Configure Leaf or Spine Replacement in ACI

Contents

Introduction

This document describes how to replace a leaf or spine switch in the Application Centric Infrastructure (ACI) fabric.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- ACI Fabric
- ACI Application Policy Infrastructure Controller (APIC) GUI
- ACI Leaf and Spine Switch CLI

Components Used

The information in this document is based on these software and hardware versions:

- ACI Leaf Switch N9K-C9372TX-E Model
- ACI Fabric Version 2.x. Some GUI updates have been added representing later releases.

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.

Background Information

Note: The procedure listed here is applicable for any model of the switch and any ACI version that runs on the fabric.

These are the steps to ensure that the switch is in ACI mode.

1. Power on the switch and connect a console.
2. Enter the command **show version** and check to see if the switch is in NxOS mode or ACI mode.
3. If it runs in NxOS mode, refer to Converting from Cisco NX-OS to ACI Boot Mode and from ACI Boot Mode Back to Cisco NX-OS in order to convert the switch to ACI mode.

Note: If you are in the USA, choose the preferred version of ACI software to be preloaded when you place the Return Material Authorization (RMA) request.

Configure

Clean Up the Replacement Switch

Once you confirm the switch is in ACI mode, these are the steps to clean up the replacement switch.

1. From the new switch console, enter the command **setup-clean-config.sh.**
2. Reload (enter the command **reload**) in order to clean up any configurations that already exist on the switch.

This prevents the issue due to some configurations that already exist in the new switch that conflicts with the current fabric, even if the new switch was configured with another ACI fabric previously.

Configuration

Step 1. Decommission/Remove the Current/Failed Switch From the Controller

1. In the ACI GUI, navigate to **Fabric** > **Inventory** > **Fabric Membership** and identify the switch to be replaced. In this example, as shown in the image, leaf 103 is replaced.

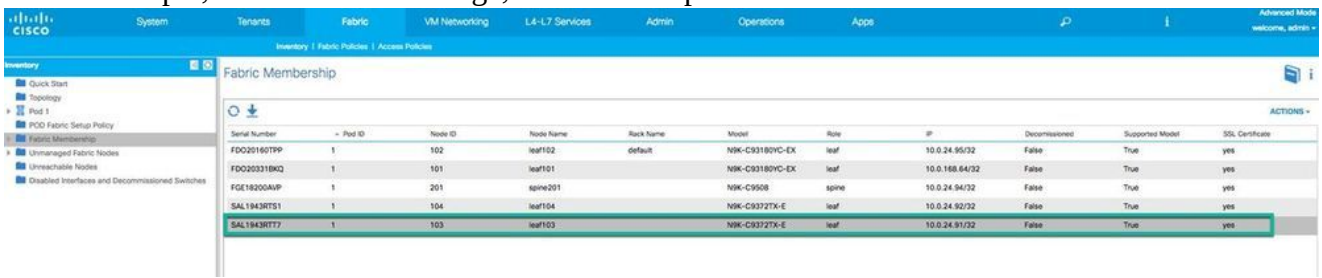

2. Right-click the switch to be replaced and from the drop-down list choose **Decommission Switch.** Now a new pop-up window opens, as shown in the image. Check point 4 to see how the GUI differs in the later release.
3. Select **Remove from Controller** and then click **Submit.**
4. As shown in the image, click **Yes** in order to confirm the decommission process. Now the switch disappears from the Fabric Membership page.

