



Cisco APIC Getting Started Guide, Release 6.1(x)

First Published: 2024-08-01

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at:

<http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and-if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)



CONTENTS

PREFACE	Trademarks iii
----------------	-----------------------

CHAPTER 1	New and Changed Information 1
	New and changed information 1

CHAPTER 2	Initial Setup 3
	For Next Steps, See... 3
	Simplified Approach to Configuring in Cisco APIC 4
	Changing the BIOS Default Password 4
	About the APIC 5
	Setting up the Cisco APIC 5
	Setup for Active and Standby APIC 8
	Bringing up the Cisco APIC Cluster Using the GUI 14
	Provisioning IPv6 Management Addresses on APICs 19
	Accessing the GUI 20
	Accessing the REST API 21
	Accessing the NX-OS Style CLI 21
	Accessing the NX-OS Style CLI from a Terminal 22
	Accessing the NX-OS Style CLI from the GUI 22
	Accessing the Object Model CLI 23

CHAPTER 3	APIC GUI Overview 25
	Overview of the GUI 25
	Menu Bar and Submenu Bar 26
	Menu Bar Tabs 27
	System Tab 27

Tenants Tab	27
Fabric Tab	28
Virtual Networking Tab	28
Admin Tab	28
Operations Tab	28
Apps Tab	29
Integrations Tab	30
Menu Bar Tools	30
Search	30
Launch the Multi-Site Manager	30
Feedback	30
Alerts	30
Tool	31
Help	31
Manage my Profile	32
Navigation Pane	33
Work Pane	33
Common Pages in the Work Pane	34
Personalizing the Interface	34
Naming the APIC GUI	34
Adding a Login Banner to the CLI or GUI	35
Single-Browser Session Management	35
Deployment Warning and Policy Usage Information	35
Graphical Configuration of Ports	36
Viewing an API Interchange in the GUI	37
GUI Icons	39
Fault, Statistics, and Health Level Icons	40
Next generation user interface	40
Preview the next generation user interface	44

CHAPTER 4	Fabric Initialization and Switch Discovery	45
	Initializing the Fabric	45
	About Fabric Initialization	45
	Fabric Topology (Example)	45

Multi-Tier Fabric Topology (Example)	46
Changing the External Routable Subnet	48
Switch Discovery	50
About Switch Discovery with the APIC	50
Switch Registration with the APIC Cluster	50
Switch Role Considerations	51
Registering an Unregistered Switch Using the GUI	51
Adding a Switch Before Discovery Using the GUI	53
Switch Discovery Validation and Switch Management from the APIC	54
Validating the Registered Switches Using the GUI	55
Validating the Fabric Topology	55
Validating the Fabric Topology Using the GUI	55
Unmanaged Switch Connectivity in VM Management	55
Troubleshooting Switch Discovery Issues	55
Finding Your Switch Inventory Using the GUI	57
Troubleshooting Switch Discovery Issues	58
Maintenance Mode	60
Maintenance Mode	60
Removing a Switch to Maintenance Mode Using the GUI	62
Inserting a Switch to Operational Mode Using the GUI	62
Cisco NX-OS to Cisco ACI POAP Auto-conversion	62
About Cisco NX-OS to Cisco ACI POAP Auto-conversion	62
Guidelines and Limitations for Cisco NX-OS to Cisco ACI POAP Auto-conversion	63
Converting a Cisco NX-OS Node to Cisco ACI With POAP Auto-conversion Using the GUI	63
Cisco Nexus 9000 Switch Secure Erase	64
About Cisco Nexus 9000 Switch Secure Erase	64
Securely Erasing User Data from a Cisco Nexus 9000 Switch Using the GUI	65
Securely Erasing User Data from a Module of a Cisco Nexus 9000 Modular Switch Line Card Using the GUI	65
Securely Erasing User Data from a Cisco Nexus 9000 Switch Using the Switch's CLI	66
CHAPTER 5	
Cisco APIC Cluster Management	69
APIC Cluster Overview	69
Expanding the Cisco APIC Cluster	69

Contracting the Cisco APIC Cluster	70
Cluster Management Guidelines	70
Expanding the APIC Cluster Size	71
Reducing the APIC Cluster Size	71
Replacing Cisco APIC Controllers in the Cluster	72
Expanding the APIC Cluster Using the GUI	74
Expanding the APIC Cluster Using the Add Node Option	74
Contracting the APIC Cluster Using the GUI	77
Contracting the APIC Cluster Using the Delete Node Option	77
Commissioning and Decommissioning Cisco APIC Controllers	78
Commissioning a Cisco APIC in the Cluster Using the GUI	78
Commissioning a Cisco APIC in the Cluster	79
Decommissioning a Cisco APIC in the Cluster Using the GUI	80
Decommissioning a Cisco APIC in the Cluster	81
Shutting Down the APICs in a Cluster	82
Shutting Down all the APICs in a Cluster	82
Bringing Back the APICs in a Cluster	83
Cold Standby	83
About Cold Standby for a Cisco APIC Cluster	83
Guidelines and Limitations for Standby Cisco APICs	83
Verifying Cold Standby Status Using the GUI	84
Switching Over an Active APIC with a Standby APIC Using the GUI	84
Migration of APICs	85
Guidelines and Limitations	85
How the migration process works	86
Migrate a Physical APIC Cluster to a Virtual APIC Cluster (or Virtual APIC Cluster to a Physical APIC Cluster)	86
Checking the Migration Status	88
Operations in case of Migration Failure	90
Basic Troubleshooting	91
<hr/>	
APPENDIX A	Configuring the Cisco APIC Using the CLI 93
	Cluster Management Guidelines 93
	Replacing a Cisco APIC in a Cluster Using the CLI 94

Reducing the APIC Cluster Size	95
Contracting the Cisco APIC Cluster	96
Switching Over Active APIC with Standby APIC Using CLI	97
Verifying Cold Standby Status Using the CLI	97
Registering an Unregistered Switch Using the CLI	98
Adding a Switch Before Discovery Using the CLI	98
Removing a Switch to Maintenance Mode Using the CLI	98
Inserting a Switch to Operation Mode Using the CLI	99
Configuring a Remote Location Using the NX-OS Style CLI	99
Finding Your Switch Inventory Using the NX-OS CLI	100
Verifying the Cisco APIC Cluster Using the CLI	102

APPENDIX B

Configuring the Cisco APIC Using the REST API	105
Expanding the APIC Cluster Using the REST API	105
Contracting the APIC Cluster Using the REST API	105
Reducing the APIC Cluster Size	107
Switching Over Active APIC with Standby APIC Using REST API	108
Registering an Unregistered Switch Using the REST API	109
Adding a Switch Before Discovery Using the REST API	109
Removing a Switch to Maintenance Mode Using the REST API	110
Inserting a Switch to Operational Mode Using the REST API	110
Configuring a Remote Location Using the REST API	111
Sending an On-Demand Tech Support File Using the REST API	111
Finding Your Switch Inventory Using the REST API	112



CHAPTER

1

New and Changed Information

This chapter contains the following sections:

- [New and changed information, on page 1](#)

New and changed information

The following tables provide an overview of the significant changes to this guide for this release. The tables do not provide an exhaustive list of all changes made to the guide or of the new features up to this release.

Table 1: New features and changed information for Cisco APIC release 6.1(1)

Feature	Description	Where documented
Preview of the next generation user interface	This release introduces a preview of the next generation Cisco APIC user interface. With this preview, you can get an idea of upcoming development of the GUI.	Preview the next generation user interface, on page 44
Migrating physical APICs to virtual APICs and virtual APICs to physical APICs	You can migrate Cisco APICs from a physical APIC cluster to a virtual APIC cluster deployed on an ESXi host (using VMware vCenter), or from a virtual APIC cluster (on an ESXi host) to a physical APIC cluster.	Migration of APICs



CHAPTER 2

Initial Setup

This chapter contains the following sections:

- [For Next Steps, See..., on page 3](#)
- [Simplified Approach to Configuring in Cisco APIC, on page 4](#)
- [Changing the BIOS Default Password, on page 4](#)
- [About the APIC, on page 5](#)
- [Setting up the Cisco APIC , on page 5](#)
- [Accessing the GUI, on page 20](#)
- [Accessing the REST API, on page 21](#)
- [Accessing the NX-OS Style CLI, on page 21](#)
- [Accessing the Object Model CLI, on page 23](#)

For Next Steps, See...

This table provides a list of additional documents that are useful references along with the *Cisco APIC Getting Started Guide*. These Cisco APIC documents and others are available at the [APIC documents landing page](#).



Tip To find documentation for a specific Cisco APIC feature, type the feature name in the **Choose a Topic** box in the [APIC documents landing page](#).

Documents
Application Centric Infrastructure Fabric Hardware Installation Guide
Cisco APIC Installation, Upgrade, and Downgrade Guide
Cisco APIC Basic Configuration Guide
Cisco APIC Layer 2 Networking Configuration Guide
Cisco APIC Layer 3 Networking Configuration Guide
Cisco APIC Security Configuration Guide
Cisco APIC System Management Configuration Guide

Documents
Cisco ACI Virtualization Guide
Cisco Application Centric Infrastructure Fundamentals
Cisco APIC Layer 4 to Layer 7 Services Deployment Guide

Most of these links take you to the section of the documentation landing page that contains the specified document. Click the arrow at the right end of the section title to expand the document list for that section, then find the document for your release.

If the document for a release does not exist, the document for the previous release applies. For example, the *Cisco APIC System Management Configuration Guide* was not republished for the 5.0 releases because there are no changes from the 4.2 releases. Therefore, you should use the document for the 4.2 releases.

Simplified Approach to Configuring in Cisco APIC

Cisco APIC supports a simplified approach to configuring the ACI with an additional NX-OS style CLI interface. The existing methods of configuration using REST API and the GUI are supported as well.

In addition to the simple approach available for network administrators and other users of the NX-OS style CLI, there is intelligence embedded in this approach as compared to the GUI or the REST API. In several instances, the NX-OS style CLI can create the ACI model constructs implicitly for a user's ease of use, and they also provide validations to ensure consistency in configuration. This functionality reduces and prevents faults.

For further details about configurations and tasks, see the *Cisco APIC Basic Configuration Guide* and the *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*.

Changing the BIOS Default Password

Cisco Application Policy Infrastructure Controller (APIC) ships with a default BIOS password. The default password is "password". When the boot process starts, the boot screen displays the BIOS information on the console server.



Note The 6.0(2) and later releases support the APIC-L4 and APIC-M4 servers. These servers have a default password of "password" or "Insieme123".

To change the default BIOS password perform the following task:

-
- Step 1** During the BIOS boot process, when the screen displays **Press <F2> Setup**, press **F2**. The **Entering Setup** message displays as it accesses the setup menu.
- Step 2** At the **Enter Password** dialog box, enter the current password.

Note The default is "password".

The 6.0(2) and later releases support the APIC-L4 and APIC-M4 servers. These servers have a default password of "password" or "Insieme123".

Step 3 In the **Setup Utility**, choose the **Security** tab, and choose **Set Administrator Password**.

Step 4 In the **Enter Current Password** dialog box, enter the current password.

Step 5 In the **Create New Password** dialog box, enter the new password.

Step 6 In the **Confirm New Password** dialog box, re-enter the new password.

Step 7 Choose the **Save & Exit** tab.

Step 8 In the **Save & Exit Setup** dialog box, choose **Yes**.

Step 9 Wait for the reboot process to complete.
The updated BIOS password is effective.

About the APIC

The Cisco Application Centric Infrastructure (ACI) is a distributed, scalable, multitenant infrastructure with external end-point connectivity controlled and grouped through application-centric policies. The Application Policy Infrastructure Controller (APIC) is the unified point of automation, management, monitoring, and programmability for the ACI. The APIC supports the deployment, management, and monitoring of any application anywhere, with a unified operations model for the physical and virtual components of the infrastructure. The APIC programmatically automates network provisioning and control that is based on the application requirements and policies. It is the central control engine for the broader cloud network; it simplifies management and allows flexibility in how application networks are defined and automated. It also provides northbound Representational State Transfer (REST) APIs. The APIC is a distributed system that is implemented as a cluster of many controller instances.

Setting up the Cisco APIC

This section describes how to establish a local serial connection to the Cisco APIC server to begin the initial basic configuration. For additional connection information, including instructions on connecting to the server remotely for setup, refer to "Initial Server Setup" in the *Cisco APIC M3/L3 Server Installation and Service Guide*.

Initial Connection

The Cisco APIC M3/L3 Server operates on a Cisco Integrated Management Controller (CIMC) platform. You can make an initial connection to the CIMC platform using one of these methods:

- Use a KVM cable (Cisco PID N20-BKVM) to connect a keyboard and monitor to the KVM connector on the front panel of the server.
- Connect a USB keyboard and VGA monitor to the corresponding connectors on the rear panel of the server.



Note You cannot use the front panel VGA and the rear panel VGA at the same time.

You can make a serial connection using one of the following methods. Two of these methods require a configuration change in the CIMC:



Note You cannot use more than one of these methods simultaneously.

- Use the DB9 connector of the KVM cable
- Use the rear panel RJ-45 console port (after enabling in the CIMC)
- Connect by Serial-over-LAN (SoL) (after enabling in the CIMC)

The default connection settings from the factory are:

- The serial port baud rate is 115200
- The RJ-45 console port located on the rear panel is disabled in the CIMC
- SoL is disabled in the CIMC

The following are additional notes about serial access:

- If you are using a Cisco Integrated Management Controller (CIMC) for your setup, setup the CIMC first, and then access the Cisco APIC through the CIMC KVM or continue to access the Cisco APIC locally through the rear panel USB/VGA port. If you choose the CIMC KVM access, you will have remote access available later which is required during operations.
- If you are using the RJ-45 console port, connect to CIMC using SSH and enable the SoL port using the following commands:

```
scope sol
  set enabled yes
  set baud-rate 115200
commit
exit
```

After enabling SoL, enter the command **connect host** to access the APIC console.



Note When using SoL, physically disconnect the rear panel RJ-45 console port.

Initial Cisco APIC Setup

When the Cisco Application Policy Infrastructure Controller (Cisco APIC) is launched for the first time, the Cisco APIC console presents a series of initial setup options. For many options, you can press **Enter** to choose the default setting that is displayed in brackets. At any point in the setup dialog, you can restart the dialog from the beginning by pressing **Ctrl-C**.

Important Notes

- If the UNIX user ID is not explicitly specified in the response from the remote authentication server, then some Cisco APIC software releases assign a default ID of 23999 to all users. If the response from the remote authentication server fails to specify a UNIX ID, all users will share the same ID of 23999 and this can result in the users being granted higher or lower privileges than the configured privileges through the RBAC policies on the Cisco APIC.
- Cisco recommends that you assign unique UNIX user IDs in the range of 16000 to 23999 for the AV Pairs that are assigned to the users when in Bash shell (using SSH, Telnet, or Serial/KVM consoles). If a situation arises where the Cisco AV Pair does not provide a UNIX user ID, the user is assigned a user ID of 23999 or similar number from the range that also enables the user's home directories, files, and processes accessible to the remote users with a UNIX ID of 23999.

To ensure that your remote authentication server does not explicitly assign a UNIX ID in its **cisco-av-pair** response, open an SSH session to the Cisco APIC and log in as an administrator (using a remote user account). Once logged in, run the following commands (replace **userid** with the username that you logged in with):

- **admin@apic1: remoteuser-userid> cd /mit/uni/userext/remoteuser-userid**
- **admin@apic1: remoteuser-userid> cat summary**

- Cisco recommends against modifying any parameters using CIMC. If there are any issues, ensure that the default setting for CIMC management node is **Dedicated Mode** and not **Shared**. If **Dedicated Mode** is not used, it can prevent the discovery of fabric nodes.
- Do not upgrade software or firmware using the CIMC user interface, XML, or SSH interfaces unless the modified property and software or firmware version are supported with your specific Cisco APIC version.
- Set the NIC mode to **Dedicated**, when setting up the CIMC, in the CIMC Configuration Utility. After the CIMC is configured, in the CIMC GUI, verify that you have the following parameters set.

Parameters	Settings
LLDP	Disabled on the VIC
TPM Support	Enabled on the BIOS
TPM Enabled Status	Enabled
TPM Ownership	Owned

- Beginning with Release 5.0(2), if you log in to your Cisco APIC using https, and then attempt to log in to the same Cisco APIC using http in the same browser window without first logging out of the Cisco APIC in the https window, you might see the following error message:

Need a valid webtoken cookie (named APIC-Cookie) or a signed request with signature in the cookie.

If this occurs, resolve the issue using either of the following methods:

- Log out of the Cisco APIC in the https window, or
- Delete the cookies in the browser window

You should be able to successfully log into the Cisco APIC using http after resolving the issue with either of the methods above.

