none"> N9K1 decides that is ready to transition to distribute state, so it changes now Time Out bit from 1 (fast) to 0 (Slow) and sets Distribute bit to 1. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x3d (Ac-1:To-0:Ag-1:Sy-1:Co-1:Ds-1:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x1f (Ac-1:To-1:Ag-1:Sy-1:Co-1:Ds-0:De-0:Ex-0) </pre>	
9	N9K1 Partner	N9K2 Actor
<ul style="list-style-type: none"> N9k2 acknowledge N9k1 PDU and changes its Time Out bit from 1 to 0 and sets distribute bite to 1. At this moment both nexus are ready to send data on the port-channel. 		 <pre> Actor: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x3d (Ac-1:To-0:Ag-1:Sy-1:Co-1:Ds-1:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x3d (Ac-1:To-0:Ag-1:Sy-1:Co-1:Ds-1:De-0:Ex-0) </pre>
10	N9K1 Actor	N9K2 Partner
<ul style="list-style-type: none"> N9K1 Acknowledge LACP PDU from N9K2. At this moment Port-Channel woll transition to up. 	 <pre> Actor: info: (8000, E8-a7-3a-39-3d-6b, 8000, 8000, 101) state:0x3d (Ac-1:To-0:Ag-1:Sy-1:Co-1:Ds-1:De-0:Ex-0) Partner: info: (8000, E8-a7-3a-2e-a3-f, 8000, 8000, 101) state:0x3d (Ac-1:To-0:Ag-1:Sy-1:Co-1:Ds-1:De-0:Ex-0) </pre>	

Verify LACP FSM log

LACP Finite State Machine has a dedicated logg where it stores all the events of the interface LACP states and LACP PDUS can be found on this log:

```
sh lacp internal info interface e1/1 detail fsmlog
```

On new versions you can also use:

```
sh lacp internal event-history interface e1/1
```

First section describes interface LACP transitions

```
>>>>FSM: <Ethernet1/1> has 61 logged transitions<<<<<
<Snipped>
```

- 58) FSM:<Ethernet1/1> Transition at 127198 usecs after Mon Aug 14 22:34:42 2023
 Previous state: [LACP_ST_WAIT_FOR_HW_TO_PROGRAM_RECEIVE_PATH]
 Triggered event: [LACP_EV_PORT_RECEIVE_PATH_ENABLED_AS_CHANNEL_MEMBER_MESSAGE]
 Next state: [LACP_ST_PORT_MEMBER_RECEIVE_ENABLED]

- 59) FSM:<Ethernet1/1> Transition at 127227 usecs after Mon Aug 14 22:34:42 2023
 Previous state: [LACP_ST_PORT_MEMBER_RECEIVE_ENABLED]
 Triggered event: [LACP_EV_PARTNER_PDU_IN_SYNC_COLLECT_ENABLED_DISTRIBUTING_DISABLED]
 Next state: [LACP_ST_WAIT_FOR_HW_TO_PROGRAM_TRANSMIT_PATH]

- 60) FSM:<Ethernet1/1> Transition at 128265 usecs after Mon Aug 14 22:34:42 2023
 Previous state: [LACP_ST_WAIT_FOR_HW_TO_PROGRAM_TRANSMIT_PATH]
 Triggered event: [LACP_EV_PERIODIC_TRANSMIT_TIMER_EXPIRED]
 Next state: [FSM_ST_NO_CHANGE]

- 61) FSM:<Ethernet1/1> Transition at 134352 usecs after Mon Aug 14 22:34:42 2023
 Previous state: [LACP_ST_WAIT_FOR_HW_TO_PROGRAM_TRANSMIT_PATH]
 Triggered event: [LACP_EV_PORT_HW_PATH_ENABLED]
 Next state: [LACP_ST_PORT_MEMBER_COLLECTING_AND_DISTRIBUTING_ENABLED]
 Curr state: [LACP_ST_PORT_MEMBER_COLLECTING_AND_DISTRIBUTING_ENABLED]

Second section shows all the LACP PDUS information that exus has sent or received.