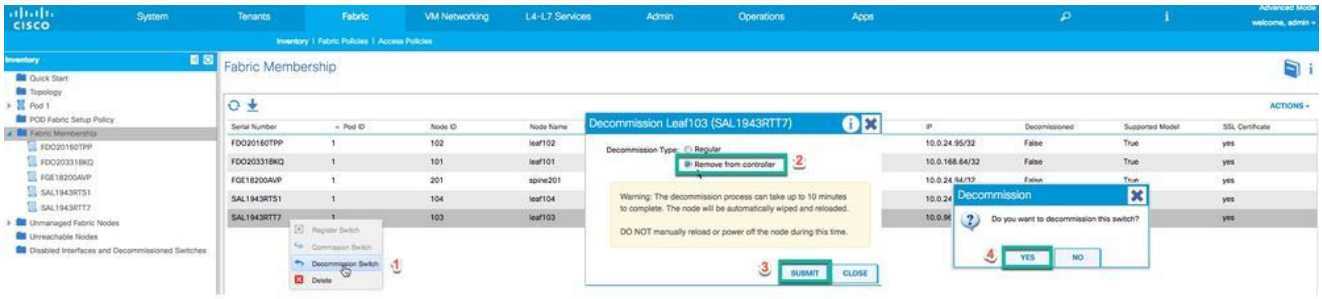

On later releases the GUI option might show up differently. Select **Remove From Controller** for switch replacement on 5.x. On 6.0.x, select **Decomission** and then click**Decomission & Remove** to proceed with the switch removal.

