

Nexus 9500 Spine Chassis Hardware Failure Recovery



Document ID: 118864

Contributed by Robert Correiro, Cisco TAC Engineer.

Mar 26, 2015

Contents

Introduction

Problem

Solution

Introduction

This document describes the process to replace the chassis of a Nexus 9500 spine switch, which failed due to a hardware issue, in Application Centric Infrastructure (ACI) mode.

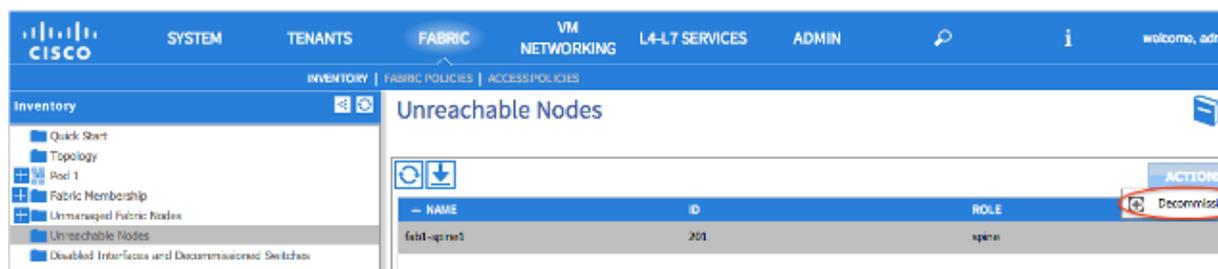
Problem

There is a current and working ACI fabric and a Nexus 9500 spine has failed. The failure has been determined to be a hardware failure and only the chassis needs to be replaced. All working components (supervisors, line cards, fabric modules, and so on) need to be transferred into the replacement.

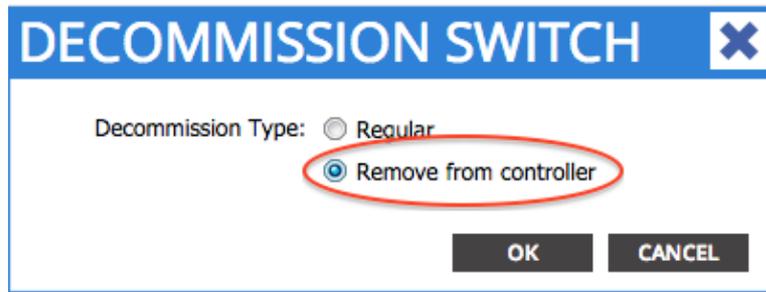
Solution

Complete these steps in order to replace the spine:

1. Disconnect all power cables in order to ensure the failed spine is powered off.
2. From the Application Policy Infrastructure Controller (APIC) GUI, choose **Fabric > Inventory > Unreachable Nodes**. The spine is listed here after it has been powered off for a few minutes. Take note of its 'Name' and 'Node ID'.
3. Decommission the failed spine and remove it from the controller.
 - a. Select the spine in the work pane. From the Actions drop-down list, choose **Decommission**.



- a. Click the **Remove from controller** radio button and click **OK**.



Tip: The 'Remove from controller' option completely removes the node from the ACI fabric and the serial number is disassociated from the Node ID. The 'Regular' option is used in order to temporarily remove the node from the ACI fabric, with the expectation that the same node will rejoin the fabric with the same Node ID. For instance, if the node needs to be temporarily powered down for maintenance.

Note: After a node is decommissioned it can take 5–10 minutes for it to be removed and disappear from the APIC GUI.

4. Remove the failed spine from the rack and install the replacement chassis. Transfer all working components into the new chassis in the same arrangement as in the failed chassis. Components to transfer include fabric modules, supervisors, system controllers, power supplies, fan trays, and line cards.
5. Console into the active supervisor in the new chassis in order to erase its configuration.
 - a. Enter the `cat /proc/cmdline` command in order to determine which ACI image currently runs on the supervisor. Look for the 'ksimg' section.

```
(none)# cat /proc/cmdline
console=ttyS0,9600n8nn card_index=21000 loader_ver="8.06" quiet ksimg=bootflash:aci-n9000-dk9.11.0.2j.bi
rw root=/dev/ram0 rdbase=0x8000000 ip=off ramdisk_size=131072 kgdboc=ttyS0,115200,B mtdparts=physmap-flash.0
12k(mtdoops),256k(RR),256k(SM_LOG),512k(KLOG),512k(EXTRA),12m(KTRACES),50m(PLOG) elevator=noop intel_idle.m
cstate=2 pcie_ports=native
```

Tip: ACI switch image names always begin with 'aci-n9000'.

- b. Use the ACI switch image name and enter these commands:

```
spine# /bin/prepare-mfg.sh <aci_image.bin>
spine# reload
```

Note: `<aci_image.bin>` is the filename of the ACI switch image located on the APIC.

- c. If the reload window was missed and you are unable to reload the device, enter the `vsh -c reload` command.

This erases the configuration on the active supervisor. Repeat Step 5 in order to erase the configuration on the standby supervisor. The configuration is automatically pushed down from the APIC after the spine joins the fabric.

6. In the APIC GUI, choose **Fabric > Inventory > Fabric Membership**. The spine is listed as a row in the work pane without a Node ID, Node Name, and IP address. Double-click the row and assign the same 'Node ID' and 'Node Name' as before.

The screenshot shows the Cisco APIC GUI with the 'Fabric Membership' page. The table below represents the data shown in the interface:

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMMISSIONED	SUPPORTED MODEL
FGE	201	fab1-spine1	spine1	N9K-C9508	spine	0.0.0.0	False	True
FGE	202	fab1-spine2		N9K-C9508	spine	192.168.0.64/32	False	True
SAL	102	fab1-leaf2		N9K-C9312RFX	leaf	192.168.0.127/..	False	True
SAL	101	fab1-leaf1		N9K-C9396FX	leaf	192.168.0.95/32	False	True

The spine receives an IP address within a few minutes, which is reflected in the APIC GUI.

The screenshot shows the Cisco APIC GUI with the 'Fabric Membership' page. The table below represents the data shown in the interface:

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMMISSIONED	SUPPORTED MODEL
FGE	201	fab1-spine1		N9K-C9508	spine	192.168.246.1/32	False	True
FGE	202	fab1-spine2		N9K-C9508	spine	192.168.0.64/32	False	True
SAL	102	fab1-leaf2		N9K-C9312RFX	leaf	192.168.0.127/32	False	True
SAL	101	fab1-leaf1		N9K-C9396FX	leaf	192.168.0.95/32	False	True

The spine has fully rejoined the ACI fabric and the APIC automatically pushes down all relevant policies. In this case, the chassis of a spine with Node ID 201 failed. The chassis was replaced and its serial number was mapped to the same Node ID. The APIC then pushes all policies relevant to Node 201 to the spine without the need for further configuration.