- During the initial setup, the system will prompt you to select IPv4, or IPv6, or dual stack configuration. Choosing dual stack will enable accessing the Cisco APIC and Cisco Application Centric Infrastructure (Cisco ACI) fabric out-of-band management interfaces with either IPv4 or IPv6 addresses. While the examples in the table below use IPv4 addresses, you can use whatever IP address configuration options you chose to enable during the initial setup.
- A minimum subnet mask of /19 is recommended.
- Connecting the Cisco APIC to the Cisco ACI fabric requires a 10G interface on the ACI-mode leaf switch. You cannot connect the Cisco APIC directly to the Cisco Nexus 9332PQ, Cisco Nexus 93180LC, or Cisco Nexus 9336C-FX2 ACI-mode leaf switches unless you use a 40G to 10G converter (part number CVR-QSFP-SFP10G), in which case the port on the leaf switches will auto-negotiate to 10G without requiring any manual configuration.



Note Starting with Cisco APIC release 2.2(1n), the Cisco Nexus 93180LC leaf switch is supported.

- The fabric ID is set during the Cisco APIC setup and it cannot be changed unless you perform a clean reload of the fabric. To change the fabric ID, export the Cisco APIC configuration, change the sam.config file, and perform a clean reload of the Cisco APIC and leaf switches. Remove the "fvFabricExtConnP" setting from the exported configuration before importing the configuration into the Cisco APIC after the Cisco APIC comes up. All Cisco APICs in a cluster must have the same fabric ID.
- All logging is enabled by default.
- For login and cluster operations, non-default HTTPS port (default is 443) is not supported for layer 3 physical and layer 3 virtual APICs (on ESXi and AWS). Virtual APICs on ESXi/ AWS are supported from release 6.0(2).

About Cold Standby for a Cisco APIC Cluster

The Cold Standby functionality for a Cisco APIC cluster enables you to operate the Cisco APICs in a cluster in an active/standby mode. In a Cisco APIC cluster, the designated active Cisco APICs share the load and the designated standby Cisco APICs can act as a replacement for any of the Cisco APICs in an active cluster.

An admin user can set up the Cold Standby functionality when the Cisco APIC is launched for the first time. We recommend that you have at least 3 active Cisco APICs in a cluster, and one or more standby Cisco APICs. An admin user must initiate the switch over to replace an active Cisco APIC with a standby Cisco APIC. See the *Cisco APIC Management, Installation, Upgrade, and Downgrade Guide* for more information.

Setup for Active and Standby APIC

Beginning with Cisco Application Policy Infrastructure Controller (APIC) release 6.0(2), for the initial set up and cluster bringup, use the GUI. For more information, see the [Bringing up the Cisco APIC Cluster Using the GUI , on page 14](#) procedure.

Table 2: Setup for Active APIC

Name	Description	Default Value
Fabric name	Fabric domain name	ACI Fabric1
Fabric ID	Fabric ID	1
Number of active controllers	Cluster size	3 Note When setting up a Cisco APIC in an active-standby mode, you must have at least 3 active Cisco APICs in a cluster.
POD ID	POD ID	1
Standby controller	Setup standby controller	NO
Controller ID	Unique ID number for the active Cisco APIC instance.	Valid range: 1-32
Standalone APIC Cluster	Is the Cisco APIC cluster not directly connected to the Fabric, but connected by a layer 3 inter-pod network (IPN). This feature is available only on Cisco APIC release 5.2(1) and later.	NO See the knowledge base article <i>Deploying APIC Cluster Connectivity to the Fabric Over a Layer 3 Network</i> for additional setup instructions.
Controller name	Active controller name	apic1
IP address pool for tunnel endpoint addresses	Tunnel endpoint address pool	10.0.0.0/16 This value is for the infrastructure virtual routing and forwarding (VRF) only. This subnet should not overlap with any other routed subnets in your network. If this subnet does overlap with another subnet, change this subnet to a different /16 subnet. The minimum supported subnet for a 3 Cisco APIC cluster is /23. If you are using Release 2.0(1) the minimum is /22. The 172.17.0.0/16 subnet is not supported for the infra TEP pool due to a conflict of address space with the docker0 interface. If you must use the 172.17.0.0/16 subnet for the infra TEP pool, you must manually configure the docker0 IP address to be in a different address space in each Cisco APIC before you attempt to put the Cisco APICs in a cluster.

Name	Description	Default Value
VLAN ID for infrastructure network ¹	<p>Infrastructure VLAN for Cisco APIC-to-switch communication including virtual switches</p> <p>Note Reserve this VLAN for Cisco APIC use only. The infrastructure VLAN ID must not be used elsewhere in your environment and must not overlap with any other reserved VLANs on other platforms.</p>	
IP address pool for bridge domain multicast address (GIPo)	<p>IP addresses used for fabric multicast.</p> <p>For Cisco APIC in a Cisco ACI Multi-Site topology, this GIPo address can be the same across sites.</p>	<p>225.0.0.0/15</p> <p>Valid range: 225.0.0.0/15 to 231.254.0.0/15, prefixlen must be 15 (128k IPs)</p>
IPv4/IPv6 addresses for the out-of-band management	<p>IP address that you use to access the Cisco APIC through the GUI, CLI, or API.</p> <p>This address must be a reserved address from the VRF of a customer</p>	—
IPv4/IPv6 addresses of the default gateway	<p>Gateway address for communication to external networks using out-of-band management</p>	—
Management interface speed/duplex mode	<p>Interface speed and duplex mode for the out-of-band management interface</p>	<p>auto</p> <p>Valid values are as follows</p> <ul style="list-style-type: none"> • auto • 10baseT/Half • 10baseT/Full • 100baseT/Half • 100baseT/Full • 1000baseT/Full

Name	Description	Default Value
Strong password check	Check for a strong password	[Y]
Password	Password of the system administrator This password must be at least 8 characters with one special character.	—

¹ To change the VLAN ID after the initial APIC setup, export your configurations, rebuild the fabric with a new infrastructure VLAN ID and import the configurations so that the fabric does not revert to the old infrastructure VLAN ID. See the KB article about *Using Export and Import to Recover Configuration State*.

Table 3: Setup for Standby APIC

Name	Description	Default Value
Fabric name	Fabric domain name	ACI Fabric1
Fabric ID	Fabric ID	1
Number of active controllers	Cluster size	3 Note When setting up Cisco APIC in an active-standby mode, you must have at least 3 active Cisco APICs in a cluster.
POD ID	ID of the POD	1
Standby controller	Setup standby controller	Yes
Standby Controller ID	Unique ID number for the standby Cisco APIC instance	Recommended range: >20
Controller name	Standby controller name	NA
IP address pool for tunnel endpoint addresses	Tunnel endpoint address pool	10.0.0.0/16 This value is for the infrastructure virtual routing and forwarding (VRF) only. This subnet should not overlap with any other routed subnets in your network. If this subnet does overlap with another subnet, change this subnet to a different /16 subnet. The minimum supported subnet for a 3 Cisco APIC cluster is /23. If you are using Release 2.0(1) the minimum is /22.

Name	Description	Default Value
VLAN ID for infrastructure network ²	<p>Infrastructure VLAN for Cisco APIC-to-switch communication including virtual switches</p> <p>Note Reserve this VLAN for Cisco APIC use only. The infrastructure VLAN ID must not be used elsewhere in your environment and must not overlap with any other reserved VLANs on other platforms.</p>	
IPv4/IPv6 addresses for the out-of-band management	<p>IP address that you use to access the Cisco APIC through the GUI, CLI, or API.</p> <p>This address must be a reserved address from the VRF of a customer</p>	—
IPv4/IPv6 addresses of the default gateway	Gateway address for communication to external networks using out-of-band management	—
Management interface speed/duplex mode	Interface speed and duplex mode for the out-of-band management interface	<p>auto</p> <p>Valid values are as follows</p> <ul style="list-style-type: none"> • auto • 10baseT/Half • 10baseT/Full • 100baseT/Half • 100baseT/Full • 1000baseT/Full
Strong password check	Check for a strong password	[Y]
Password	<p>Password of the system administrator</p> <p>This password must be at least 8 characters with one special character.</p>	—

- ² To change the VLAN ID after the initial APIC setup, export your configurations, rebuild the fabric with a new infrastructure VLAN ID and import the configurations so that the fabric does not revert to the old infrastructure VLAN ID. See the KB article about *Using Export and Import to Recover Configuration State*.

Example

The following example output shows the initial setup dialog as displayed on the console.



Note Instead of using the **APIC Cluster Bringup** GUI, you can bootstrap and bringup the cluster using REST APIs. For more information, see the *Cisco APIC REST API Configuration Guide*.

Beginning with Cisco APIC release 6.0(2), questions in the example output are not included. For bootstrapping, and bringing up the Cisco APIC cluster, use the GUI. For details, see the [Bringing up the Cisco APIC Cluster Using the GUI](#), on page 14 procedure.

```
Cluster configuration ...
Enter the fabric name [ACI Fabric1]:
Enter the fabric ID (1-128) [1]:
Enter the number of active controllers in the fabric (1-9) [3]:
Enter the POD ID (1-9) [1]:
Is this a standby controller? [NO]:
Enter the controller ID (1-3) [1]:
Enter the controller name [apic1]: apic-1
Enter address pool for TEP addresses [10.0.0.0/16]:
Note: The infra VLAN ID should not be used elsewhere in your environment
      and should not overlap with any other reserved VLANs on other platforms.
Enter the VLAN ID for infra network (2-4094): 3914
Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:

Out-of-band management configuration ...
Enable IPv6 for Out of Band Mgmt Interface? [N]:
Enter the IPv4 address [192.168.10.1/24]: 172.31.1.2/24
Enter the IPv4 address of the default gateway [None]: 172.31.1.1
Enter the interface speed/duplex mode [auto]:

admin user configuration ...
Enable strong passwords? [Y]:
Enter the password for admin:

Reenter the password for admin:

Cluster configuration ...
Fabric name: ACI Fabric1
Fabric ID: 1
Number of controllers: 3
Controller name: apic-1
POD ID: 1
Controller ID: 1
TEP address pool: 10.0.0.0/16
Infra VLAN ID: 3914
Multicast address pool: 225.0.0.0/15

Out-of-band management configuration ...
Management IP address: 172.31.1.2/24
Default gateway: 172.31.1.1
Interface speed/duplex mode: auto

admin user configuration ...
```

```
Strong Passwords: Y
User name: admin
Password: *****
```

The above configuration will be applied ...

Warning: TEP address pool, Infra VLAN ID and Multicast address pool cannot be changed later, these are permanent until the fabric is wiped.

Would you like to edit the configuration? (y/n) [n]:

Bringing up the Cisco APIC Cluster Using the GUI

Beginning with Cisco APIC release 6.0(2), the initial cluster set up and bootstrapping procedure has been simplified with the addition of GUI screen(s) for cluster bring up. The **APIC Cluster Bringup** GUI supports virtual and physical APIC platforms. The virtual APICs (deployed using ESXi or AWS), and physical APICs can be connected to the ACI fabric directly to the leaf switches or remotely attached through a Layer 3 network. The GUI supports both the scenarios. A major advantage of using the **APIC Cluster Bringup** GUI is that, you do not need to enter the parameters for every APIC in a cluster. One APIC can relay the information to the other APICs of the cluster.

Alternatively, you can perform the initial setup and cluster bringup using the REST APIs. See the *Getting Started* section of the [APIC REST API Configuration Procedures](#) guide.

Before you begin

- For virtual APIC on ESXi, ensure to complete the deployment of the Cisco APIC VM using the OVF template on the VMware vCenter GUI. For a three-node cluster, configure three VMs with the management IP address, gateway, and admin passwords. The number of VMs is dependent on the size of the Cisco APIC cluster.
- For virtual APIC on AWS, ensure to complete the deployment of the Cisco APIC VM using the cloud formation template (CFT) on the AWS GUI. AWS allocates IP addresses dynamically from the out-of-band (OOB)/infra/inband subnets accordingly, to correspond with the network adapters of the virtual APIC's EC2 instance.
- For virtual APICs (deployed using AWS/ ESXi), ensure that the admin password(s) are the same for all the Cisco APICs in a cluster.
- For the physical APIC cluster, configure the OOB address for APIC 1. Ensure that the CIMC addresses of APICs 2 to N (where N is the cluster size) are reachable via the OOB address of APIC 1.
- Connectivity between out-of-band and the CIMC is mandatory.

Limitations:

- No support for IPv6 addresses on virtual APICs deployed using AWS.
- For login and cluster operations, non-default HTTPS port (default is 443) is not supported for remotely-attached Cisco APICs (physical and virtual).

-
- Step 1** Log in to the APIC 1 using `https://APIC1-IP`.
- a) For virtual APICs:

If you have completed the deployment of virtual APICs using ESXi (OVF template) or remote AWS (CFT), then you see output on the VM console similar to the following example:

```
System pre-configured successfully.  
Use: https://172.31.1.2 to complete the bootstrapping.
```

The IP address to access the bootstrapping GUI (**APIC Cluster Bringup**) is explicitly indicated, as shown in the example. You can proceed to step 2.

After deploying Cisco APIC on AWS, keep the OOB management IP address handy to access the **Cluster Bringup** GUI. You can get the OOB management IP address from the **Stacks Outputs** tab on the AWS GUI.

b) For physical APICs:

Log in to the APIC 1 KVM console using the CIMC; you will see a screen as shown below:

```
APIC Version: 6.0(2a)  
Welcome to Cisco APIC Setup Utility  
Press Enter Or Input JSON string to bootstrap your APIC node.
```

If you see only a black screen on the KVM, connect to the CIMC using SSH and use serial over LAN (SoL) ("connect host") to connect to the console.

On APIC 1, press **Enter** and provide the requested information. The IP address to access the bootstrapping GUI (**APIC Cluster Bringup**) is explicitly indicated.

```
admin user configuration ...  
  Enter the password for admin [None]:  
  Reenter the password for admin [None]:  
Out-of-band management configuration ...  
  Enter the IP Address [192.168.10.1/24]: 172.20.7.79/23  
  Enter the IP Address of default gateway [192.168.10.254]: 172.20.6.1  
Would you like to edit the configuration? (y/n) [n]:  
System pre-configured successfully.  
Use: https://172.20.7.79 to complete the bootstrapping
```

The IP addresses displayed above are examples. The IP addresses will vary based on your deployment.

Step 2 Using the OOB address, log in to the **APIC Cluster Bringup** GUI. The GUI screen has four parts. Enter the details in the following screens:

- Connection Type
- Cluster Details
- Controller Registration
- Summary

Each of the above screens are discussed in detail in the subsequent steps. The screens are marked as steps with sequential numbers: 1, 2, 3, and 4; after you have entered and saved the required details in each of these screens, the number is replaced with a tick-mark.

Step 3 The first step is entering the **Connection Type** information. In the **Connection Type** screen, choose the type of connection between the APIC and the fabric.

The options are:

- Directly connected to leaf switches (ACI fabric)
- Remotely attached through a Layer 3 network

If it is virtual APIC using AWS, the system detects that the APIC is remotely-attached through a Layer 3 network and proceeds directly to the **Cluster Details** screen.

Step 4 Click **Next**.

Step 5 The second step is entering the **Cluster Details**. Enter the fabric-level details in the **Cluster Details** screen.

- **Fabric Name:** Enter a name for the fabric.
- **Cluster Size:** The default cluster size displayed is "3", which is the recommended minimum cluster size. You can modify this value, based on your cluster size. The supported values are 1, 3, 4, 5, 6, 7, 8, and 9.
- **GiPo Pool:** Enter the IP address used for fabric multicast. The default address is 225.0.0.0/15. The range is from 225.0.0.0/15 to 231.254.0.0/15. The prefixlen must be 15 (128k IP addresses).
You cannot change this value after you have completed the configuration. Having to modify this value requires a wipe of the fabric.
- **Pod ID:** (applicable only for directly connected APICs (virtual and physical)) the pod ID is displayed. If this is your first APIC, "1" is auto-populated. Subsequent APICs of the cluster can be associated with any pod number.
For a remotely-attached APICs, pod is 0.
- **TEP Pool:** (applicable only for directly connected APICs (ESXi virtual APIC and physical APIC)) enter the subnet of addresses used for internal fabric communication. The size of the subnet used will impact the scale of your pod.
You can not change this value after you have completed the configuration. Having to modify this value requires a wipe of the fabric.
- **Infrastructure VLAN:** Enter the VLAN ID for fabric connectivity (infra VLAN). This VLAN ID should be allocated solely to ACI, and not used by any other legacy device(s) in your network. Default value is 3914. The range is from 0 to 4093.
You can not change this value after you have completed the configuration. Having to modify this value requires a wipe of the fabric.
- **Enable IPv6 on APICs** (not applicable for virtual APIC on AWS): Put a check in this box if you want to enable IPv6 addresses for out-of-band management.