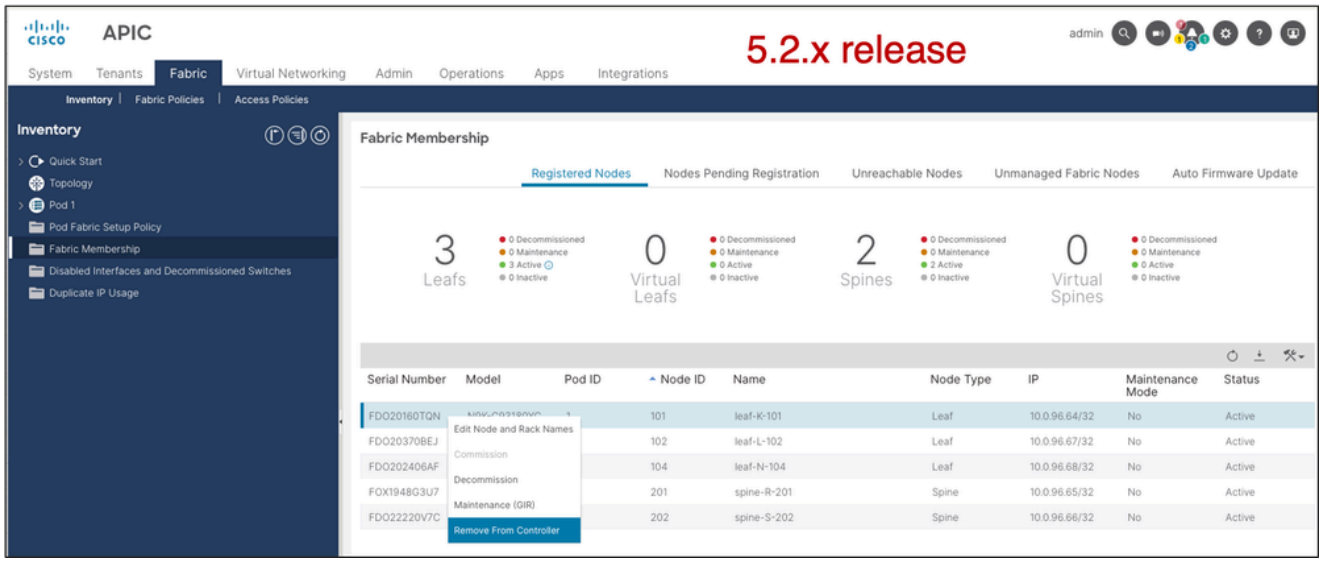

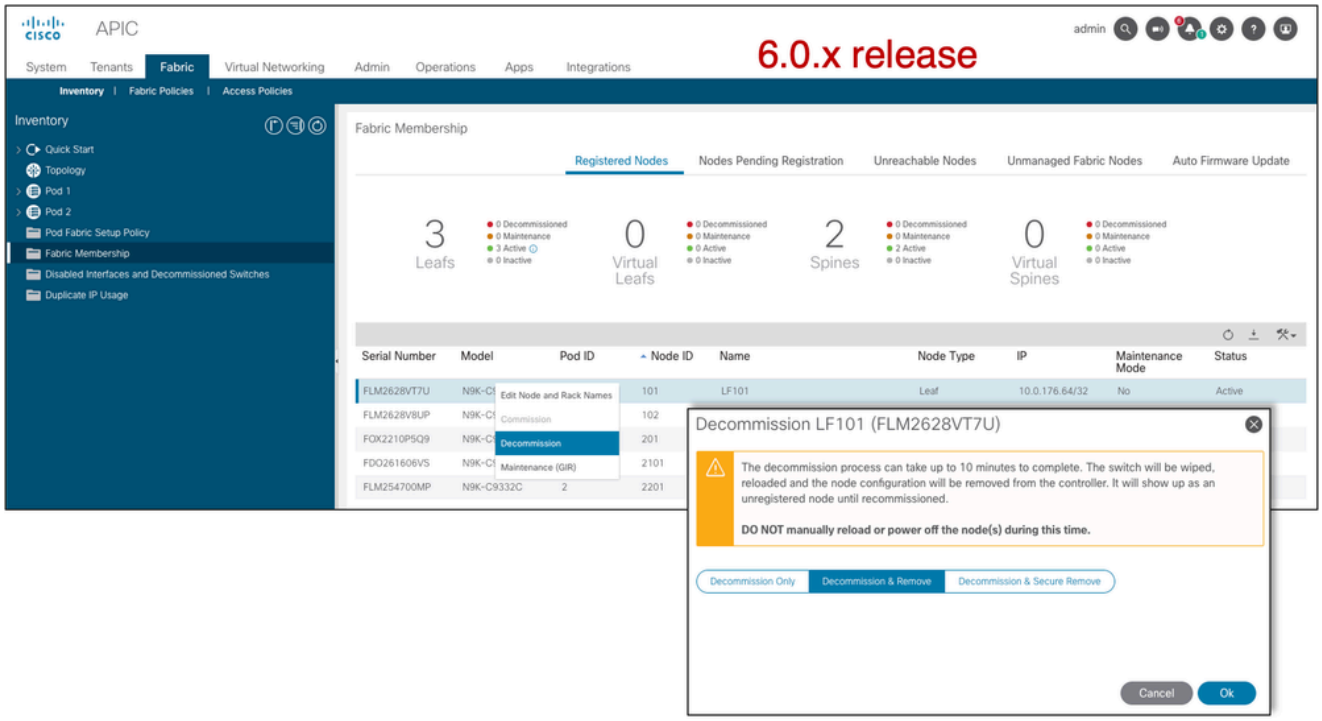

5. Disconnect the switch to be replaced from the fabric and disconnect the power cable.
6. Unmount the old switch and mount the new switch.

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 (in the earlier release) is used in order to temporarily remove the node from the ACI fabric, with the expectation that the same node rejoins the fabric with the same Node ID in the future. For instance, if the node needs to be temporarily powered down for maintenance.

Step 2. Commission the New Switch

Note: Ensure that the new leaf/spine switch is connected to all the spine/leaf switches in the fabric. If you replace a leaf switch, connect only the uplink cables to your spines. Wait for the leaf switch to be active (step 5) in the fabric before you connect the downlink cables.

Note: Before you add the new replacement switch to the fabric you have to upgrade it manually to the target image or an image that has a direct upgrade path to the target image (in case you would like the last upgrade step done by a policy upgrade to make sure the BIOS/FPGA is updated properly). When you add a switch with an image that has multiple upgrade steps to the target image, it causes multiple issues and impacts your production environment.

If the switch is in ACI mode and you have connected it to the fabric, the new switch, once powered on, can get discovered automatically through Link Layer Discovery Protocol (LLDP).

1. Power on the new switch and connect the new switch to the fabric.
2. Navigate back to **GUI > Fabric > Inventory > Fabric Membership** and look for a new switch which does not have an IP address assigned (0.0.0.0) and no node ID assigned, as shown in the image. Cross verify the switch with its serial number.

