

Leaf Hardware Failure Recovery

TAC

Document ID: 118865

Contributed by Robert Correiro, Cisco TAC Engineer.

Apr 02, 2015

Contents

Introduction

Problem

Solution

Introduction

This document describes the process to replace a leaf switch, which has failed due to a hardware issue, in Application Centric Infrastructure (ACI) mode.

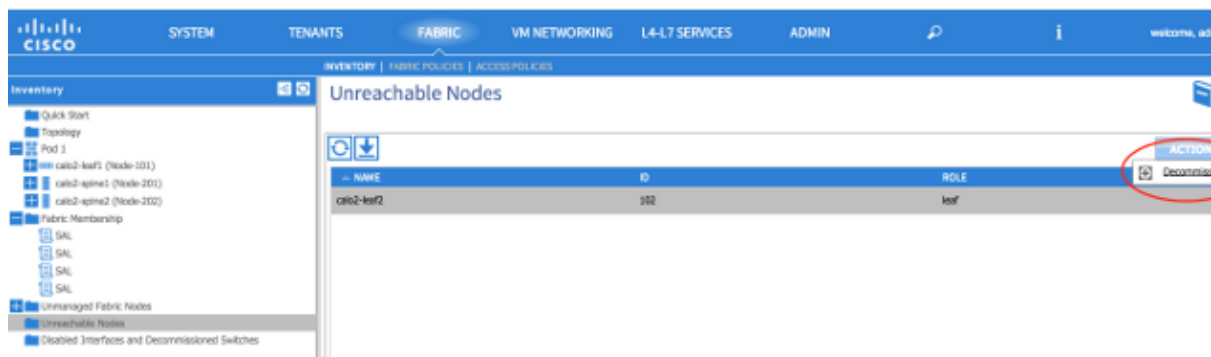
Problem

There is current and working ACI fabric. A leaf has failed, which was previously functional and part of the fabric. This is due to a hardware failure and the leaf needs to be replaced.

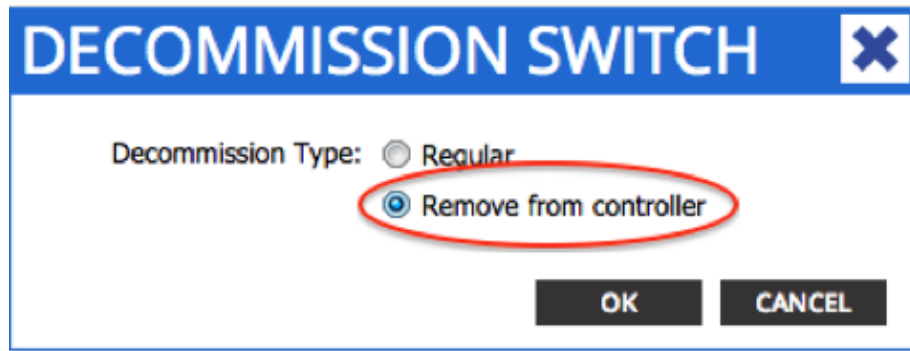
Solution

Complete these steps in order to replace the leaf:

1. If the failed leaf is currently powered on, remove the power cables in order to power it off.
2. From the Application Policy Infrastructure Controller (APIC) GUI, choose **Fabric > Inventory > Unreachable Nodes**. The failed leaf 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 leaf and remove it from the controller.
 - a. Select the leaf in the work pane. From the Actions drop-down list, choose **Decommission**.



- b. 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 leaf from the rack and install the replacement. The new leaf boots into standalone mode that runs NX-OS.
5. Complete one of these three options in order to load an ACI image into the leaf's bootflash.
 - ◆ Option A – Copy the ACI image from a USB drive into bootflash. List the contents of the USB drive (usb1: or usb2:) for the ACI image name in order to complete this:

```
switch# dir usb1:
switch# copy usb1:aci-n9000-dk9.11.0.3f.bin bootflash:
```

- ◆ Option B – Enable Secure Copy (SCP) services in order to push the ACI image from an APIC. Configure the management 0 interface with an IP address, set up a default gateway for the management virtual routing and forwarding (VRF) instance, and enable the SCP server feature.

```
switch# configure terminal
switch(config)# interface mgmt 0
switch(config-if)# ip address ipv4-address { [/length] | [subnet-mask]}
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vrf context management
switch(config-vrf)# ip route 0.0.0.0/0 default-gw-ip
switch(config-vrf)# exit
switch(config)# feature scp-server
switch(config)# exit
switch# copy running-config startup-config
```

Note: Before you proceed, verify that you can ping one of your APICs from the management VRF.