Step 6 Click **Next**.

Step 7 The third step is entering the **Controller Registration** details. Click **Add Controller** to add the first APIC (of the cluster). Enter the following details:

- **Controller Type:** The bootstrapping procedure auto-detects the deployment for which the configuration is being carried out. Based on that, either **Virtual** or **Physical** is chosen. The options displayed for the virtual and physical controller types are discussed in substeps (a) and (b), respectively. Follow either of these substeps based on the controller type.
 - a) When the Controller Type is **Virtual**:
 - **Virtual Instance:** The management IP used to access the APIC cluster bringup GUI. Only for the first APIC, this IP address is auto-populated. For the nodes that you subsequently add to the cluster, you will need to enter the management IP address and click **Validate**.
The management IP addresses are defined during the deployment of the VMs using ESXi/AWS. As mentioned in the prerequisites, keep all the required IP addresses handy while bringing up the cluster.
 - **General pane**

- Name: User-defined name for the controller.
 - Controller ID: The ID is auto-populated. If this is the first APIC of the cluster, the ID is "1". If you are adding the second controller of the cluster, "2" is auto-populated (and so on).
 - Pod ID: (Applicable only for *directly connected* virtual APIC on ESXi) The pod ID is auto-populated for APIC 1 of the cluster. For subsequent controllers of the cluster, enter a value. The range is from 1 to 128.
 - Serial Number: The serial number of the virtual machine is auto-populated.
- Out of Band Network pane
 - IPv4 Address: The IP address is displayed (as defined during the deployment).
 - IPv4 Gateway: The IP address is displayed (as defined during the deployment).

If you have enabled IPv6 addresses for OOB management earlier (step 5), enter the IPv6 address and gateway.

- Infra L3 Network pane (this pane is displayed only if the **Connection Type** that you chose earlier is *Remotely attached through an L3 network*.
 - IPv4 Address: Enter the infra network address.
 - IPv4 Gateway: Enter the IP address of the gateway.
 - VLAN: (Applicable only for *remotely attached* virtual APIC- ESXi) Enter the interface VLAN ID to be used.

The Infra L3 Network pane does not display when you deploy the virtual APIC using AWS.

After you have entered and saved the first APIC details, click **Add Controller** on the **Controller Registration** screen to add another APIC to the cluster.

b) When the Controller Type is **Physical**:

- CIMC Details pane
 - IP Address: The CIMC IP address. Only for the first Cisco APIC, this IP address is auto-populated. When you add more controllers to the cluster, you need to enter the CIMC IP addresses.
 - Username: The username to access the CIMC. The username is auto-populated (for the first controller and subsequent controllers).
 - Password: Enter the password to access CIMC. For the first controller, the password is auto-populated. For the subsequent controllers, enter the password.
 - Click **Validate**. *Validation success* is displayed on successful authentication.

If the CIMC is unreachable from the Cisco APIC out of band management IP address due to the CIMC NIC mode settings, change the NIC mode or enter JSON strings to perform the bootstrap.

- General pane
 - Name: Enter a name for the controller.
 - Controller ID: If it is the first controller of the cluster, "1" is auto-populated. If it is the second controller, "2" is auto-populated, and so on (increasing order).

- **Pod ID:** (applicable only for a directly-connected APIC) the pod ID is auto-populated for APIC 1 of the cluster. For subsequent controllers of the cluster, enter a value. The range is from 1 to 128.
- **Serial Number:** The serial number is auto-populated (for APICs 1 to N, where N is the cluster size) after CIMC validation.

APIC 1 verifies the reachability of the CIMC IP addresses and also captures the serial number of the new APICs.

- **Out of Band Network pane**
 - **IPv4 Address:** For APIC 1, the address is auto-populated. For subsequent APICs, enter the IP address (as defined during the deployment).
 - **IPv4 Gateway:** For APIC 1, the gateway address is auto-populated. For subsequent APICs, enter the IP address (as defined during the deployment).

If you have enabled IPv6 addresses for OOB management earlier (step 5), enter the IPv6 address and gateway.

- **Infra L3 Network pane** (this pane is displayed only if the **Connection Type** that you chose earlier is remotely attached through a Layer 3 network).
 - **IPv4 Address:** Enter the infra network IP address.
 - **IPv4 Gateway:** Enter the infra network IP address of the gateway.
 - **VLAN:** Enter a VLAN ID.

On the **Controller Registration** screen, after you have entered and saved the first APIC details, click **Add Controller** to add another APIC to the cluster.

(Optional, applicable only for virtual APICs) On the **Controller Registration** screen, put a check in the **Import existing security certificates** box to import existing security certificates for fabric recovery in virtual APICs. After putting a check in the box, enter the required details in the following fields:

- The **Remote Server IP Address** which contains the configuration file.
- The **Remote Path** which contains the configuration file.
- The configuration **File Name**.
- The **AES Encryption Passphrase** which was earlier used while backing up the configuration. The backup configuration file is linked to this key (passphrase).
- Choose the **Protocol**. The choices are:
 - **FTP**
 - **SFTP**
 - **SCP**
- **Remote Port**
- (applicable only for SFTP and SCP **Protocols**) Choose the **Authentication Type**. The choices are:
 - **Use Password**
 - **Use SSH Private Key Files**

- The **Username** to access the remote server.
- The **Password** to authenticate access to the remote server.
- (applicable only for Use SSH Private Key Files **Authentication Type**) Enter the **SSH Key Contents** here.
- (applicable only for Use SSH Private Key Files **Authentication Type**) Specify the **SSH Key Passphrase** used for encrypting the private key.

For details about the Import/Export procedure, see the [Cisco ACI Configuration Files: Import and Export](#) document.

The **Import existing security certificates** is applicable only for virtual APICs (deployed using AWS/ ESXi). Physical APICs have in-built certificates. However, in case of virtual APICs, when you are restoring using backup configuration to recover the fabric, the existing security certificates can be re-used.

Step 8 Click **Next**.

The **Next** button is disabled until all the controllers for a cluster are added. This is defined by the value you have entered for **Cluster Size** in the **Cluster Details** screen.

You can use the **Back** button to navigate to an earlier screen. After adding an APIC, click **Edit Details** to edit the information for an APIC. Except the first APIC, you can delete the other controllers, if required, by clicking the delete icon.

Step 9 In the **Summary** screen, review the updates, and click **Deploy**.

Step 10 The **Cluster Status** page is displayed, which shows the current status of the cluster formation. Wait for a few minutes after which you will be automatically redirected to the standard Cisco APIC GUI.

Provisioning IPv6 Management Addresses on APICs

IPv6 management addresses can be provisioned on the Cisco Application Policy Infrastructure Controller (APIC) at setup time or through a policy once the Cisco APIC is operational. Pure IPv4, pure IPv6, or dual stack (that is, both IPv6 and IPv4 addresses) are supported. A snippet of a typical setup screen that describes how to set up dual stack (IPv6 and IPv4) addresses for out-of-band management interfaces during the setup is given below. However, the following questionnaire is applicable for releases prior to 6.0(2). From Cisco APIC release 6.0(2), the cluster bringup is using the GUI as detailed above.

Cluster configuration ...

```
Enter the fabric name [ACI Fabric1]:
Enter the number of controllers in the fabric (1-9) [3]:
Enter the controller ID (1-3) [1]:
Enter the controller name [apic1]: infraipv6-ifc1
Enter address pool for TEP addresses [10.0.0.0/16]:
Note: The infra VLAN ID should not be used elsewhere in your environment
and should not overlap with any other reserved VLANs on other platforms.
Enter the VLAN ID for infra network (1-4094): 3914
Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:
```

Out-of-band management configuration ...

```
Enable IPv6 for Out of Band Mgmt Interface? [N]: Y (Enter Y to Configure IPv6 Address for
Out of Band Management Address)
Enter the IPv6 address [0:0:0:0:ffff:c0a8:a01/40]: 2001:420:28e:2020:0:ffff:ac1f:88e4/64
(IPv6 Address)
Enter the IPv6 address of the default gateway [None]: 2001:420:28e:2020:acc:68ff:fe28:b540
```

```
(IPv6 Gateway)
Enable IPv4 also for Out of Band Mgmt Interface? [Y]: (Enter Y to Configure IPv4 Address
for Out of Band Management Address)
Enter the IPv4 address [192.168.10.1/24]: 172.31.136.228/21 (IPv4 Address)
Enter the IPv4 address of the default gateway [None]: 172.31.136.1 (IPv4 Gateway)
Enter the interface speed/duplex mode [auto]:

admin user configuration ...
Enable strong passwords? [Y]:
Enter the password for admin:

Reenter the password for admin:
```



Note While using the **APIC Cluster Bringup** GUI, you can select the **Enable IPv6** option to use IPv6 addresses.

Accessing the GUI

Step 1 Open one of the supported browsers:

- Chrome version 59 (at minimum)
- Firefox version 54 (at minimum)
- Internet Explorer version 11 (at minimum)
- Safari version 10 (at minimum)

Note A known issue exists with the Safari browser and unsigned certificates. Read the information presented here before accepting an unsigned certificate for use with WebSockets. When you access the HTTPS site, the following message appears:

“Safari can’t verify the identity of the website APIC. The certificate for this website is invalid. You might be connecting to a website that is pretending to be an APIC, which could put your confidential information at risk. Would you like to connect to the website anyway?”

To ensure that WebSockets can connect, you must do the following:

Click **Show Certificate**.

Choose **Always Trust** in the three drop-down lists that appear.

If you do not follow these steps, WebSockets will not be able to connect.

Step 2 Enter the URL: **https://mgmt_ip-address**

Use the out-of-band management IP address that you configured during the initial setup. For example, https://192.168.10.1.

Note Only https is enabled by default. By default, http and http-to-https redirection are disabled.

Note If you see the following error message when logging into your Cisco APIC:

Need a valid webtoken cookie (named APIC-Cookie) or a signed request with signature in the cookie.

This is due to a known issue that occurs when you are logging into a Cisco APIC using both https and http. See the "Important Notes" section in [Setting up the Cisco APIC](#), on page 5 for more information on this issue and the workaround.

Step 3 When the login screen appears, enter the administrator name and password that you configured during the initial setup.

Step 4 In the **Domain** field, from the drop-down list, choose the appropriate domain that is defined.

If multiple login domains are defined, the **Domain** field is displayed. If the user does not choose a domain, the DefaultAuth login domain is used for authentication by default. This may result in login failure if the username is not in the DefaultAuth login domain.

What to do next

To learn about the features and operation of the Application Centric Infrastructure fabric and the Application Policy Infrastructure Controller, see the available white papers and the *Cisco Application Centric Infrastructure Fundamentals Guide*.

Accessing the REST API

By using a script or a browser-based REST client, you can send an API POST or GET message of the form:

https://apic-ip-address/api/api-message-url

Use the out-of-band management IP address that you configured during the initial setup.

- Note**
- Only https is enabled by default. By default, http and http-to-https redirection are disabled.
 - You must send an authentication message to initiate an API session. Use the administrator login name and password that you configured during the initial setup.

Accessing the NX-OS Style CLI

You can access the APIC NX-OS style CLI either directly from a terminal or through the APIC GUI.

For information about using the NX-OS style CLI commands, see the *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide* and the *Cisco APIC NX-OS Style CLI Command Reference*.

Guidelines and Restrictions for the APIC NX-OS Style CLI

- The CLI is supported only for users with administrative login privileges.

- The APIC NX-OS style CLI uses similar syntax and other conventions to the Cisco NX-OS CLI, but the APIC operating system is not a version of Cisco NX-OS software. Do not assume that a Cisco NX-OS CLI command works with or has the same function on the APIC CLI.
- If FIPS is enabled in the Cisco ACI setups, then SHA256 support is mandatory on the SSH Client. Additionally, to have the SHA256 support, the openssh-client must be running version 6.6.1 or higher.
- In releases earlier than Cisco APIC Release 1.2, the default CLI was a Bash shell with commands to directly operate on managed objects (MOs) and properties of the Management Information Model. Beginning with Cisco APIC Release 1.2, the default CLI is a NX-OS style CLI. The object model CLI is available by typing the **bash** command at the initial CLI prompt.

Accessing the NX-OS Style CLI from a Terminal

- Step 1** From a secure shell (SSH) client, open an SSH connection to APIC at *username@ip-address*.
Use the administrator login name and the out-of-band management IP address that you configured during the initial setup. For example, *admin@192.168.10.1*.
- Step 2** When prompted, enter the administrator password.
-

What to do next

When you enter the NX-OS style CLI, the initial command level is the EXEC level. You can stay in EXEC mode or you can type **configure** to enter global configuration mode. In any mode, type **?** to see the available commands.

For information about using the NX-OS style CLI commands, see the *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide* and the *Cisco APIC NX-OS Style CLI Command Reference*.

Accessing the NX-OS Style CLI from the GUI

- Step 1** From the menu bar, choose **System > Controllers**.
- Step 2** In the navigation pane, click **Controllers**.
- Step 3** Right-click the desired APIC and choose **Launch SSH**.
- Step 4** Follow the displayed instructions to open an SSH session to the selected controller.
-

What to do next

When you enter the NX-OS style CLI, the initial command level is the EXEC level. You can stay in EXEC mode or you can type **configure** to enter global configuration mode. In any mode, type **?** to see the available commands.

For information about using the NX-OS style CLI commands, see the *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide* and the *Cisco APIC NX-OS Style CLI Command Reference*.

Accessing the Object Model CLI



Note In releases earlier than Cisco APIC Release 1.2, the default CLI was a Bash shell with commands to directly operate on managed objects (MOs) and properties of the Management Information Model. Beginning with Cisco APIC Release 1.2, the default CLI is a NX-OS style CLI. The object model CLI is available by typing the **bash** command at the initial CLI prompt.

Step 1 From a secure shell (SSH) client, open an SSH connection to *username@ip-address*.

Use the administrator login name and the out-of-band management IP address that you configured during the initial setup. For example, `ssh admin@192.168.10.1`.

Step 2 When prompted, enter the administrator password that you configured during the initial setup.

You are now in the NX-OS style CLI for APIC.

Step 3 Type **bash** to enter the object model CLI.

Step 4 To return to the NX-OS style CLI, type **exit**.

This example shows how to enter the object model CLI and how to return to the NX-OS style CLI:

```
$ ssh admin@192.168.10.1
Application Policy Infrastructure Controller
admin@192.168.10.1's password: cisco123
apic# <---- NX-OS style CLI prompt
apic# bash
admin@apic1:~> <---- object model CLI prompt
admin@apic1:~> exit
apic#
```

What to do next

Every user must use the shared directory called `/home`. This directory gives permissions for a user to create directories and files; files created within `/home` inherit the default umask permissions and are accessible by the user and by root. We recommend that users create a `/home/userid` directory to store files, such as `/home/jsmith`, when logging in for the first time.

For more information about accessing switches using the ACI CLI using modes of operation such as BASH or VSH, see the *Cisco APIC Command Line Interface User Guide* and the *Cisco ACI Switch Command Reference*.

For detailed information about configuring the APIC CLI, see the *Cisco APIC Object Model Command Line Interface User Guide*.



CHAPTER 3

APIC GUI Overview

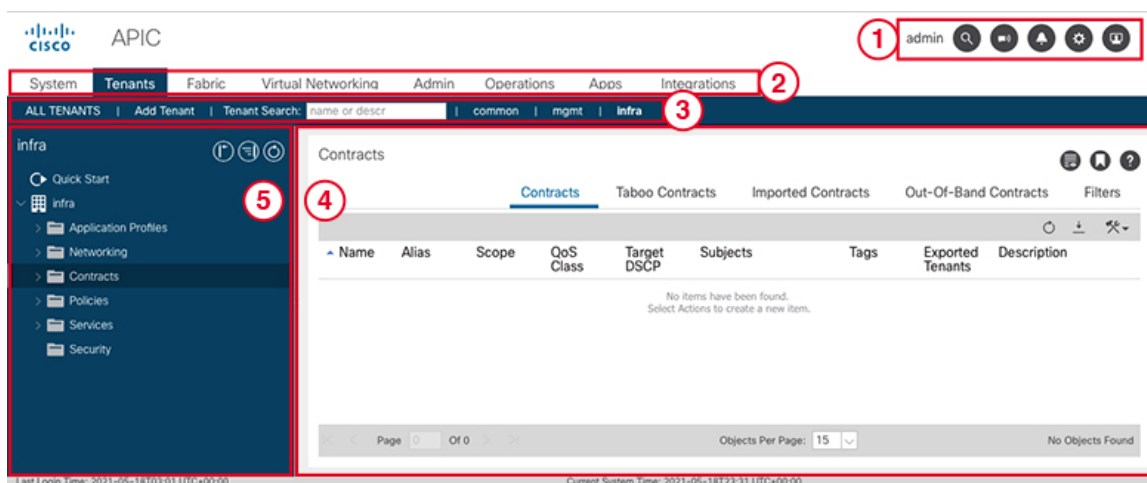
This chapter contains the following sections:

- [Overview of the GUI, on page 25](#)
- [Menu Bar and Submenu Bar, on page 26](#)
- [Navigation Pane, on page 33](#)
- [Work Pane, on page 33](#)
- [Personalizing the Interface, on page 34](#)
- [Single-Browser Session Management, on page 35](#)
- [Deployment Warning and Policy Usage Information, on page 35](#)
- [Graphical Configuration of Ports, on page 36](#)
- [Viewing an API Interchange in the GUI, on page 37](#)
- [GUI Icons, on page 39](#)
- [Next generation user interface, on page 40](#)

Overview of the GUI

The APIC GUI is a browser-based graphical interface for configuring and monitoring the ACI fabric. The GUI is organized to provide hierarchical navigation to all components, logical and physical, of the overall system. The primary control regions of the GUI are shown in the following figure.