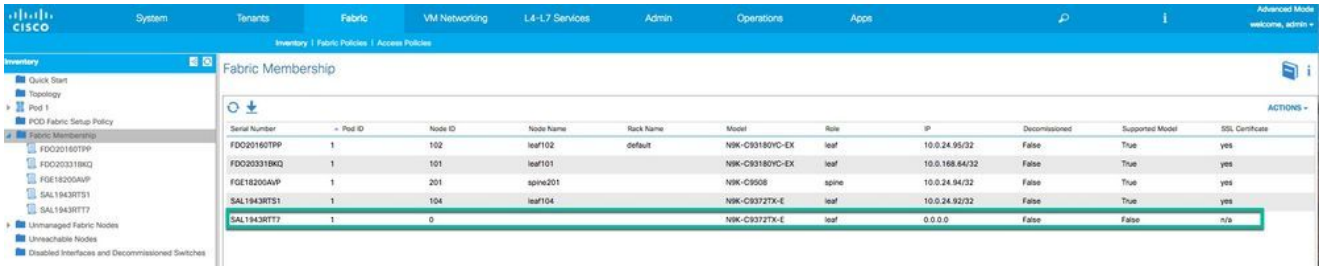

3. As shown in the image, right-click the new switch and from the drop-down list choose **Register Switch.**

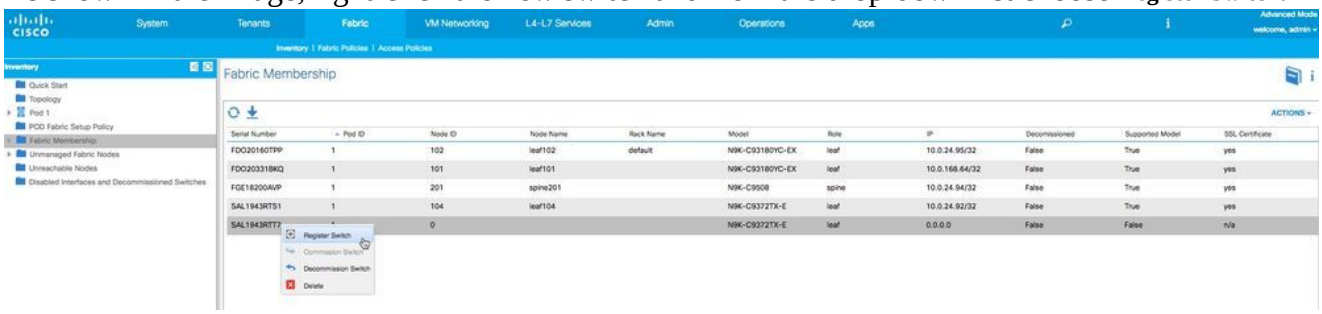

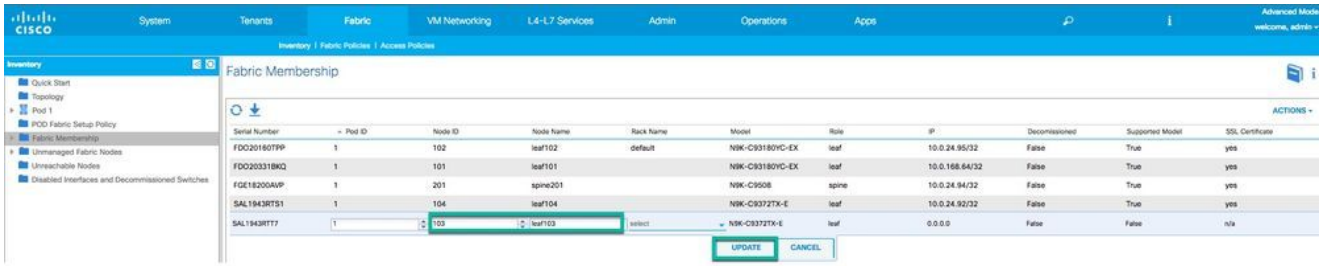

4. The fields, as shown in the image, are to be filled with the required information.

- **POD ID**: Default is 1. If you have a multi-pod fabric use the correct POD ID.
- **Node ID**: It is very important to configure the correct node ID. Enter the same node ID as the previous switch because the APIC pushes the configuration based on the node ID. Once you assign and it gets registered, you cannot change this without decommissioning the switch.
- **Node Name**: Enter the same name for the node as before.

