Troubleshoot ACI Dual Role Switch and Intel VIC Card

Contents

Introduction

Problem

Solution

1. Spine Switch N9K-C93600CD-GX is not Coming up in The Cluster

2. APIC GUI Shows Wrong Interface Numbers Connected

3. APIC GUI Shows the Same MAC Address for All Interfaces (eth2-1, eth2-2, eth2-3, eth2-4)

Introduction

This document describes the use of dual-role switches in the Application Centric Infrastructure (ACI) fabric and the use of Intel Virtual Interface Cards (VICs).

Problem

The issues reported:

1. Spine Switch N9K-C93600CD-GX is not coming up in the cluster.

2. Application Policy Infrastructure Controller (APIC) GUI shows the wrong interface numbers connected.

3. APIC GUI shows the same MAC address for all interfaces (eth2-1, eth2-2, eth2-3, eth2-4).

Solution

1. Spine Switch N9K-C93600CD-GX is not Coming up in The Cluster

It has been observed that a very generic approach has been used when these kinds of issues are reported. These are basic troubleshooting steps which can be performed for isolation but must be performed after checking the installation guide for the product and making sure that the current settings and requirements match.

i. Moving the connections either on the switch side or the APIC side is done.

ii. Reloading the switch or APIC is done.

iii. Additional CLI commands are collected or sometimes tech support logs are collected to further investigate the issue.

All these steps are correct and must be followed. But there is another step that can be checked whenever there is a discovery issue with one specific Part Identifier (PID). That basic check is to go through the hardware installation guide of that specific switch.

For example, a user had an issue with switch PID N9K-C93600CD-GX and the user was trying to bring it

up as a spine and It was connected to the leaf switch via its own port number 20. This spine switch never came up.

In the installation guide,this informatio can be found:

- This switch default role is as a leaf switch.
- The default fabric links (ports 29-36) must be used for initial switch discovery via another switch.
- In order to change the switch from the default role, you must proceed as follows: the node appears as a discovered device in the fabric inventory view, you must set the role of the switch (spine or leaf), and the switch automatically goes for reboot to come up in the configured role.
- If you connect a default spine (a dual role switch that by default is a spine, such as Nexus 9316D-GX) directly to an APIC, the change of the role to leaf is performed automatically by APIC as well as the reboot. After that, the node appears in Nodes pending registration and you need to register the node.

Always check the sections like Leaf/Spine role considerations and discovery considerations before performing any additional checks.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/hw/aci-93600cd-gx/guide/b_c93600CD-GX-aci-mode-hardware-installation-guide/b_c93600CD-GX-aci-mode-hardware-installation-guide_chapter_01.html.

Similarly, for PID C9316D-GX, the default role is the spine. But it can also work as a leaf in the fabric.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/hw/aci_9316D-GX_hig/guidebook/b_C9316D-GX_aci_hardware_installation_guide/m_overview_nx-os.html.

More examples are PID 9332D-GX2B and 9364C-GX, the default role is leaf but can work as a spine.

9332D-GX2B must be connected via its own ports range of 25-32.

9364C-GX must be connected via its own ports range of 49-62.

So, always check the hardware installation guide before proceeding with any other troubleshooting step, as it saves time.

2. APIC GUI Shows Wrong Interface Numbers Connected

It was reported that the physical cable connections were made on APIC interfaces eth2-1 and eth2-3 but in the APIC GUI, it was found that interfaces eth2-2 and eth2-4 were showing up. Similar behavior was seen for all the three APICs in the cluster.



eth2-1	1500	50:7C:6F:31:6F:C8	down
eth2-2	1500	50:7C:6F:31:6F:C8	up
eth2-3	1500	50:7C:6F:31:6F:C8	down
eth2-4	1500	50:7C:6F:31:6F:C8	up

The user was using the Peripheral Component Interconnect Express (PCIe) slot - APIC-PCIE-IQ10GC Intel

X710 quad Port 10GBase-T network interface, not the Cisco VIC cards.

Verify in Cisco Integrated Management Controller (CIMC), chassis > Inventory > PCI Adapters.



It was confirmed that the NIC mode in CIMC is dedicated to all three CIMCs, along with the Trusted Platform Module (TPM) status enabled and ownership owned. All other outputs look fine too. Later, a software defect (Cisco bug ID CSCwd21587) was filed to further troubleshoot the issue.

It was found that:

Intel X710-T4 Quad-port 10GBase-T NIC, it has the port numbering scheme starting on the right port and incrementing going towards the left port.

Generally, the connections in any fabric setup are made with the assumption that numbering starts from the left, which is not applicable here.

Ports are assigned numbers from right to left in this manner:

```
| eth2-4 | eth2-3 | eth2-2 | eth2-1 |
```

With this numbering scheme in mind, APIC is detecting correct ports and showing ports on GUI as expected.

It has been updated in the document as well: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/server/M3-L3-server/APIC-M3-L3-Server.pdf.

Points to remember:

- APIC-PCIE-IQ10GC or UCSC-PCIE-IQ10GC must always be installed in PCIe slot 1 for APIC M3/L3.
- APIC-PCIE-IQ10GC or UCSC-PCIE-IQ10GC can use any port or any pair of ports in order to connect to a leaf node.
- APIC-PCIE-IQ10GC or UCSC-PCIE-IQ10GC have port numbering in the order | eth2-4 | eth2-3 | eth2-2 | eth2-1 | and the numbering on chassis is not valid.
- From Release 4.2(5) the UCSC-PCIE-IQ10GC Intel X710 Quad Port 10GBase- network interface card is supported for 10GBast-T connectivity to Cisco ACI leaf nodes.

3. APIC GUI Shows the Same MAC Address for All Interfaces (eth2-1, eth2-2, eth2-3, eth2-4)

It was observed that APIC GUI was not showing the correct MAC address for each Interface. All the MAC addresses were the same.



Always remember that there is an active/backup teaming between the interfaces, so you must always see the MAC address of the active interface which can be assigned to the bond0 interface, hence you see the same MAC address.

Here, you are seeing the MAC address of the down interface as per the output:

```
Slave Interface: eth2-1
MII Status: down
Speed: Unknown
Duplex: Unknown
Link Failure Count: 6
Permanent HW addr: 50:7c:6f:31:72:50
Slave queue ID: 0

Slave Interface: eth2-2
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 6
Permanent HW addr: 50:7c:6f:31:72:51
Slave queue ID: 0

Slave Interface: eth2-3
MII Status: down
Speed: Unknown
Duplex: Unknown
Link Failure Count: 5
Permanent HW addr: 50:7c:6f:31:72:52
Slave queue ID: 0

Slave Interface: eth2-4
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 7
Permanent HW addr: 50:7c:6f:31:72:53
Slave queue ID: 0
```

This is a software defect that is documented under the Cisco bug ID CSCwd21587.

Ideally, you must see the MAC address of eth2-2 which is the active and eth2-4 is the backup here.

This is a backend issue as the list in UI is made based on the moquery cnwPhysIf. This moquery cnwPhysIf also shows the MAC address of the down interface eth2-1.

The workaround is to use the command cat /proc/net/bonding/bond0 on APIC in order to check the correct MAC addresses. For a permanent fix, check the software defect page.