Figure 1: APIC GUI Regions



The functions of these regions are described in the following links:

1. Menu bar tools—See [Menu Bar and Submenu Bar, on page 26](#)
2. Menu bar—See [Menu Bar and Submenu Bar, on page 26](#)
3. Submenu bar—See [Menu Bar and Submenu Bar, on page 26](#)
4. Work pane—See [Work Pane, on page 33](#)
5. Navigation pane—See [Navigation Pane, on page 33](#)

Below the Navigation pane, the Last Login is displayed, showing the date and time of the last instance of the current user's login.

As you operate the GUI to make configuration changes and retrieve information, the GUI communicates with the underlying operating system by exchanging REST API messages. You can observe these API messages using the API Inspector tool described in [Viewing an API Interchange in the GUI, on page 37](#).








Menu Bar and Submenu Bar

The menu bar is displayed across the top of the APIC GUI. The menu bar provides access to the main configuration tabs, along with access to tools such as search, notifications, and preferences. Immediately below the menu bar is the submenu bar, which presents specific configuration areas for each selected menu bar tab. The submenu bar tabs are different for each menu bar tab and might also differ depending upon your specific configuration or privilege level.



Tip In the APIC GUI configuration instructions, you will see notation such as **Fabric > Fabric Policies**. In this example, you are asked to click the **Fabric** tab in the menu bar followed by the **Fabric Policies** tab in the submenu bar.

At the far right side of the menu bar are the following menu bar tools:

Menu Bar Tools	Description
<i>username</i>	The name of the currently logged in local user.
	Search, on page 30
	Launch the Multi-Site Manager, on page 30
	Feedback, on page 30
	Alerts, on page 30
	Tool, on page 31
	Help, on page 31
	Manage my Profile, on page 32

The individual menu bar tabs and tools are described in the following sections.

Menu Bar Tabs

System Tab

Use the **System** tab to collect and display a summary of the overall system health, its history, and a table of system-level faults.

In addition, the **System** tab provides the following functions:

- You can configure global system policies in the **System Settings** submenu.
- You can view your licensing status in the **Smart Licensing** submenu.
- You can view user sessions in the **Active Sessions** submenu.

Tenants Tab

Use the **Tenants** tab in the menu bar to perform tenant management. The submenu bar provides a list of all tenants, an **Add Tenant** link, and links to three built-in tenants plus up to two of the most recently used tenants.

- A tenant contains policies that enable qualified users domain-based access control. Qualified users can access privileges such as tenant administration and networking administration.
- A user requires read/write privileges for accessing and configuring policies in a domain. A tenant user can have specific privileges into one or more domains.
- In a multitenancy environment, a tenant provides group user access privileges so that resources are isolated from one another (such as for endpoint groups and networking). These privileges also enable different users to manage different tenants.

The built-in tenants are:

- The **common** tenant is preconfigured for defining policies that provide common behavior for all the tenants in the fabric. A policy defined in the common tenant is usable by any tenant.
- The **infra** tenant is preconfigured for configuration related to the fabric infrastructure
- The **mgmt** tenant is preconfigured for inband and out-of-band connectivity configurations of hosts and fabric nodes (leafs, spines, and controllers).



Note For Layer 2 configuration of ports, you can type into the node and path fields to filter ports.

Fabric Tab

The **Fabric** tab contains the following tabs in the submenu bar:

- **Inventory** tab—Displays the individual components of the fabric.
- **Fabric Policies** tab—Displays the monitoring and troubleshooting policies and fabric protocol settings or fabric maximum transmission unit (MTU) settings.
- **Access Policies** tab—Displays the access policies that apply to the edge ports of the system. These ports are on the leaf switches that communicate externally.

Virtual Networking Tab

Use the **Virtual Networking** tab to view and configure the inventory of the various virtual machine (VM) managers. You can configure and create various management domains under which connections to individual management systems (such as VMware vCenters or VMware vShield) can be configured. Use the **Inventory** tab in the submenu bar to view the hypervisors and VMs that are managed by these VM management systems (also referred to as controllers in API).

Admin Tab

Use the **Admin** tab to perform administrative functions such as authentication, authorization, and accounting functions, scheduling policies, retaining and purging records, upgrading firmware, and controlling features such as syslog, Call Home, and SNMP.

Operations Tab

The **Operations** tab provides the following built-in tools for planning and monitoring fabric resources.

- **Visibility & Troubleshooting**—Shows the location of specified end points in the fabric and displays the traffic path, including any L4-L7 devices.
- **Capacity Dashboard**—Displays the available capacity of configurable resources such as end points, bridge domains, tenants, and contexts.
- **EP Tracker**—Enables you to view virtual and bare metal endpoint connections and disconnections to leaf switches and FEXes.
- **Visualization**—Provides visualization of traffic maps.

Capacity Dashboard

The capacity dashboard displays the available capacity of configurable resources such as endpoints, bridge domains, tenants, and contexts. The dashboard contains the following tabs:

- **Fabric Capacity**: Displays the capacity of the managed objects within the fabric. Each tile provides the current and maximum capacity of each object, as well as the percentage of the maximum capacity that is used. You can hover your cursor over some of the tiles to see more information.
- **Leaf Capacity**: Displays the capacity of the managed objects for each leaf switch that the Cisco Application Policy Infrastructure Controller (APIC) manages.
 - For all of the objects, the GUI displays the current resource usage and maximum resource capacity, as well as the percentage of the maximum resource capacity that is used.
 - The data for some of the objects is split into subcategories, such as local and remote for ESG MAC addresses.
 - The data for MAC, IPv4, and IPv6 addresses shows the total number of local and remote addresses.
 - The data for /32 routes and /128 routes provides the following information:
 - **UC**: The total IPv4 /32 or /128 unicast routes. This value persists through each interval without resetting to zero.
 - **EP**: The total IPv4 /32 or /128 endpoints. This value persists through each interval without resetting to zero.
 - **MCast**: The total IPv4 /32 or /128 multicast routes. This value persists through each interval without resetting to zero.
 - You can click the **Configure Profile** button in the **Switch** column to configure the forward scale profile for that switch.
 - You can click any other part of a row to see detailed capacity usage information for that switch. For resources with the **Absolute** entry, this is the current resource usage. In the case of /32 and /128 routes, **Absolute** is the total of unicast routes, endpoints, and multicast routes being used. **Percentage** is the percentage of the maximum resource capacity that being used.

Apps Tab

The **Apps** tab displays all the applications installed or uploaded to APIC. The tab allows an APIC administrator to upload, enable, upgrade, install, or uninstall a packaged application in APIC.

Integrations Tab

Use the **Integrations** tab to view all third-party integrations.

Menu Bar Tools

Search

Click the Search icon to display the search field. The search field enables you to locate objects by name or other distinctive fields.

Figure 2: Search



The search function allows the use of wildcards (*).

Launch the Multi-Site Manager

Click the Multi-Site Manager icon to launch the multi-site manager. The multi-site manager allows you can launch site APICs.

Figure 3: Launch the Multi-Site Manager



Feedback

Click the feedback menu bar icon to send comments to Cisco.

Figure 4: Feedback



Alerts

Click the alert menu bar icon to view a list of active alerts. When system alerts are available, a numeric badge will appear on the alert icon indicating the number of active alerts. When critical system notifications are available, the alert icon will blink red. To view the alerts, click the following icon.

Figure 5: Alerts



To disable blinking of the alert icon, remove all critical alerts from the alert list. A disabled **Close** button on a critical alert indicates that you must first resolve the underlying issue before the alert can be cleared.

Tool

To access the system tools, click the following menu bar icon and select an item from the drop-down list.

Figure 6: Tool



The following selections are available:

- **ACI Fabric Setup**—Open the ACI Fabric Setup. This panel help you set up the basic APIC infrastructure.
- **Show API Inspector**—Open the API Inspector, which is a built-in tool of the APIC that allows you to view the internal API messages between the GUI and the APIC operating system to execute tasks. For more information, see [Viewing an API Interchange in the GUI, on page 37](#).
- **Start Remote Logging**—Forward logging information to a remote URL.
- **Object Store Browser**—Open the Managed Object Browser, or Visore, which is a utility built into APIC that provides a graphical view of the managed objects (MOs) using a browser.
- **Show Debug Info**—Open a status bar at the bottom of the GUI to display information such as current managed object (MO) and system time. When the status bar is open, this selection changes to **Hide Debug Info**.
- **Config Sync Issues**— Open the Configuration Objects Pending Resolution panel. This panel shows if there are any transactions involving user-configurable objects that have yet to take effect in APIC. You can use information in the panel to help with debugging.



Note Global system settings are configured in **System > System Settings**.

Help

To access the Help tools, click the following menu bar icon and select an item from the drop-down list.

Figure 7: Help



The following selections are available:

- **Help**—Display links to API documentation and to the APIC
- **What's New** —Display splash screen showing recent features.
- **About**—Display the APIC version.

Manage my Profile

To configure settings and preferences for the logged in user, click the following menu bar icon and select an item from the drop-down list.

Figure 8: Manage my Profile



The following selections are available:

- **Bookmarks**—Display links to menus bookmarked by the user.
 - Menus that display the Favorites icon (★) can be bookmarked by clicking the icon.
- **Change My Password**—Change the password of the currently logged in local user.
- **Change My SSH Keys**—Change the user's public SSH key used for certificate-based login.
- **Change My X509 Certificate**—Change the user's X.509-format certificate for login.
- **View My Permissions**—Display the user's role-based read and write privileges for domains and accessible objects.
- **Settings**—Change general GUI settings.
 - **Remember Tree Selection**—Enable the GUI to keep the navigation tree expanded when returning to a window. For example, if you enable this property and expand the navigation tree in the Tenants tab, click on the Fabric tab, then return to the Tenants tab, the tree will remain expanded.
 - **Preserve Tree Divider Position**—Enable the GUI to keep the position of the tree divider after dragging the tree divider to the desired location.
 - **Disable Notification on Success**—Suppress the success dialog box notification.
 - **Disable Deployment Warning at Login**—Disable the Deployment Warning dialog box when logging in. See [Deployment Warning and Policy Usage Information, on page 35](#).
 - **Default Page Size for Tables**—Set the GUI table size.
 - **Show All UI Sections**—Display hidden UI configuration options.
 - **Show What's New at Login**—Display splash screen at login, showing recent features.
 - **Enable Single-Browser Session (SBS)**—Allows logging in to the APIC GUI and then opening additional browser tabs or windows to the same APIC without being required to log in from each new tab or window. See [Single-Browser Session Management, on page 35](#).
- **Change Deployment Settings**—Enable and set the scope of the deployment notification. See [Deployment Warning and Policy Usage Information, on page 35](#).
- **Logout**—Exit the APIC configuration GUI.

Navigation Pane

Use the **Navigation** pane, which is on the left side of the APIC GUI below the submenu bar, to navigate to all elements of the submenu category.

For each submenu category, the **Navigation** pane is organized as a hierarchical tree of objects, logical and physical, related to that category. These objects typically represent ports, policies, or groupings of other objects. When you select an object in the **Navigation** pane, details of the object display in the **Work** pane.

When you right-click an object in the **Navigation** pane, you might be presented with a menu of possible actions related to the object, such as one or more of the following actions:

- **Delete**—Delete the object.
- **Create <type of object>**—Create a new object.
- **Save as...**—Download the object and its properties in JSON or XML format to a local file.
- **Post...**—Export the object and its properties to an existing local file.
- **Share**—Displays the URL of the object. You can copy the URL and send it to others.
- **Open In Object Store Browser**—Open the object in Visore, a built-in utility that displays an object and its properties. This information may be useful in troubleshooting or for developing API tools.
- **Clone**—Create a copy of the object. This action is useful for deriving a new contract or policy based on an existing contract or policy.



Note If any container in the **Navigation** pane, for example **Application Profiles** under a **Tenant**, contains more than 40 profiles, you cannot click on a profile and expand it in the **Navigation** pane. You must select the desired profile from the **Work** pane and expand it.

Work Pane

Use the **Work** pane, which is on the right side of the APIC GUI, to display details about the component that you selected in the **Navigation** pane.

The **Work** pane includes the following elements:

- A content area that displays tabs. These tabs enable you to access information that is related to the component that you chose in the **Navigation** pane. The tabs displayed in the content area depend upon the selected component.
- For some components, a link to conceptual information related to the component, represented by a list

icon in the upper right corner.



- You can bookmark almost any page, which enables you to go back to that page easily by choosing the bookmark from your list of bookmarks.

Bookmarked links are accessible from the **User Profile and Preferences** icon in the Menu Bar.

- You can mark a tab as the "favorite" on a page. Whenever you navigate to that page, that tab will be the default tab that is displayed. This feature is enabled only for the tabs in the **Work** pane; you cannot mark a menu bar tab as a favorite.

Common Pages in the Work Pane

In addition to displaying specific task menus, the Work pane also displays several types of special-purpose menus described in this section.

Quick Start Pages

Many APIC menu and submenu tabs open an initial Quick Start page, which summarizes the purpose of the tab, provides links to step-by-step instructions and videos for commonly-used procedures, and provides shortcut links to commonly-used subsections within the tab. An overall Quick Start page at **System > QuickStart** assists you in performing common and basic procedures, providing step-by-step instructions, available concept information, and links to main functional areas in the GUI.

Dashboard Pages

Dashboard pages provide at-a-glance summaries of the status of the ACI system and major system components, including health score trends, components with below-threshold health scores, and fault counts. You can configure health score thresholds to determine when components will appear in the dashboard. The system dashboard page at **System > Dashboard** summarizes the health of the overall ACI system, while switch dashboard pages at **Fabric > Inventory > Pod n > component > Dashboard** summarize the health and faults of each spine and leaf switch.

Summary Pages

Many top-level folders in the Navigation pane display tile-based Summary pages in the Work pane that link to subfolders. Some Summary pages, such as those in **Fabric > Inventory > Pod n**, contain tiles summarizing major components along with brief health and fault information for each component. Other Summary pages, such as those in **Fabric > Fabric Policies > Policies**, contain tiles that describe the configuration areas served by the contained folders.

Personalizing the Interface

Naming the APIC GUI

An ACI controller cluster comprises three or more APICs. In some cases, it might be helpful to know which APIC you are viewing. Perform the following steps to add a custom name to the heading of the APIC GUI.

-
- Step 1** On the APIC menu bar, choose **System > System Settings**.
 - Step 2** In the **Navigation** pane, click **APIC Identification Preferences**.
 - Step 3** In the work pane, type the desired APIC name in the **GUI Alias** box.
 - Step 4** Click **Submit**.

The APIC name appears in parentheses at the top left of the GUI.

Adding a Login Banner to the CLI or GUI

You can define banners to be displayed when the user is prompted to log in to the CLI or GUI. The CLI banner is a simple text string to be printed at the terminal before the password prompt. You can define a banner for the APIC CLI and a separate banner for the switch CLI. The GUI banner displays at the APIC URL before user login authentication. The GUI banner is defined as a URL of a site hosting the desired HTML.

Step 1 On the APIC menu bar, choose **System > System Settings**.

Step 2 In the **Navigation** pane, click **APIC Identification Preferences**.

Step 3 In the work pane, complete the following fields as desired:

- a) To configure an APIC CLI banner, type the banner text into the **Controller CLI Banner** textbox.
- b) To configure a switch CLI banner, type the banner text into the **Switch CLI Banner** textbox.
- c) To configure an APIC GUI banner, type the URL of a site hosting the desired HTML into the **GUI Banner (URL)** textbox.

Note The URL site owner must allow the site to be placed in an iFrame to display the informational banner. If the owner of the site sets the `x-frame-option` to `deny` or `sameorigin`, the site the URL points to will not appear.

Step 4 Click **Submit**.