<Snipped>

```
(1) Send LACP PDU: len:110 at 492243 usecs after Tue Aug 15 00:02:13 2023
01010114 8000f8a7 3a393d6b 80008000 01013d00 00000214 8000f8a7 3a2ea30f
80008000 01013d00 00000310 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00000000 0000
```

Actor:

```
info: (8000, f8-a7-3a-39-3d-6b, 8000, 8000, 101)
state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0)
```

Partner:

```
info: (8000, f8-a7-3a-2e-a3-f, 8000, 8000, 101)
state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0)
```


```
(2) Recv LACP PDU: len:124 at 708749 usecs after Tue Aug 15 00:02:12 2023
0180c200 0002f8a7 3a2ea310 88090101 01148000 f8a73a2e a30f8000 80000101
3d000000 02148000 f8a73a39 3d6b8000 80000101 3d000000 03100000 00000000
00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000 00000000 00000000 00000000
```

Actor:

```
info: (8000, f8-a7-3a-2e-a3-f, 8000, 8000, 101)
state:0x0f (Ac-1:To-1:Ag-1:Sy-1:Co-0:Di-0:De-0:Ex-0)
```

Partner:

```
info: (8000, f8-a7-3a-39-3d-6b, 8000, 8000, 101)
state:0x07 (Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
```



Note: Please take in consideration that LACP PDU event history time line goes from newer logs to older logs however LACP states event-history time line goes from older logs to newer logs.

Configure and verify LACP Ethanalyzer

LACP PDUS needs to be proceced by CPU , nexus has installed in hardware internal access list to redirect LACP packets to CPU, all LACP PDUS can be observed with ethanalyzer, in order to filter them wireshark filter "**slow**" can be used.

```
ethanalyzer local interface inband display-filter "slow and eth.addr==04:76:b0:b2:00:20 and eth.addr==0  
Capturing on inband  
2023-07-03 23:37:14.420839 04:76:b0:b2:00:20 -> 01:80:c2:00:00:02 LACP Link Aggregation Control Protocol
```

Detailed:

```
Frame 19 (124 bytes on wire, 124 bytes captured)  
Arrival Time: Jul 3, 2023 23:38:14.425502000  
[Time delta from previous captured frame: 0.836575000 seconds]
```


Configure and verify LACP Elam

First mac address from remote interface needs to be identified:

```
N9K2#sh int e1/1 | i i addr
Hardware: 1000/10000/25000/40000/50000/100000 Ethernet, address: 0476.b0b2.0020 (bia 0476.b0b2.0020)
```

Now on nexus N9K1 elam is configured.

```
N9K1#debug platform internal tah elam
N9K1(TAH-elam)#trigger init
N9K1(TAH-elam-insel6)#set outer l2 src_mac 04:76:b0:b2:00:20 dst_mac 01:80:c2:00:00:02 <<<<<Dest mac is
N9K1(TAH-elam-insel6)#tart
N9K1(TAH-elam-insel6)#report
SUGARBOWL ELAM REPORT SUMMARY
slot - 1, asic - 0, slice - 1
=====Incoming Interface: Eth1/49
Src Idx : 0x601, Src BD : 1
Outgoing Interface Info: met_ptr 0Packet Type: CE

Dst MAC address: 01:80:C2:00:00:02
Src MAC address: 04:76:B0:B2:00:20
Sup hit: 1, Sup Idx: 2627. <<<<<Traffic needs to be punted to the CPU.
Drop Info:
-----LUA:
LUB:
LUC:
LUD:
Final Drops:vntag:
vntag_valid : 0
vntag_vir : 0
vntag_svif : 0
```

In order to decode the sup redirect index, comand sh system internal access-list sup-redirect-stats can be excecuted:

```
sh system internal access-list sup-redirect-stats | i i 2627
2627 LACP 0
2627 LACP 103
```