5. As shown in the image, the new leaf gets an IP assigned from the APIC DHCP pool.

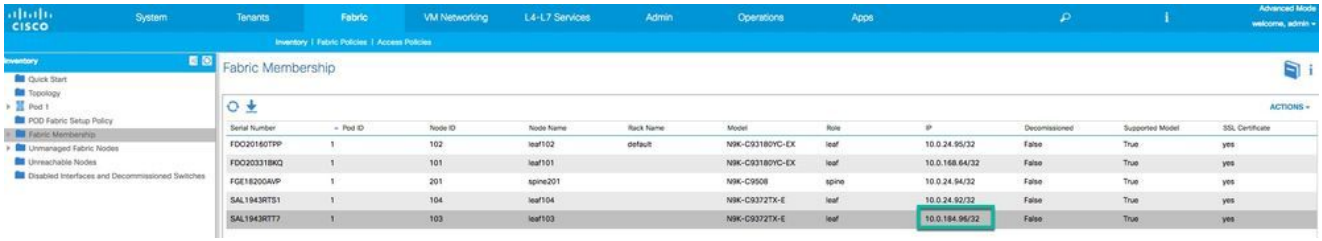

6. If you replace the leaf switch, connect the downlink cables now and confirm all ports are up.

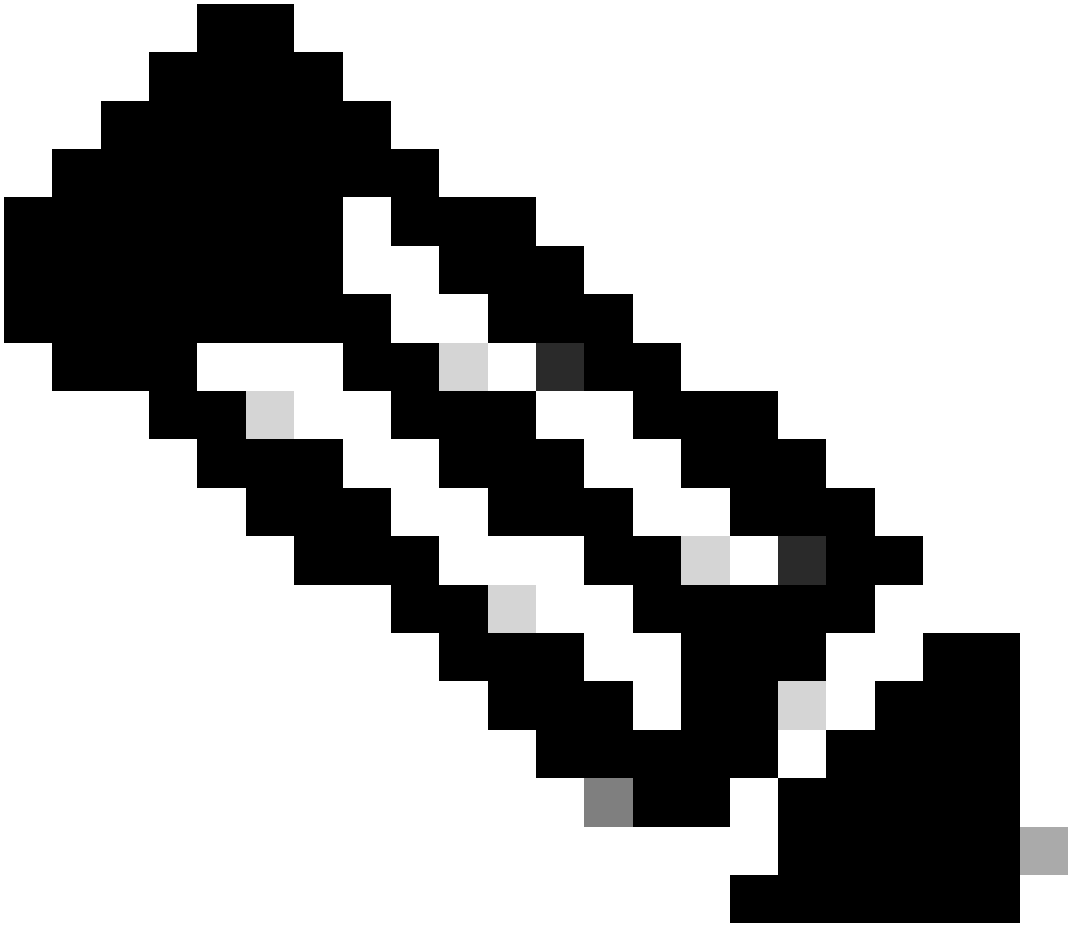

Note: If the decommissioned node has Port Profile deployed on it, an additional reload is necessary in the commissioned node in order to apply the configuration in the ports.

Verify

Use this section in order to confirm that your configuration works properly.