Single-Browser Session Management

Beginning with Cisco APIC Release 4.0(1), you can log in to the APIC GUI and then open additional browser tabs or windows to the same APIC without being required to log in from each new tab or window. This behavior is disabled by default and can be enabled by checking the **Enable Single-Browser Session (SBS)** checkbox in the **User Profile and Preferences > Settings** menu from the main menu bar tools.

If you want to log in to APIC from different tabs or windows of a browser using different credentials, make sure the single-browser session management feature is disabled.

Deployment Warning and Policy Usage Information

By configuring **Deployment Warning Settings**, you can enable the automatic display of policy usage information whenever you modify or delete policies that might affect other resources or policies. The policy usage information allows you to identify which resources and policies are being used by the policy that you are currently modifying or deleting. Tables display the nodes where the given policy is used and other policies that use this policy. By default, usage information is displayed within a dialog box whenever you attempt to modify a policy. Also, at any time, you can click the **Show Usage** button at the bottom of the screen to view the same information.

The **Deployment Warning Settings** dialog box allows you to enable and alter the scope of deployment notification that displays policy usage information. You can access this dialog box by selecting **Change**

Deployment Settings in the menu bar tool **User Settings and Preferences** drop-down list or through a button on the **Policy Usage Information** dialog box.

When the **Policy** tab is selected in the upper right corner of the **Deployment Warning Settings** dialog box, you can configure the following policy options:

- **(Global) Show Deployment Warning on Delete/Modify**—Enable the **Deployment Warning** notification for every policy deletion or modification across the APIC.
- **(Local) Show Deployment Warning on Delete/Modify**—Set the rule for the **Deployment Warning** notification for specific policy configuration.
 - **Use Global Settings**—Use the setting selected for **(Global) Show Deployment Warning on Delete/Modify**.
 - **Yes**—Display the **Deployment Warning** notification before submitting configuration modifications on any policy change. Valid for this browser session only.
 - **No**—Do not display the **Deployment Warning** notification before submitting configuration modifications on any policy change. Valid for this browser session only.

When the **History** tab is selected in the upper right corner of the **Deployment Warning Settings** dialog box, you can view tables of **Events** and **Audit Log** entries for previous deployment warnings.

Graphical Configuration of Ports

The APIC GUI provides a graphical method for configuring ports, port channels, and virtual port channels on the leaf switches in the fabric, configure ports for dynamic breakout, and link interfaces to FEX switches. This configuration capability is present in the following GUI locations:

- **Fabric > Inventory > Topology**
- **Fabric > Inventory > Pod**
- **Fabric > Inventory > Pod > Leaf**
- **Fabric > Inventory > Pod > Spine**

In the Work pane's **Interface** tab, click on the + button (at the top left), select one or more switches to configure, and click **Add Selected**. To select multiple switches, use **Ctrl+Click** or **Shift+Click**.

The switches are graphically displayed with their ports and links. If you have configured a breakout port, a block containing the sub ports is displayed below the leaf diagram.



Note If you accessed the **Interface** tab from a leaf switch, the leaf switch is automatically added.

Select the interfaces to configure. When interfaces are selected, the available configuration buttons appear. Depending on the number of selected interfaces and where they are located, you can then click one of the following buttons at the top of the page:

- **L2**—Layer 2. Visible when you click one or more leaf interfaces on the switch diagrams.
- **PC**—Port Channel. Visible when you click one or more leaf interfaces on the switch diagrams.

- **VPC**—Virtual Port Channel. Visible when you click at least one interface on two switch diagrams.
- **FEX**—Fabric Extender. Visible when you click one or more leaf interfaces on the switch diagrams.
- **Breakout**—Breakout mode. Visible when you click one or more leaf interfaces on the switch diagrams.
- **Fabric**—Add policies to a fabric interface. Visible when you click a port that is eligible to be a fabric port.
- **Uplink and Downlink**—Convert eligible uplinks to downlinks and vice versa.
- **Spine**—Visible when you click one or more leaf interfaces on the switch diagrams.

Viewing an API Interchange in the GUI

When you perform a task in the APIC graphical user interface (GUI), the GUI creates and sends internal API messages to the operating system to execute the task. By using the API Inspector, which is a built-in tool of the APIC, you can view and copy these API messages. A network administrator can replicate these messages in order to automate key operations, or you can use the messages as examples to develop external applications that will use the API.

Step 1 Log in to the APIC GUI.

Step 2 In the upper right corner of the APIC window, click the System Tools icon to view the drop-down list.

Step 3 In the drop-down list, choose the **Show API Inspector**.

The **API Inspector** opens in a new browser window.

Step 4 In the **Filters** toolbar of the **API Inspector** window, choose the types of API log messages to display.

The displayed messages are color-coded according to the selected message types. This table shows the available message types:

Name	Description
trace	Displays trace messages.
debug	Displays debug messages. This type includes most API commands and responses.
info	Displays informational messages.
warn	Displays warning messages.
error	Displays error messages.
fatal	Displays fatal messages.
all	Checking this checkbox causes all other checkboxes to become checked. Unchecking any other checkbox causes this checkbox to be unchecked.

Step 5 In the **Search** toolbar, you can search the displayed messages for an exact string or by a regular expression.

This table shows the search controls:

Name	Description
Search	In this text box, enter a string for a direct search or enter a regular expression for a regex search. As you type, the first matched field in the log list is highlighted.
Reset	Click this button to clear the contents of the Search text box.
Regex	Check this checkbox to use the contents of the Search text box as a regular expression for a search.
Match case	Check this checkbox to make the search case sensitive.
Disable	Check this checkbox to disable the search and clear the highlighting of search matches in the log list.
Next	Click this button to cause the log list to scroll to the next matched entry. This button appears only when a search is active.
Previous	Click this button to cause the log list to scroll to the previous matched entry. This button appears only when a search is active.
Filter	Check this checkbox to hide nonmatched lines. This checkbox appears only when a search is active.
Highlight all	Check this checkbox to highlight all matched fields. This checkbox appears only when a search is active.

Step 6 In the **Options** toolbar, you can arrange the displayed messages.

This table shows the available options:

Name	Description
Log	Check this checkbox to enable logging.
Wrap	Check this checkbox to enable wrapping of lines to avoid horizontal scrolling of the log list
Newest at the top	Check this checkbox to display log entries in reverse chronological order.
Scroll to latest	Check this checkbox to scroll immediately to the latest log entry.
Clear	Click this button to clear the log list.
Close	Click this button to close the API Inspector.

Example







This example shows two debug messages in the API Inspector window:








```
13:13:36 DEBUG - method: GET url: http://192.0.20.123/api/class/infraInfra.json
response: {"imdata":{"infraInfra":{"attributes":{"instanceId":"0:0","childAction":"","dn":"uni/infra","lcOwn":"local","name":"","replTs":"never","status":""}}}}
```

```
13:13:40 DEBUG - method: GET url: http://192.0.20.123/api/class/l3extDomP.json?
query-target=subtree&subscription=yes
response: {"subscriptionId":"72057598349672459","imdata":[]}
```


GUI Icons





Table 4: Frequently Displayed Icons in the APIC GUI

Icons	Description
	Search, on page 30
	Alerts, on page 30
	Manage my Profile, on page 32
	Tool, on page 31
	Bookmark this page
	Displays concept information for the current menu page
	Quick Start
	Plays a Quick Start video
	Displays a Quick Start procedure
	Link to related section
	Topology
	Pod

Icons	Description
	Collapse Tree View
	Expand Tree View
	Collapse All Nodes
	Displays a drop-down list of actions
	Refresh the displayed information
	Download to a file
	Upload a file

Fault, Statistics, and Health Level Icons

Table 5: Severity Levels of Faults Displayed in the APIC GUI

Icons	Description
	Critical—This icon displays a fault level with critical severity.
	Major—This icon displays a fault level with major severity.
	Minor—This icon displays a fault level with minor severity.
	Warning—This icon displays a fault level that requires a warning.

Next generation user interface

This release introduces a preview of the next generation user interface (UI), which includes a modern and simplified Cisco Application Policy Infrastructure Controller (APIC) GUI. With this preview, you can get an idea of upcoming development of the GUI. Any selection that is grayed out in the new GUI does not exist in this preview, but will be added in a future release. You cannot make any configuration changes using the preview; the new GUI will allow you to make configuration changes in a future release.

Figure 9: The new GUI, on page 41 shows the new GUI.

Figure 9: The new GUI




Table 6: Areas of the new GUI, on page 41 describes the areas of the new GUI.

Table 6: Areas of the new GUI

Callout	Description
1	Tools area
2	Navigation area
3	Tabs area
4	Work area

Tools area

This area contains basic tools that appear regardless of the GUI page.

- **Help:** Opens a menu with these choices:
 - **About:** Shows the Cisco APIC release, copyright year, and current system time.
 - **APIC Fabric Setup:** Helps you determine what you have not yet configured for your fabric.
 - **Help Center:** Provides links to some of the more commonly-used documentation.
 - **Open Object Store Browser:** Opens the Managed Object Browser, or Visore, which is a utility built into Cisco APIC that provides a graphical view of the managed objects (MOs) using a browser.
 - **Show API Inspector:** Opens the API Inspector, which is a built-in tool of the Cisco APIC that allows you to view the internal API messages between the GUI and the Cisco APIC operating system to execute tasks.

- **Notifications:** Shows notifications about key issues with the Cisco APIC. Using the drop-down menu, you can choose to show all notifications or only notifications of a specific severity. The menu includes only fault severities that are currently raised. You can also click **View All** to see all faults in a table.
- **Profile:** Opens a menu with these choices:
 - **Switch to existing UI:** Switches back to the existing (old) GUI.
 - **User Preferences:** Configures your user settings.
 - **Change Password:** Changes your password.
 - **Change User Certificate:** Changes your login certificate.
 - **Change SSH Key:** Changes the public SSH key that you use for certificate-based login.
 - **Logout:** Logs you out from the Cisco APIC.

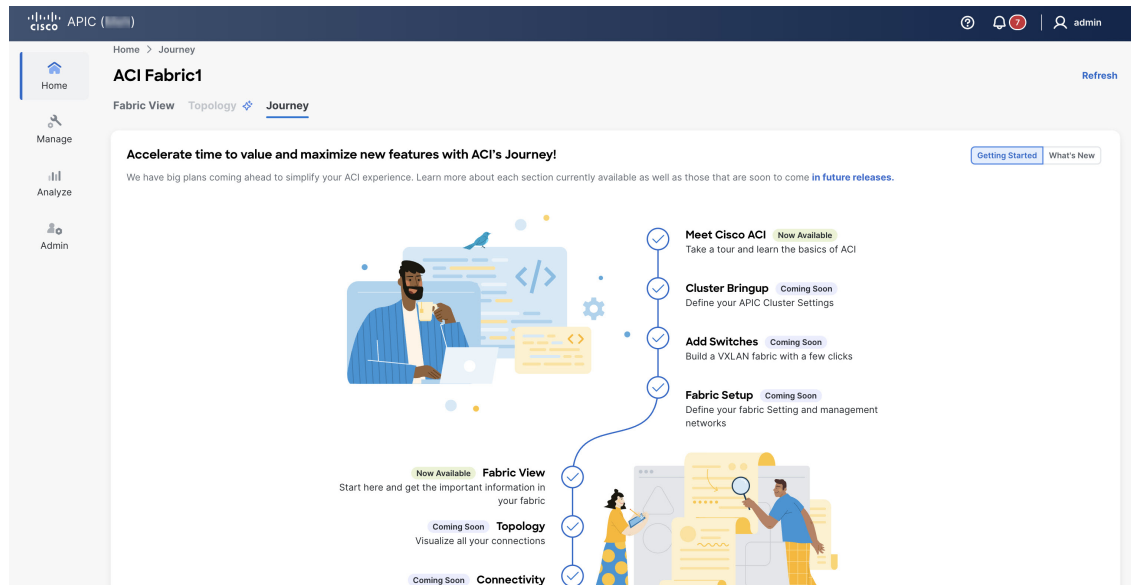
Navigation area

This is the core navigation for the Cisco APIC, which is present on all GUI pages.

- **Home:** Enables you to view information about the fabric. This area has these tabs:
 - **Fabric View:** Shows various information about your the fabric. You can click blue text as well as **Controllers, Switches, and Pods** to see more information about those things.
 - **Topology:** Shows a graphical representation of your fabric's topology.
 - **Journey:** Shows the main functionalities of Cisco APIC, displays the execution plan for the new GUI, and provides a preview of upcoming features and their functionality.

Most of the stages have a **Tell Me More** button, which opens a dialog that gives you more information about the part of the GUI that is associated with that stage, including a screenshot. For a stage that is labeled as "Coming Soon," the dialog outlines the upcoming development for that part of the GUI and the screenshot shows how the GUI might look.

Figure 10: Cisco ACI journey



- **Manage:** Enables you to view information about and configure these things:
 - **Inventory:** Shows information about the controllers, switches, and pods.
 - **Tenants:** Shows information about the tenants.
 - **Policy Groups:** Shows information about the policy groups.
 - **Policies:** Shows information about the configured policies.
 - **Domains and Pools:** Shows information about the domains and pools.
 - **Software Management:** Enables you to upgrade the firmware of the Cisco APICs and switches.
 - **Monitoring Destinations:** Shows information about the monitoring destinations, such as for callhome, SNMP, or TACACS.
 - **Security:** Shows information about the security configuration of your fabric.
 - **Switches and Interfaces:** Shows information about the switches and interfaces.
- **Analyze:** Enables you to view these things:
 - **Faults:** Shows the faults that Cisco APIC raised due to issues with your fabric.
 - **History and Logs:** Shows the various records and logs, such as fault records, event records, and audit logs.
 - **Tech Support:** Enables you to export the tech support logs, which you provide when you contact the Cisco Technical Assistance Center (TAC).
 - **Apps:** Shows the apps that you have installed and enables you to use them.
- **Admin:** Enables you to use or configure these things:
 - **Backup and Restore:** Enables you to backup and restore the Cisco APIC configurations.

- **Authentication:** Enables you to configure how users must authenticate to log into the Cisco APIC.
- **Active Sessions:** Shows the users who are logged into the Cisco APIC.
- **Integrations:** Shows information about third-party products that you have integrated with Cisco APIC.
- **Licensing:** Shows information about your active licenses for using Cisco APIC and its features.
- **System Settings:** Enables you to configure the Cisco APIC.

Tabs area

This area contains tabs that change depending on the core navigation that you chose. The tabs enable you to navigate to subsections of that core navigation.


Work area

This area changes depending on which page you are viewing. The work area shows status, contains links to view configured objects, and much more. Any text in blue is a link that you can click to see more information about or change the configuration of something.

Preview the next generation user interface


This release introduces a preview of the next generation user interface (UI), which includes a modern and simplified Cisco Application Policy Infrastructure Controller (APIC) GUI. Follow these steps to preview the next generation user interface.

- Step 1** Click the profile icon in the upper right of any page and choose **Switch to new UI**.



You cannot view anything in the new GUI that is grayed out. These pages will be added in a future release.

- Step 2** To change back to the original GUI so that you can make configuration changes, click the profile icon in the upper right of any page and choose **Switch to existing UI**.





CHAPTER 4

Fabric Initialization and Switch Discovery

This chapter contains the following sections:

- [Initializing the Fabric, on page 45](#)
- [Switch Discovery, on page 50](#)
- [Maintenance Mode, on page 60](#)
- [Cisco NX-OS to Cisco ACI POAP Auto-conversion, on page 62](#)
- [Cisco Nexus 9000 Switch Secure Erase, on page 64](#)

Initializing the Fabric

About Fabric Initialization

You can build a fabric by adding switches to be managed by the APIC and then validating the steps using the GUI, the CLI, or the API.



Note Before you can build a fabric, you must have already created an APIC cluster over the out-of-band network.


Fabric Topology (Example)

An example of a fabric topology is as follows:

- Two spine switches (spine1, spine2)
- Two leaf switches (leaf1, leaf2)
- Three instances of APIC (apic1, apic2, apic3)

The following figure shows an example of a fabric topology.

Figure 11: Fabric Topology Example



Connections: Fabric Topology

An example of the connection details for the fabric topology is as follows:

Name	Connection Details
leaf1	eth1/1 = apic1 (eth2/1) eth1/2 = apic2 (eth2/1) eth1/3 = apic3 (eth2/1) eth1/49 = spine1 (eth5/1) eth1/50 = spine2 (eth5/2)
leaf2	eth1/1 = apic1 (eth 2/2) eth1/2 = apic2 (eth 2/2) eth1/3 = apic3 (eth 2/2) eth1/49 = spine2 (eth5/1) eth1/50 = spine1 (eth5/2)
spine1	eth5/1 = leaf1 (eth1/49) eth5/2 = leaf2 (eth1/50)
spine2	eth5/1 = leaf2 (eth1/49) eth5/2 = leaf1 (eth1/50)

Multi-Tier Fabric Topology (Example)

3-tier Core-Aggregation-Access architectures are common in data center network topologies. As of the Cisco APIC Release 4.1(1), you can create a multi-tier ACI fabric topology that corresponds to the

Core-Aggregation-Access architecture, thus mitigating the need to upgrade costly components such as rack space or cabling. The addition of a tier-2 leaf layer makes this topology possible. The tier-2 leaf layer supports connectivity to hosts or servers on the downlink ports and connectivity to the leaf layer (aggregation) on the uplink ports.