Enter this command on the APIC:

```
admin@apic:~>
scp /firmware/fwrepos/fwrepo/<aci_image.bin> admin@<node-mgmt-ip>:<aci_image.bin>
```

Note: *<aci_image.bin>* is the filename of the ACI switch image located on the APIC. *<node-mgmt-ip>* is the management IP address which was previously configured on the switch.

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

- ◆ Option C – Copy the image from an accessible SCP/FTP/TFTP server. Configure the management 0 interface with an IP address, set up a default gateway for the management VRF instance, and copy the image into bootflash. This example uses SCP.

```
switch# configure terminal
switch(config)# interface mgmt 0
switch(config-if)# ip address ipv4-address{ [/length] | [subnet-mask] }
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vrf context management
switch(config-vrf)# ip route 0.0.0.0/0 default-gw-ip
switch(config-vrf)# end
```

Verify the server can be pinged from the management VRF and then copy the image from the server into bootflash.

```
switch#
copy scp://scpuser@10.0.0.10/path/to/aci-n9000-dk9.11.0.3f.bin bootflash: vrf manage
```

Enter the **dir bootflash:** command in order to verify the ACI image has transferred successfully to the standalone switch.

```
switch# dir bootflash:
 6626   Nov 18 14:22:33 2014 20141118_142200_poap_6132_init.log
500237761 Nov 14 18:24:12 2014 aci-n9000-dk9.11.0.2j.1.0-2j.bin
328541633 Nov 18 14:13:02 2014 auto-s
      2   Nov 18 14:15:24 2014 diag_bootup
      53  Nov 18 14:15:01 2014 disk_log.txt
 4096   Nov 14 19:43:26 2024 lost+found/
 3093   Nov 18 14:10:01 2014 mem_log.txt.old.gz
309991424 Nov 18 14:10:52 2014 n9000-dk9.6.1.2.I2.1.bin
 4096   Nov 08 14:28:49 2014 scripts/
```

6. Convert the Standalone Leaf to ACI mode. If the leaf runs NX-OS version **6.1(2)I3(3)** or later, use Method 1. Otherwise, use Method 2. Method 2 might also be used as a backup if Method 1 fails.

- ◆ Method 1

NXOS software tested: n9000-dk9.6.1.2.I3.3a.bin
ACI software tested: aci-n9000-dk9.11.0.3f.bin

Enter these commands on the standalone node:

```
switch# configure terminal
switch(config)# boot aci bootflash:<aci_image.bin>
switch(config)# reload
```

- ◆ Method 2

NXOS software tested: n9000-dk9.6.1.2.i2.1.bin
ACI software tested: aci-n9000-dk9.11.0.2j.1.0-2j.bin

Enter these commands on the standalone node:

```
switch# configure terminal
switch(config)# no boot nxos
switch(config)# end
switch# copy running-config startup-config
switch# reload
```

This causes the switch to boot into the 'loader>' prompt. Enter the *dir* command in order to list the contents of bootflash for the ACI image name.

```
Loader Version 8.06

loader > dir

bootflash:

  auto-s
  mem_log.txt
  disk_log.txt
  mem_log.txt.old.gz
  lost+found
  .patch
  aci-n9000-dk9.11.0.2j.bin
  .patch-issu
  scripts
  20141121_003542_poap_6330_init.log
  n9000-dk9.6.1.2.i2.2b.bin
  20141121_005455_poap_5924_init.log
```

Boot the ACI image with the *boot<aci_image_name.bin>* command.