1. You can verify the switch status in **GUI > Fabric > Inventory > Topology .** The new switch is part of the topology, as shown in the image.

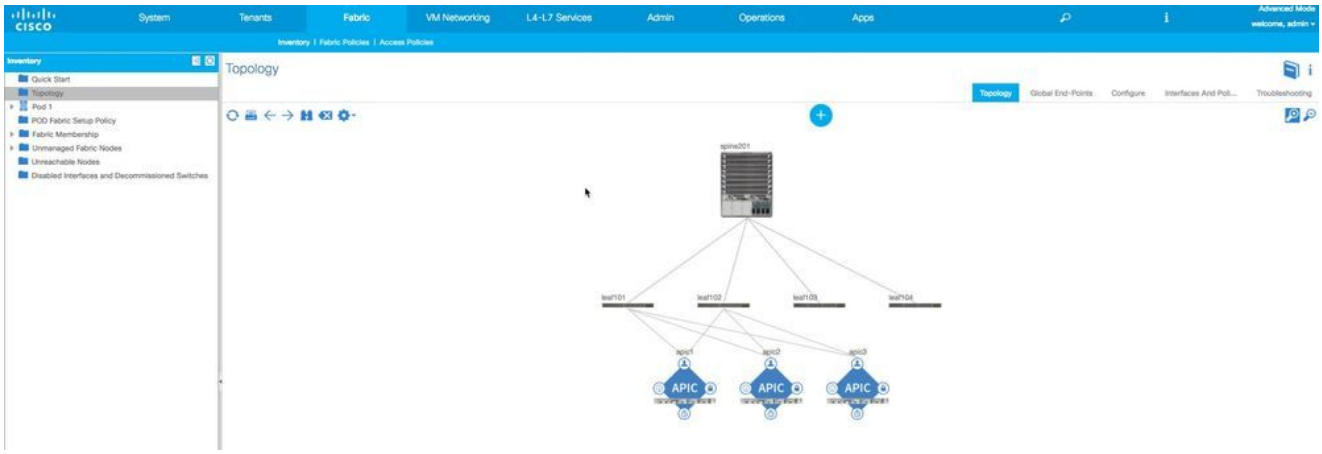

2. Connect to the APIC IP address through SSH and enter the command **acidiag fnvread** in order to confirm the new switch state which shows up as **active**.

```
apic1# acidiag fnvread
      ID   Pod ID              Name    Serial Number         IP Address    Role        State   LastUpdMsgId
-------------------------------------------------------------------------------------------------------------
     101        1           leaf101      FDO20331BKQ     10.0.168.64/32    leaf       active   0
     102        1           leaf102      FDO20160TPP      10.0.24.95/32    leaf       active   0
     103        1           leaf103      SAL1943RTT7     10.0.184.96/32    leaf       active   0
     104        1           leaf104      SAL1943RTS1      10.0.24.92/32    leaf       active   0
     201        1          spine201      FGE18200AVP      10.0.24.94/32   spine       active   0

Total 5 nodes

apic1#
```

Troubleshoot

This section provides information you can use in order to troubleshoot your configuration.

Scenario 1. The New Node is Not Discovered in the Fabric

1. Connect a console and enter the command **show version**.
2. If it is in NxOS mode, convert to ACI mode.
3. Enter the command **show lldp neighbors** and check if it discovers the directly connected switch.
4. If it is not listed, check and confirm the cable is good. Otherwise, open a case with the Technical Assistance Center (TAC) for help.

Note: For the procedure to convert NxOS mode to ACI mode, refer to the Background Information section.

Scenario 2. The Newly Added Switch is Shown as NOT SUPPORTED

1. Navigate to **GUI > Fabric > Inventory > Fabric Membership**.
2. Check whether the new switch is listed as **No** under the **Supported Model** column.
3. If **No**, it could be the issue of your APIC catalogue firmware which is too old. Thus, the model of the new switch is not listed in the catalogue.

In order to solve this, upgrade the APIC to the same code version as the new switch. After which, the new switch can join the fabric.

Scenario 3. SSL Certificate Issue

If the switch fails to get registered with the fabric after you assign a node ID and node name, there could be an SSL certificate issue. In order to verify this, from the console enter the command **netstat -an | grep <TEP ip of APIC>** and check for an **ESTABLISHED** session with APIC on port 12215. This session can be established with any of the APICs in your fabric. In order to verify, enter the command again with different APIC IP addresses.