In the multi-tier topology, the leaf switches initially have uplink connectivity to the spine switches and downlink connectivity to the tier-2 leaf switches. To make the entire topology an ACI fabric, all ports on the leaf switches connecting to tier-2 leaf fabric ports must be configured as fabric ports (if not already using the default fabric ports). After APIC discovers the tier-2 leaf switch, you can change the downlink port on the tier-2 leaf to a fabric port and connect to an uplink port on the middle layer leaf.



Note If you are not using the default fabric ports to connect leaf switches to tier-2 leaf, you must convert the leaf ports from downlink to uplink (leaf switch reload required). For more information about changing port connectivity, see the Access Interfaces chapter of the *Cisco APIC Layer 2 Networking Configuration Guide*.

The following figure shows an example of a multi-tier fabric topology.

Figure 12: Multi-Tier Fabric Topology Example



While the topology in the above image shows the Cisco APIC and L3Out/EPG connected to the leaf aggregation layer, the tier-2 leaf access layer also supports connectivity to APICs and L3Out/EPGs.



Note Only Cisco Nexus 9000 Series switches with model numbers that end in EX, and later are supported as tier-2 leaf switches and as leaf switches, if there are tier-2 leaf switches attached to them. See the table below. Tier-2 leaf switches attached to remote leaf switches are not supported.

Table 7: Supported Switches and Port Speeds for Multi-Tier Architecture

Switch	Maximum supported downlink port (as tier-2 leaf)	Maximum supported fabric ports (as tier-2 leaf)	Maximum supported fabric ports (as tier-1 leaf)
Nexus 93180YC-EX	48x1/10/25-Gbps 4x40/100-Gbps	48 x 10/25-Gbps 6 x 40/100-Gbps	48 x 10/25-Gbps 6 x 40/100-Gbps
Nexus 93108TC-EX	48x100M/1/10G BASE-T 4x40/100-Gbps	6 x 40/100-Gbps	6 x 40/100-Gbps
N9K-9348GC-FXP**	48 x 100M/1G BASE-T	4 x 10/25-Gbps 2 x 40/100-Gbps	4 x 10/25-Gbps 2 x 40/100-Gbps
N9K-93180YC-FX	48 x 1/10/25-Gbps 4x40/100-Gbps	48 x 10/25-Gbps 6 x 40/100-Gbps	48 x 10/25-Gbps 6 x 40/100-Gbps
N9K-93108TC-FX	48 x 100M/1/10G BASE-T 4x40/100-Gbps	6 x 40/100-Gbps	6 x 40/100-Gbps
N9K-93240YC-FX2	48x1/10/25-Gbps 10x40/100-Gbps	48x1/10/25-Gbps 12x40/100-Gbps	48x10/25-Gbps fiber ports 12x40/100-Gbps
N9K-C9336C-FX2	34 x 40/100-Gbps	36 x 40/100-Gbps	36 x 40/100-Gbps
N9K-C93216TC-FX2***	96 x 10G BASE-T 10 x 40/100-Gbps	12 x 40/100-Gbps	12 x 40/100-Gbps
N9K-C93360YC-FX2***	96 x 10/25-Gbps 10 x 40/100-Gbps	52 x 10/25Gbps 12 x 40/100Gbps	52 x 10/25Gbps 12 x 40/100Gbps
N9K-C9364C-GX	62 x 40/100-Gbps	62 x 40/100-Gbps	62 x 40/100-Gbps

* Last 2 original fabric ports cannot be used as downlink ports.

** If tier-2 leaf does not require much bandwidth, it can be used as tier-1 though it has fewer fiber ports. Copper port cannot be used as a fabric port.

*** Supported beginning with Cisco APIC Release 4.1(2).

Changing the External Routable Subnet

These procedures describe how to change the external routable subnet, if you find that you have to make changes to the information in the subnets or TEP table after you've made those configurations.



Note Changing an external routable subnet configuration using multiple subnets is not supported.

-
- Step 1** Navigate to the area where you originally configured the external routable subnet.
- On the menu bar, click **Fabric > Inventory**.
 - In the Navigation pane, click **Pod Fabric Setup Policy**.
 - On the **Fabric Setup Policy** panel, double-click the pod where you originally configured the external routable subnet.
The **Fabric Setup Policy for a POD** page for this pod appears.
 - Locate the information for the subnets or TEP table, depending on the release of your APIC software:
 - For releases prior to 4.2(3), locate the **Routable Subnets** table.
 - For 4.2(3) only, locate the **External Subnets** table.
 - For 4.2(4) and later, locate the **External TEP** table.
- Step 2** Locate the external routable subnet that you want to delete in the table and determine if the state of that subnet is set to **active** or **inactive**.
- If the state is set to **active**, change the state to **inactive**:
- Double-click on the entry in the subnets or TEP table for the existing external routable subnet that you want to delete.
 - Change the state for the subnet to **inactive**, then click **Update**.
- Step 3** Delete the existing external routable subnet.
- Click on the entry in the subnets or TEP table for the existing external routable subnet that you want to delete.
 - Click the trashcan icon at the top of the table, then click **Yes** in the pop-up confirmation window to delete the external routable subnet.
- Step 4** Wait for at least 30 seconds, then configure a new external routable subnet.
- Click + in the subnets or TEP table to configure a new external routable subnet.
 - Enter the IP address and Reserve Address, if necessary, and set the state to **active** or **inactive**.
 - The IP address is the subnet prefix that you wish to configure as the routeable IP space.
 - The Reserve Address is a count of addresses within the subnet that must not be allocated dynamically to the spine switches and remote leaf switches. The count always begins with the first IP in the subnet and increments sequentially. If you wish to allocate the Unicast TEP from this pool, then it must be reserved.
 - Click **Update** to add the new external routable subnet to the subnets or TEP table.
 - On the **Fabric Setup Policy** panel, click **Submit**.
- Step 5** Verify that the new routable IP address is configured correctly.
- Log into the APIC controller through the CLI and enter the following command:
- ```
apic1# avread | grep routableAddress
```
- Output similar to the following should appear:
- ```
routableAddress 14.3.0.228 14.3.0.229 14.3.1.228
```
- Step 6** Check the NAT entries created on the spine switch.
- Log into the spine switch through the CLI and enter the following command:
- ```
spine1# show natable
```

Output similar to the following should appear:

```
-----NAT TABLE-----
Private Ip Routable Ip

10.0.0.2 14.3.0.229
10.0.0.1 14.3.0.228
10.0.0.3 14.3.1.228
```

---

## Switch Discovery

### About Switch Discovery with the APIC

The APIC is a central point of automated provisioning and management for all the switches that are part of the ACI fabric. A single data center might include multiple ACI fabrics; each data center might have its own APIC cluster and Cisco Nexus 9000 Series switches that are part of the fabric. To ensure that a switch is managed only by a single APIC cluster, each switch must be registered with that specific APIC cluster that manages the fabric.

The APIC discovers new switches that are directly connected to any switch it currently manages. Each APIC instance in the cluster first discovers only the leaf switch to which it is directly connected. After the leaf switch is registered with the APIC, the APIC discovers all spine switches that are directly connected to the leaf switch. As each spine switch is registered, that APIC discovers all the leaf switches that are connected to that spine switch. This cascaded discovery allows the APIC to discover the entire fabric topology in a few simple steps.

### Switch Registration with the APIC Cluster

After a switch is registered with the Cisco Application Policy Infrastructure Controller (APIC), the switch is part of the Cisco APIC-managed fabric inventory. With the Cisco Application Centric Infrastructure (ACI) fabric, the Cisco APIC is the single point of provisioning, management, and monitoring for switches in the infrastructure.

The following guidelines and limitations apply:

- Before you begin registering a switch, make sure that all switches in the fabric are physically connected and booted in the desired configuration. For information about the installation of the chassis, see <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-installation-guides-list.html>.
- The infrastructure IP address range must not overlap with other IP addresses used in the ACI fabric for in-band and out-of-band networks.
- When a switch is power cycled or upgraded, downlink interfaces will be in the admin-down state until the switch can download the configurations again from the Cisco APICs to prevent external devices from sending traffic to the switch that is not yet ready. Fabric links and down links for Cisco APIC connectivity are exempt from being changed to the admin-down state. To achieve this exemption, the leaf switch remembers the downlink interface that was connected to the Cisco APICs prior to the power cycle or

upgrade. Because of this, you must not change the Cisco APIC connectivity until the switches are fully operational again after the power cycle or upgrade.

## Switch Role Considerations

- The default fabric links must be used for initial switch discovery from another switch.
- If a default spine switch is connected to the Cisco Application Policy Infrastructure Controller (APIC) directly, the switch will be converted to a leaf switch automatically. During the conversion period, the fault will be present in the Cisco APIC, which is a normal behaviour. The fault will be removed after the switch conversion is finished.
- For a leaf switch, you can configure a port profile to convert a port to be a downlink or fabric link after the port is registered to the Cisco APIC. For more information, see the *Cisco APIC Layer 2 Networking Configuration Guide* at the following site:

[https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html#Configuration\\_Guides](https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html#Configuration_Guides)

The following table specifies the default role for the switches for which you are able to change their role:

**Table 8: Default Switch Roles**

| Switch Product ID | Default Role | First Release to Support a Role Change <sup>1</sup> |
|-------------------|--------------|-----------------------------------------------------|
| N9K-C93600CD-GX   | Leaf         | 5.2(1)                                              |
| N9K-C9364C-GX     | Leaf         | 5.1(3)                                              |
| N9K-C9316D-GX     | Spine        | 5.1(4)                                              |

<sup>1</sup> Specifies the first release to support changing the role change for the indicated switch. Role changing for that switch is supported in all subsequent releases.

## Registering an Unregistered Switch Using the GUI



**Note** The infrastructure IP address range must not overlap with other IP addresses used in the ACI fabric for in-band and out-of-band networks.

### Before you begin

Make sure that all switches in the fabric are physically connected and booted.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
- Step 2** In the Navigation pane, choose **Fabric Membership**.
- Step 3** In the Work pane, click the **Nodes Pending Registration** tab.

Switches in the **Nodes Pending Registration** tab table can have the following conditions:

- A newly discovered but unregistered node has a node ID of 0 and has no IP address.
- A manually entered (in Cisco Application Policy Infrastructure Controller (APIC)) but unregistered switch has an original status of **Undiscovered** until it is physically connected to the network. Once connected, the status changes to **Discovered**.

**Step 4** In the **Nodes Pending Registration** table, locate a switch with an ID of 0 or a newly connected switch with the serial number you want to register.

**Step 5** (Optional) Double-click the row of that node if you want to see more information about the node.

A dialog appears that shows you various node properties, such as the ACI-mode switch release and information about LLDP neighbors.

**Step 6** Right-click the row of that switch, choose **Register**, and perform the following actions:

- Verify the displayed Serial Number to determine which switch is being added.
- Configure or edit the following settings:

| Field              | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Pod ID</b>      | Identifier of the pod where the node is located.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Node ID</b>     | <p>A number greater than 100. The first 100 IDs are reserved for Cisco APIC appliance nodes.</p> <p><b>Note</b> We recommend that leaf nodes and spine nodes be numbered differently. For example, number spines in the 100 range (such as 101, 102) and number leaves in the 200 range (such as 201, 202).</p> <p>After the node ID is assigned, it cannot be updated. After the node has been added to the <b>Registered Nodes</b> tab table, you can update the node name by right-clicking the table row and choosing <b>Edit Node and Rack Name</b>.</p> |
| <b>RL TEP Pool</b> | Tunnel endpoint (TEP) pool identifier for the node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Node Name</b>   | The node name, such as leaf1 or spine3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Role</b>        | <p>The assigned node role. The options are:</p> <ul style="list-style-type: none"> <li>• <b>spine</b></li> <li>• <b>leaf</b></li> <li>• <b>virtualleaf</b></li> <li>• <b>virtualspine</b></li> <li>• <b>remote leaf</b></li> <li>• <b>tier-2-leaf</b></li> </ul> <p>If you choose a role other than the default role for the node, the node automatically reboots during the registration to change the role.</p>                                                                                                                                             |

| Field            | Setting                                                                                                                                  |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rack Name</b> | The name of the rack in which the node is installed. Choose <b>Default</b> , or choose <b>Create Rack</b> to add a name and description. |

c) Click **Register**.

Cisco APIC assigns an IP address to the node and the node is added to the **Registered Nodes** tab table. Next and if applicable, other nodes that are connected to this node are discovered and appear in the **Nodes Pending Registration** tab table.

**Step 7** Continue to monitor the **Nodes Pending Registration** tab table. As more nodes appear, repeat these steps to register each new node until all installed nodes are registered.

## Adding a Switch Before Discovery Using the GUI

You can add a switch description before the switch is physically connected to the network by following these steps:

### Before you begin

Make sure that you know the serial number of the switch.

**Step 1** On the menu bar, choose **Fabric > Inventory**.

**Step 2** In the Navigation pane, choose **Fabric Membership**.

**Step 3** On the **Registered Nodes** or **Nodes Pending Registration** work pane, click the Actions icon, then click **Create Fabric Node Member**.

The **Create Fabric Node Member** dialog appears.

**Step 4** Configure the following settings:

| Field                | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Pod ID</b>        | Identify the pod where the node is located.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Serial Number</b> | Required: Enter the serial number of the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Node ID</b>       | <p>Required: Enter a number greater than 100. The first 100 IDs are reserved for Cisco Application Policy Infrastructure Controller (APIC) appliance nodes.</p> <p><b>Note</b> We recommend that you number leaf nodes and spine nodes differently. For example, number leaf nodes in the 100 range (such as 101, 102) and number spine nodes in the 200 range (such as 201, 202).</p> <p>After the node ID is assigned, it cannot be updated. After the node has been added to the <b>Registered Nodes</b> tab table, you can update the node name by right-clicking the table row and choosing <b>Edit Node and Rack Name</b>.</p> |
| <b>Switch Name</b>   | The node name, such as leaf1 or spine3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Field                | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Node Type</b>     | <p>Choose the type (role) for the node. The options are:</p> <ul style="list-style-type: none"> <li>• <b>leaf</b></li> </ul> <p>Put a check in one of the following boxes if applicable:</p> <ul style="list-style-type: none"> <li>• <b>Is Remote:</b> Specifies that the node is a remote leaf switch.</li> <li>• <b>Is Virtual:</b> Specifies that the node is virtual.</li> <li>• <b>Is Tier-2 Leaf:</b> The fabric node member (leaf switch) being created will take on the characteristics of a tier-2 leaf switch in a multi-tier architecture.</li> </ul> <ul style="list-style-type: none"> <li>• <b>spine</b></li> </ul> <p>Put a check in the following box if applicable:</p> <ul style="list-style-type: none"> <li>• <b>Is Virtual:</b> Specifies that the node is virtual.</li> </ul> <ul style="list-style-type: none"> <li>• <b>unknown</b></li> </ul> <p>If you choose a role other than the default role for the node, the node automatically reboots during the registration to change the role.</p> |
| <b>VPC Pair</b>      | Optional. If the node is part of a vPC pair, choose the ID of the node with which to pair this node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>VPC Domain ID</b> | Enter the vPC domain ID for the vPC pair. The range is from 1 to 1000. This field only appears if you entered a value for <b>VPC Pair</b> , and is required in that case.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

The Cisco APIC adds the new node to the **Nodes Pending Registration** tab table.

### What to do next

Connect the physical switch to the network. Once connected, the Cisco APIC matches the serial number of the physical switch to the new entry. Monitor the **Nodes Pending Registration** tab table until the **Status** for the new switch changes from **Undiscovered** to **Discovered**. Follow the steps in the [Registering an Unregistered Switch Using the GUI, on page 51](#) section to complete the fabric initialization and discovery process for the new switch.

## Switch Discovery Validation and Switch Management from the APIC

After the switches are registered with the APIC, the APIC performs fabric topology discovery automatically to gain a view of the entire network and to manage all the switches in the fabric topology.

Each switch can be configured, monitored, and upgraded from the APIC without having to access the individual switches.