The leaf boots into ACI mode in a fabric discovery state. Log in with the username *admin* with no password.

```
User Access Verification
(none) login: Certificate verification passed

User Access Verification
(none) login: admin
*****
Fabric discovery in progress, show commands are not fully functional
Logout and Login after discovery to continue to use show commands.
*****
(none)# █
```

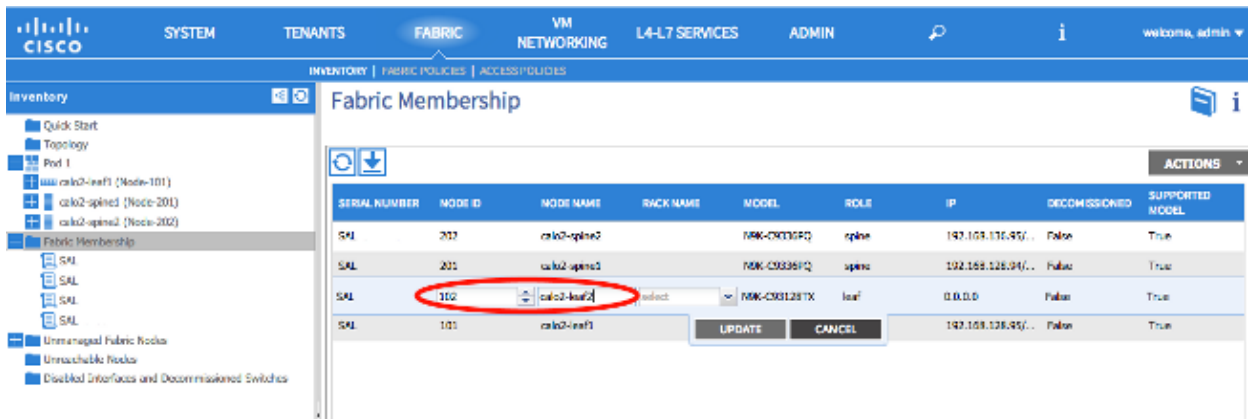
7. Use the image name from the *dir* command in order to set the boot variables on the leaf:

```
(none)# dir bootflash
(none)# setup-bootvars.sh <aci_image_name.bin>
```

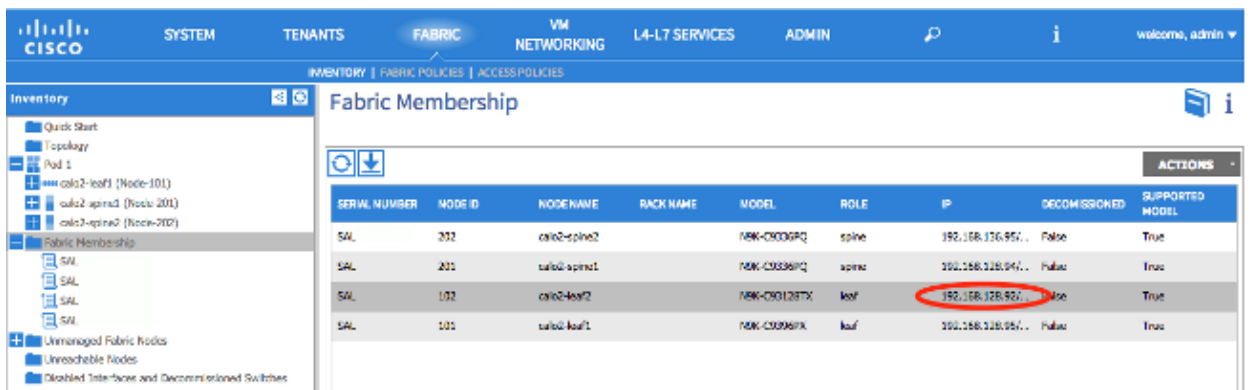
8. Enter these commands in order to verify the boot variables were set properly:

```
leaf# cat /mnt/cfg/0/boot/grub/menu.lst.local
boot aci-n9000-dk9.11.0.3f.bin
leaf# cat /mnt/cfg/1/boot/grub/menu.lst.local
boot aci-n9000-dk9.11.0.3f.bin
```

9. In the APIC GUI, choose *Fabric > Inventory > Fabric Membership*. The new leaf is listed as a row in the work pane. Double-click the row and assign the same 'Node ID' and 'Node Name' from the previous leaf.



The leaf receives an IP address, which is reflected in the APIC GUI.



The new leaf has now fully joined the ACI fabric and the APIC automatically pushes down all relevant policies. In this case, a leaf with Node ID 102 failed. It was replaced by another leaf, which was given the same Node ID. The APIC then pushes all policies relevant to Node 102 down to the new leaf without the need for further configuration.

On the command line the hostname changes in order to reflect what was entered into the APIC GUI, which verifies policies have been pushed down to the new node.