Example:

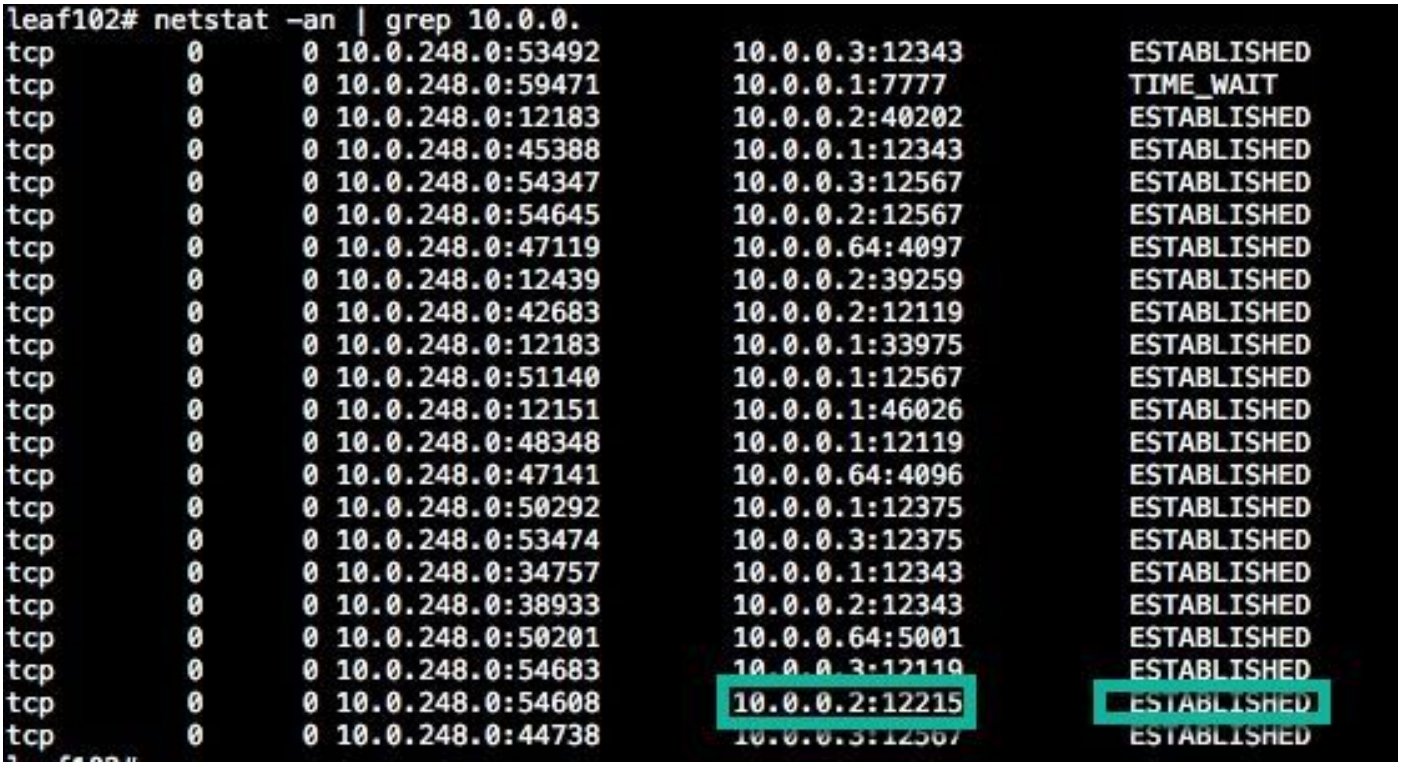

An established session with any of the APICs on port 12215 means that the new switch is able to communicate with the APIC policy manager. If you do not see this session with any of the APICs, it could be an SSL certificate issue. Open a case with TAC for further assistance.

Scenario 4. New Switch Does Not Get a TEP IP Address Assigned

If the new switch does not get a TEP IP address assigned after you register the switch, it can be because of an issue in DHCP IP address allocation from the APIC. Open a case with TAC for assistance.