## Validating the Registered Switches Using the GUI

- 
- Step 1** On the menu bar, navigate to **Fabric > Inventory > Fabric Membership**.
- Step 2** In the **Fabric Membership** work pane, click the **Registered Nodes** tab.  
The switches in the fabric are displayed in the **Registered Nodes** tab table with their node IDs. In the table, all the registered switches are displayed with the IP addresses that are assigned to them.
- 

## Validating the Fabric Topology

After all the switches are registered with the APIC cluster, the APIC automatically discovers all the links and connectivity in the fabric and discovers the entire topology as a result.

## Validating the Fabric Topology Using the GUI

- 
- Step 1** On the menu bar, navigate to **Fabric > Inventory > Pod number**.
- Step 2** In the **Work** pane, click the **Topology** tab.  
The displayed diagram shows all attached switches, APIC instances, and links.
- Step 3** (Optional) Hover over any component to view its health, status, and inventory information.
- Step 4** (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
- Step 5** (Optional) To refresh the topology diagram, click the  icon in the upper right corner of the **Work** pane.
- 

## Unmanaged Switch Connectivity in VM Management

The hosts that are managed by the VM controller (for example, a vCenter), can be connected to the leaf port through a Layer 2 switch. The only prerequisite required is that the Layer 2 switch must be configured with a management address, and this management address must be advertised by Link Layer Discovery Protocol (LLDP) on the ports that are connected to the switches. Layer 2 switches are automatically discovered by the APIC, and they are identified by the management address. To view the unmanaged switches in APIC, navigate to **Fabric > Inventory > Fabric Membership** and click the **Unmanaged Fabric Nodes** tab.

## Troubleshooting Switch Discovery Issues

The ACI-mode switch software includes a comprehensive leaf and spine switch discovery validation program. The validation program is launched with a switch CLI command when a switch is stuck in the discovery mode.

The validation program performs the following tests:

1. System state—Checks the state of the `topSystem` managed object (MO).
  - a. If the state is "out-of-service," checks for any scheduled upgrades.

- b. If the state is "downloading bootscript," a failure has occurred in the downloading bootscript. The failure is reported. If the switch is an L3out spine, the program additionally checks the bootstrap download state and reports any failure.
2. DHCP status—Checks for DHCP status and information, such as the TEP IP, node Id, and name assigned from the `dhcpResp` MO.
3. AV details—Checks whether the APICs are registered and whether they have valid IP addresses.
4. IP reachability—Uses the **iping** command to verify IP reachability to the address assigner APIC. To retest this condition, use the **show discoveryissues apic ipaddress** command.
5. infra VLAN received—Checks for the presence of the infra VLAN details in the `lldpInst` MO. If this switch belongs to a pod that has no APIC, no infra VLAN details are present, and this section of the test result can be ignored.
6. LLDP adjacency—Checks for the presence of LLDP adjacencies and for any wiring mismatch issues. LLDP issues can generate fault reports such as infra VLAN mismatch, chassis ID mismatch, or no connection to the front end ports.
7. Switch version—Reports the running firmware version of the switch. Also reports the version of the APIC, if available.
8. FPGA/BIOS—Checks for any FPGA/BIOS version mismatch on the switch.
9. SSL validation—Checks for validity of the SSL certificate details using the **acidiag verifyssl -s serialNumber** command.
10. Policy downloads—Checks the `pconsBootStrap` MO to see whether registration to APIC (PM shards) is complete and whether all policies were downloaded successfully.
11. Time—Reports the current time on the switch.
12. Hardware status—Checks the module, power, and fan status from the `eqptCh`, `eqptFan`, `eqptPsu`, `eqptFt` and `eqptLC` MOs.

### Running the Test Manually

To run the switch discovery validation program, log in to the spine or leaf switch CLI console and execute the following command:

```
show discoveryissues [apic ipaddress]
```

### Example of a Successful Test

The following example shows the switch discovery validation program output for a successful test.

```
spine1# show discoveryissues

Checking the platform type.....SPINE!
Check01 - System state - in-service [ok]
Check02 - DHCP status [ok]
 TEP IP: 10.0.40.65 Node Id: 106 Name: spine1
Check03 - AV details check [ok]
Check04 - IP rechability to apic [ok]
 Ping from switch to 10.0.0.1 passed
Check05 - infra VLAN received [ok]
```

```

infra vLAN:1093
Check06 - LLDP Adjacency [ok]
 Found adjacency with LEAF
Check07 - Switch version [ok]
 version: n9000-14.2(0.167) and apic version: 5.0(0.25)
Check08 - FPGA/BIOS out of sync test [ok]
Check09 - SSL check [check]
 SSL certificate details are valid
Check10 - Downloading policies [ok]
Check11 - Checking time [ok]
 2019-08-21 17:15:45
Check12 - Checking modules, power and fans [ok]

```

### Example of a Failed Test

The following example shows the switch discovery validation program output for a switch with discovery issues.

```

spine1# show discoveryissues

Checking the platform type.....SPINE!
Check01 - System state - out-of-service [FAIL]
 Upgrade status is notscheduled
 Node upgrade is notscheduled state
Check02 - DHCP status [FAIL]
 ERROR: discover not being sent by switch
 Ignore this, if the IP is already known by switch
 ERROR: node Id not configured
 ERROR: Ip not assigned by dhcp server
 ERROR: Address assigner's IP not populated
 TEP IP: unknown Node Id: unknown Name: unknown
Check03 - AV details check [ok]
Check04 - IP reachability to apic [FAIL]
 please rerun the CLI with argument apic Ip
 (show discoveryissues apic <ip>) to check its reachability from switch
Check05 - infra VLAN received [FAIL]
 Please ignore if this switch is part of a pod with no apic
Check06 - LLDP Adjacency [FAIL]
 Error: spine not connected to any leaf
Check07 - Switch version [ok]
 version: n9000-14.2(0.146) and apic version: unknown
Check08 - FPGA/BIOS out of sync test [ok]
Check09 - SSL check [ok]
 SSL certificate details are valid
Check10 - Downloading policies [FAIL]
 Registration to all PM shards is not complete
 Policy download is not complete
 Pcons bootstrap is in triggered state
Check11 - Checking time [ok]
 2019-07-17 19:26:29
Check12 - Checking modules, power and fans [FAIL]
 Line card state is testing

```

## Finding Your Switch Inventory Using the GUI

This section explains how to find your switch model and serial numbers using the Cisco APIC GUI.

**Before you begin**

You must have access to the Cisco APIC GUI

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
- Step 2** In the navigation pane, click a **Pod** icon.  
Your switch icons appear in the navigation pane.
- Step 3** In the navigation pane, click on a switch icon.  
A list of tabs appears at the top of the work pane.
- Step 4** Click the **General** tab.  
Your switch information appears in the work pane.
- 

## Troubleshooting Switch Discovery Issues

The ACI-mode switch software includes a comprehensive leaf and spine switch discovery validation program. The validation program is launched with a switch CLI command when a switch is stuck in the discovery mode.

The validation program performs the following tests:

1. System state—Checks the state of the `topSystem` managed object (MO).
  - a. If the state is "out-of-service," checks for any scheduled upgrades.
  - b. If the state is "downloading bootscript," a failure has occurred in the downloading bootscript. The failure is reported. If the switch is an L3out spine, the program additionally checks the bootstrap download state and reports any failure.
2. DHCP status—Checks for DHCP status and information, such as the TEP IP, node Id, and name assigned from the `dhcpResp` MO.
3. AV details—Checks whether the APICs are registered and whether they have valid IP addresses.
4. IP reachability—Uses the **iping** command to verify IP reachability to the address assigner APIC. To retest this condition, use the **show discoveryissues apic ipaddress** command.
5. infra VLAN received—Checks for the presence of the infra VLAN details in the `lldpInst` MO. If this switch belongs to a pod that has no APIC, no infra VLAN details are present, and this section of the test result can be ignored.
6. LLDP adjacency—Checks for the presence of LLDP adjacencies and for any wiring mismatch issues. LLDP issues can generate fault reports such as infra VLAN mismatch, chassis ID mismatch, or no connection to the front end ports.
7. Switch version—Reports the running firmware version of the switch. Also reports the version of the APIC, if available.
8. FPGA/BIOS—Checks for any FPGA/BIOS version mismatch on the switch.
9. SSL validation—Checks for validity of the SSL certificate details using the **acidiag verifyssl -s serialNumber** command.
10. Policy downloads—Checks the `pconsBootStrap` MO to see whether registration to APIC (PM shards) is complete and whether all policies were downloaded successfully.

11. Time—Reports the current time on the switch.
12. Hardware status—Checks the module, power, and fan status from the `eqptCh`, `eqptFan`, `eqptPsu`, `eqptFt` and `eqptLC` MOs.

### Running the Test Manually

To run the switch discovery validation program, log in to the spine or leaf switch CLI console and execute the following command:

```
show discoveryissues [apic ipaddress]
```

### Example of a Successful Test

The following example shows the switch discovery validation program output for a successful test.

```
spine1# show discoveryissues

Checking the platform type.....SPINE!
Check01 - System state - in-service [ok]
Check02 - DHCP status [ok]
 TEP IP: 10.0.40.65 Node Id: 106 Name: spine1
Check03 - AV details check [ok]
Check04 - IP reachability to apic [ok]
 Ping from switch to 10.0.0.1 passed
Check05 - infra VLAN received [ok]
 infra vLAN:1093
Check06 - LLDP Adjacency [ok]
 Found adjacency with LEAF
Check07 - Switch version [ok]
 version: n9000-14.2(0.167) and apic version: 5.0(0.25)
Check08 - FPGA/BIOS out of sync test [ok]
Check09 - SSL check [check]
 SSL certificate details are valid
Check10 - Downloading policies [ok]
Check11 - Checking time [ok]
 2019-08-21 17:15:45
Check12 - Checking modules, power and fans [ok]
```

### Example of a Failed Test

The following example shows the switch discovery validation program output for a switch with discovery issues.

```
spine1# show discoveryissues

Checking the platform type.....SPINE!
Check01 - System state - out-of-service [FAIL]
 Upgrade status is notscheduled
 Node upgrade is notscheduled state
Check02 - DHCP status [FAIL]
 ERROR: discover not being sent by switch
 Ignore this, if the IP is already known by switch
 ERROR: node Id not configured
 ERROR: Ip not assigned by dhcp server
 ERROR: Address assigner's IP not populated
 TEP IP: unknown Node Id: unknown Name: unknown
Check03 - AV details check [ok]
```

```

Check04 - IP reachability to apic [FAIL]
 please rerun the CLI with argument apic Ip
 (show discoveryissues apic <ip>) to check its reachability from switch
Check05 - infra VLAN received [FAIL]
 Please ignore if this switch is part of a pod with no apic
Check06 - LLDP Adjacency [FAIL]
 Error: spine not connected to any leaf
Check07 - Switch version [ok]
 version: n9000-14.2(0.146) and apic version: unknown
Check08 - FPGA/BIOS out of sync test [ok]
Check09 - SSL check [ok]
 SSL certificate details are valid
Check10 - Downloading policies [FAIL]
 Registration to all PM shards is not complete
 Policy download is not complete
 Pcons bootstrap is in triggered state
Check11 - Checking time [ok]
 2019-07-17 19:26:29
Check12 - Checking modules, power and fans [FAIL]
 Line card state is testing

```

## Maintenance Mode

### Maintenance Mode

Following are terms that are helpful to understand when using maintenance mode:

- **Maintenance mode:** Used to isolate a switch from user traffic for debugging purposes. You can put a switch in **maintenance mode** by enabling the **Maintenance (GIR)** field in the **Fabric Membership** page in the APIC GUI, located at **Fabric > Inventory > Fabric Membership** (right-click on a switch and choose **Maintenance (GIR)**).

If you put a switch in **maintenance mode**, that switch is not considered as a part of the operational ACI fabric infra and it will not accept regular APIC communications.

You can use maintenance mode to gracefully remove a switch and isolate it from the network in order to perform debugging operations. The switch is removed from the regular forwarding path with minimal traffic disruption.

In graceful removal, all external protocols are gracefully brought down except the fabric protocol (IS-IS) and the switch is isolated from the network. During maintenance mode, the maximum metric is advertised in IS-IS within the Cisco Application Centric Infrastructure (Cisco ACI) fabric and therefore the leaf switch in maintenance mode does not attract traffic from the spine switches. In addition, all front-panel interfaces on the switch are shutdown except for the fabric interfaces. To return the switch to its fully operational (normal) mode after the debugging operations, you must recommission the switch. This operation will trigger a stateless reload of the switch.

In graceful insertion, the switch is automatically decommissioned, rebooted, and recommissioned. When recommissioning is completed, all external protocols are restored and maximum metric in IS-IS is reset after 10 minutes.

The following protocols are supported:

- Border Gateway Protocol (BGP)

- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)
- Link Aggregation Control Protocol (LACP)

Protocol Independent Multicast (PIM) is not supported.

### Important Notes

- If a border leaf switch has a static route and is placed in maintenance mode, the route from the border leaf switch might not be removed from the routing table of switches in the ACI fabric, which causes routing issues.  
To work around this issue, either:
  - Configure the same static route with the same administrative distance on the other border leaf switch, or
  - Use IP SLA or BFD for track reachability to the next hop of the static route
- While the switch is in maintenance mode, the Ethernet port module stops propagating the interface related notifications. As a result, if the remote switch is rebooted or the fabric link is flapped during this time, the fabric link will not come up afterward unless the switch is manually rebooted (using the **acdiag touch clean** command), decommissioned, and recommissioned.
- While the switch is in maintenance mode, CLI 'show' commands on the switch show the front panel ports as being in the up state and the BGP protocol as up and running. The interfaces are actually shut and all other adjacencies for BGP are brought down, but the displayed active states allow for debugging.
- For multi-pod / multi-site, **IS-IS metric for redistributed routes** should be set to less than 63 to minimize the traffic disruption when bringing the node back into the fabric. To set the **IS-IS metric for redistributed routes**, choose **Fabric > Fabric Policies > Pod Policies > IS-IS Policy**.
- Existing GIR supports all Layer 3 traffic diversion. With LACP, all the Layer 2 traffic is also diverted to the redundant node. Once a node goes into maintenance mode, LACP running on the node immediately informs neighbors that it can no longer be aggregated as part of port-channel. All traffic is then diverted to the vPC peer node.
- The following operations are not allowed in maintenance mode:
  - **Upgrade**: Upgrading the network to a newer version
  - **Stateful Reload**: Restarting the GIR node or its connected peers
  - **Stateless Reload**: Restarting with a clean configuration or power-cycle of the GIR node or its connected peers
  - **Link Operations**: Shut / no-shut or optics OIR on the GIR node or its peer node
  - **Configuration Change**: Any configuration change (such as clean configuration, import, or snapshot rollback)
  - **Hardware Change**: Any hardware change (such as adding, swapping, removing FRU's or RMA)

## Removing a Switch to Maintenance Mode Using the GUI

Use this procedure to remove a switch to maintenance mode using the GUI. During the removal of a switch to maintenance mode, the out-of-band management interfaces will remain up and accessible.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
  - Step 2** In the navigation pane, click **Fabric Membership**.
  - Step 3** In the **Registered Nodes** table in the work pane, right-click the row of the switch to be removed to maintenance mode and select **Maintenance (GIR)**.
  - Step 4** Click **OK**.
- The gracefully removed switch displays **Maintenance** in the **Status** column.
- 

## Inserting a Switch to Operational Mode Using the GUI

Use this procedure to insert a switch to operational mode using the GUI.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
  - Step 2** In the navigation pane, click **Fabric Membership**.
  - Step 3** In the **Registered Nodes** table in the work pane, right-click the row of the switch to be inserted to operational mode and select **Commision**.
  - Step 4** Click **Yes**.
- 

## Cisco NX-OS to Cisco ACI POAP Auto-conversion

### About Cisco NX-OS to Cisco ACI POAP Auto-conversion

Beginning with the 5.2(3) release, Cisco NX-OS to Cisco Application Centric Infrastructure (ACI) power-on auto-provisioning (POAP) auto-conversion automates the process of upgrading software images and installing configuration files on nodes that are being deployed in the network for the first time. When a Cisco NX-OS node with the POAP auto-conversion feature boots and does not find the startup configuration, the node enters the POAP mode and starts DHCP discovery on all ports. The node locates a DHCP server and bootstraps itself with its interface IP address, gateway, and DNS server IP addresses. The device also obtains the IP address of a TFTP server and downloads a configuration script that enables the node to download and install the appropriate software image and configuration file. This process converts the Cisco NX-OS node from the standalone mode to the Cisco ACI-mode.

To auto-convert a Cisco NX-OS node to a Cisco ACI node using POAP, you need to specify an interface on a Cisco ACI switch node that is connected to the Cisco NX-OS node that needs to be auto-converted. The interface specified on the Cisco ACI switch enables the handling of POAP and allows the Cisco NX-OS node to use the Cisco Application Policy Infrastructure Controller (APIC) as its DHCP server for auto-conversion. The Cisco ACI switch node must be already registered to the Cisco ACI fabric and be active, meaning that



the node is reachable from the Cisco APIC cluster. This auto-conversion can be used both when adding a new switch to the fabric or when replacing an existing Cisco ACI switch.

## Guidelines and Limitations for Cisco NX-OS to Cisco ACI POAP Auto-conversion

The following guidelines and limitations apply when using Cisco NX-OS to Cisco Application Centric Infrastructure (ACI) power-on auto-provisioning (POAP) auto-conversion:

- Because a Cisco NX-OS node that is being converted starts to send discover packets on all interfaces including management, any external DHCP server (apart from the Cisco Application Policy Infrastructure Controller's (APIC's) server) should be removed, as they may intercept POAP discover packets and disrupt the conversion.
- Cisco NX-OS to Cisco ACI POAP auto-conversion is supported when the NX-OS device to be converted is connected to an existing Cisco ACI switch node that has reachability to the Cisco APIC cluster. Due to this reason, the following scenarios are not supported:
  - When discovering the first Cisco ACI switch from APIC 1.
  - When replacing a Cisco ACI leaf node when a Cisco APIC is singled-homed to the leaf node.
  - When adding or replacing a Cisco ACI switch that reaches to the Cisco APIC cluster only through an IPN device. That is, when adding a Cisco NX-OS node as a new remote leaf node, adding a Cisco NX-OS node as a first spine node in a new pod, replacing a remote leaf node, or replacing a spine node in a Cisco ACI Multi-Pod setup with only one spine node in the pod. This scenario is supported beginning with the Cisco APIC 5.2(4) release with the required configurations on the IPN device.
- Modular spine node supervisor replacement is not supported.
- POAP supports switches that have -EX, -FX, -GX, or a later suffix in the product IDs (PIDs), as well as the Cisco N9K-C9364C and N9K-C9332C switches.
- After you auto-convert a spine or leaf node, the **show system reset-reason** CLI command does not display any information regarding conversion. The output only states the following:

```
reset-requested-by-cli-command-reload
```
- You must use optical cables between Cisco ACI switches and Cisco NX-OS switches. You cannot use copper cables in this case.
- The Cisco ACI switch image that needs to be used for auto-conversion must be present on the Cisco APIC cluster's firmware repository. You can use the GUI to check that the image is present by going to **Admin > Firmware > Images**.

## Converting a Cisco NX-OS Node to Cisco ACI With POAP Auto-conversion Using the GUI

The following procedure converts an existing Cisco NX-OS node from the standalone mode to the Cisco ACI mode using power-on auto-provisioning (POAP) auto-conversion. This process does not decommission the node.

### Before you begin

You must have enabled **Auto Firmware Update on Switch Discovery** with the target Cisco ACI firmware version. For more information, see the *Cisco APIC Getting Started Guide*.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
- Step 2** In the Navigation pane, choose **Fabric Membership**.
- Step 3** In the Work pane, choose the **Registered Nodes** tab.
- Step 4** (Optional) When replacing an existing Cisco ACI switch node with a new switch that may be running NX-OS, right-click the node to be replaced and choose **Remove From Controller** as you would do for a regular replacement scenario.
- Step 5** In the action menu at the top right of the table, choose **Add with NXOS to ACI Conversion**.
- In the replacement scenario, if the switch node to be replaced is decommissioned or inactive, you can alternatively right-click the node and choose **Replace with NXOS to ACI Conversion**. This will perform actions **Remove From Controller** from step 4 and **Add with NXOS to ACI Conversion** from step 5 at the same time.
- Step 6** In the dialog, fill out the fields as follows:
- **Node ID:** Choose the ID of a node that is connected to the node that you want to convert. You can click the trash can to delete a node or + to add another node. Specify at least one node. You can click **Hide Interfaces** to hide the interface information if you need more space in the GUI when configuring additional nodes.
  - **Interface ID:** Choose the ID of one of the node's interfaces that is connected to the node that you want to convert. You can click the trash can to delete an interface or + to add another interface. Configure only one interface in each node to handle POAP for POAP auto-conversion.
- Step 7** Click **Submit**.
- Step 8** Choose the **Nodes Pending Registration** tab.
- After the node appears in this tab, the node registration procedure is the same as for regular Cisco ACI switches.
- Step 9** (Optional) After the switch is registered and has joined the fabric with the `Active` status, you may delete the POAP auto-conversion setting on the interface that you configured in step 6. After the conversion completes, delete the POAP settings from the connected node:
- a) Choose the **Registered Nodes** tab.
  - b) Double-click the row of the node from which you want to delete the POAP settings.
  - c) In the dialog, choose the **NXOS Conversion Policy** tab.
  - d) Select the pathname that you want to delete, then click the delete icon (the trashcan).
- 

## Cisco Nexus 9000 Switch Secure Erase

### About Cisco Nexus 9000 Switch Secure Erase

Cisco Nexus 9000 switches utilize persistent storage to maintain system software images, switch configuration, software logs, and operational history. Each of these areas can contain user-specific information such as details on network architecture and design, and potential target vectors for would-be attackers. The secure erase feature enables you comprehensively to erase this information, which you can do when you return a switch

with return merchandise authorization (RMA), upgrade or replace a switch, or decommission a system that has reached its end-of-life.

Secure erase is supported from Cisco APIC release 6.0(x). All the leaf and spine switches in the fabric must be APIC release 6.0(x) or later.

This feature erases user data in the following storage devices:

- SSD
- EMMC
- MTD
- CMOS
- NVRAM



---

**Note** Not every switch model has all these storage devices.

---

## Securely Erasing User Data from a Cisco Nexus 9000 Switch Using the GUI

Use the following procedure to securely erase user data from a Cisco Nexus 9000 switch using the GUI.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
  - Step 2** In the Navigation pane, choose **Fabric Membership**.
  - Step 3** In the Work pane, right-click the switch (node) that you want to securely erase and choose **Decommission**.
  - Step 4** In the **Decommission** dialog, choose **Decommission & Secure Remove**.
  - Step 5** Click **OK**.
- 

The decommission process takes from 2 to 8 hours, depending on the switch and SSD type. The process securely erases the switch and removes the switch configuration from the Cisco Application Policy Infrastructure Controller (APIC). The secure erase process does not remove the NX-OS image from the bootflash. The switch cannot join the fabric until you manually re-register the switch.

The switch reboots after the secure erase operation completes. To connect to the switch, you must use the terminal console because the IP address is not reachable.

## Securely Erasing User Data from a Module of a Cisco Nexus 9000 Modular Switch Line Card Using the GUI

Use the following procedure to securely erase user data from a module of a Cisco Nexus 9000 modular switch line card using the GUI.

- 
- Step 1** On the menu bar, choose **Fabric > Inventory**.
  - Step 2** In the Navigation pane, choose **pod\_id > node\_id > Chassis > Line Modules > slot\_id**.

**Step 3** Right-click the slot ID and choose **Disable**.

**Step 4** In the **Disable** dialog, click **Secure Erase**.

---

The decommission process takes from 30 minutes to 2 hours, depending on the switch and SSD type. The process securely erases the data from the module of the switch and removes the module's configuration from the Cisco Application Policy Infrastructure Controller (APIC). The process does not remove the NX-OS image from the bootflash.

After the secure erase operation completes, the module will be in the powered-down state. To connect to the switch, you must use the terminal console because the IP address is not reachable.

## Securely Erasing User Data from a Cisco Nexus 9000 Switch Using the Switch's CLI

Use the following procedure to securely erase user data from a Cisco Nexus 9000 switch using the switch's CLI. You cannot use the Cisco Application Policy Infrastructure Controller (APIC)'s CLI for this procedure.

### Before you begin

Decommission the switch or disconnect the switch physically from the fabric before performing the secure erase operation using CLI. If you do not decommission the switch or disconnect the switch physically from the fabric, after the secure erase process completes, the Cisco APIC pushes the configuration back to the switch.

---

**Step 1** Log into the switch's CLI.

**Step 2** Enter the virtual shell.

```
leaf1# vsh
```

**Step 3** Disable the terminal's session timeout.

```
leaf1# terminal session-timeout 0
```

If you do not disable the timeout, the VSH session can time out and exit before the secure erase completes and can provide status.

**Step 4** Reset the switch to the factory settings, which also securely erases your data from the switch.

```
leaf1# factory-reset [preserve-image] [module module_number]
```

- **preserve-image**: Specify this flag to retain the NX-OS image in the bootflash of the switch. If you do not specify this flag, the NX-OS image will be erased as well and the switch boots to the loader prompt.
- **module module\_number**: For modular switch line cards and fabric modules, you must specify the number of the module on which to perform the secure erase.

---

For nonmodular switches, the decommission process takes from 2 to 8 hours, depending on the switch and SSD type. The process securely erases the switch and removes the switch configuration from the Cisco Application Policy Infrastructure Controller (APIC). The secure erase process does not remove the NX-OS image from the bootflash. The switch cannot join the fabric until you manually re-register the switch.

The switch reboots after the secure erase operation completes. To connect to the switch, you must use the terminal console because the IP address is not reachable.

For modular switch line cards or fabric modules, the decommission process takes from 30 minutes to 2 hours, depending on the switch and SSD type. The process securely erases the data from the module of the switch and removes the module's configuration from the Cisco APIC. The process does not remove the NX-OS image from the bootflash.

After the secure erase operation completes, the module will be in the powered-down state. To connect to the switch, you must use the terminal console because the IP address is not reachable.





## CHAPTER 5

# Cisco APIC Cluster Management

---

- [APIC Cluster Overview, on page 69](#)
- [Expanding the Cisco APIC Cluster, on page 69](#)
- [Contracting the Cisco APIC Cluster, on page 70](#)
- [Cluster Management Guidelines, on page 70](#)
- [Expanding the APIC Cluster Using the GUI, on page 74](#)
- [Expanding the APIC Cluster Using the Add Node Option, on page 74](#)
- [Contracting the APIC Cluster Using the GUI, on page 77](#)
- [Contracting the APIC Cluster Using the Delete Node Option, on page 77](#)
- [Commissioning and Decommissioning Cisco APIC Controllers, on page 78](#)
- [Shutting Down the APICs in a Cluster, on page 82](#)
- [Cold Standby, on page 83](#)
- [Migration of APICs, on page 85](#)

## APIC Cluster Overview

The Cisco Application Policy Infrastructure Controller (APIC) appliance is deployed in a cluster. A minimum of three controllers are configured in a cluster to provide control of the Cisco ACI fabric. The ultimate size of the controller cluster is directly proportionate to the size of the ACI deployment and is based on transaction-rate requirements. Any controller in the cluster can service any user for any operation, and a controller can be transparently added to or removed from the cluster.

This section provides guidelines and examples related to expanding, contracting, and recovering the APIC cluster.

## Expanding the Cisco APIC Cluster

Expanding the Cisco APIC cluster is the operation to increase any size mismatches, from a cluster size of N to size N+1, within legal boundaries. The operator sets the administrative cluster size and connects the APICs with the appropriate cluster IDs, and the cluster performs the expansion.

During cluster expansion, regardless of in which order you physically connect the APIC controllers, the discovery and expansion takes place sequentially based on the APIC ID numbers. For example, APIC2 is discovered after APIC1, and APIC3 is discovered after APIC2 and so on until you add all the desired APICs to the cluster. As each sequential APIC is discovered, a single data path or multiple data paths are established,

and all the switches along the path join the fabric. The expansion process continues until the operational cluster size reaches the equivalent of the administrative cluster size.

## Contracting the Cisco APIC Cluster

Contracting the Cisco APIC cluster is the operation to decrease any size mismatches, from a cluster size of N to size N -1, within legal boundaries. As the contraction results in increased computational and memory load for the remaining APICs in the cluster, the decommissioned APIC cluster slot becomes unavailable by operator input only.

During cluster contraction, you must begin decommissioning the last APIC in the cluster first and work your way sequentially in reverse order. For example, APIC4 must be decommissioned before APIC3, and APIC3 must be decommissioned before APIC2.

## Cluster Management Guidelines

The Cisco Application Policy Infrastructure Controller (APIC) cluster comprises multiple Cisco APICs that provide operators a unified real time monitoring, diagnostic, and configuration management capability for the Cisco Application Centric Infrastructure (ACI) fabric. To assure optimal system performance, use the following guidelines when making changes to the Cisco APIC cluster:

- Prior to initiating a change to the cluster, always verify its health. When performing planned changes to the cluster, all controllers in the cluster should be healthy. If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding. Also, assure that cluster controllers added to the Cisco APIC are running the same version of firmware as the other controllers in the Cisco APIC cluster.
- We recommend that you have at least 3 active Cisco APICs in a cluster, along with additional standby Cisco APICs. Cisco APIC clusters can have from 3 to 7 active Cisco APICs. Refer to the [Verified Scalability Guide](#) to determine how many active Cisco APICs are required for your deployment.
- Disregard cluster information from Cisco APICs that are not currently in the cluster; they do not provide accurate cluster information.
- Cluster slots contain a Cisco APIC `ChassisID`. Once you configure a slot, it remains unavailable until you decommission the Cisco APIC with the assigned `ChassisID`.
- If a Cisco APIC firmware upgrade is in progress, wait for it to complete and the cluster to be fully fit before proceeding with any other changes to the cluster.
- When moving a Cisco APIC, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to shut down. After the Cisco APIC has shut down, move the Cisco APIC, re-connect it, and then turn it back on. From the GUI, verify that the all controllers in the cluster return to a fully fit state.




---

**Note** Only move one Cisco APIC at a time.

---

- When moving a Cisco APIC that is connected to a set of leaf switches to another set of leaf switches or when moving a Cisco APIC to different port within the same leaf switch, first ensure that you have a



healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to move and decommission it from the cluster. After the Cisco APIC is decommissioned, move the Cisco APIC and then commission it.

- Before configuring the Cisco APIC cluster, ensure that all of the Cisco APICs are running the same firmware version. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.
- Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.
- When you decommission a Cisco APIC, the Cisco APIC loses all fault, event, and audit log history that was stored in it. If you replace all Cisco APICs, you lose all log history. Before you migrate a Cisco APIC, we recommend that you manually backup the log history.

## Expanding the APIC Cluster Size

Follow these guidelines to expand the APIC cluster size:

- Schedule the cluster expansion at a time when the demands of the fabric workload will not be impacted by the cluster expansion.
- If one or more of the APIC controllers' health status in the cluster is not "fully fit", remedy that situation before proceeding.
- Stage the new APIC controller(s) according to the instructions in their hardware installation guide. Verify in-band connectivity with a PING test.
- Increase the cluster target size to be equal to the existing cluster size controller count plus the new controller count. For example, if the existing cluster size controller count is 3 and you are adding 3 controllers, set the new cluster target size to 6. The cluster proceeds to sequentially increase its size one controller at a time until all new the controllers are included in the cluster.



---

**Note** Cluster expansion stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster expansion.

---

- Depending on the amount of data the APIC must synchronize upon the addition of each appliance, the time required to complete the expansion could be more than 10 minutes per appliance. Upon successful expansion of the cluster, the APIC operational size and the target size will be equal.



---

**Note** Allow the APIC to complete the cluster expansion before making additional changes to the cluster.

---

## Reducing the APIC Cluster Size

Follow these guidelines to reduce the Cisco Application Policy Infrastructure Controller (APIC) cluster size and decommission the Cisco APICs that are removed from the cluster:




---

**Note** Failure to follow an orderly process to decommission and power down Cisco APICs from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized Cisco APICs to remain connected to the fabric.

---

- Reducing the cluster size increases the load on the remaining Cisco APICs. Schedule the Cisco APIC size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC one by one until the cluster reaches the new lower target size.

Upon the decommissioning and removal of each controller, the Cisco APIC synchronizes the cluster.




---

**Note** After decommissioning a Cisco APIC from the cluster, promptly power it down and disconnect it from the fabric to prevent its rediscovery. Before returning it to service, do a wiped clean back to factory reset.

If the disconnection is delayed and a decommissioned controller is rediscovered, follow these steps to remove it:

1. Power down the Cisco APIC and disconnect it from the fabric.
  2. In the list of Unauthorized Controllers, reject the controller.
  3. Erase the controller from the GUI.
- 

- Cluster synchronization stops if an existing Cisco APIC becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the Cisco APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.




---

**Note** Complete the entire necessary decommissioning steps, allowing the Cisco APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

---

## Replacing Cisco APIC Controllers in the Cluster

Follow these guidelines to replace Cisco APIC controllers:

- If the health status of any Cisco APIC controller in the cluster is not **Fully Fit**, remedy the situation before proceeding.
- Schedule the Cisco APIC controller replacement at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- Make note of the initial provisioning parameters and image used on the Cisco APIC controller that will be replaced. The same parameters and image must be used with the replacement controller. The Cisco APIC proceeds to synchronize the replacement controller with the cluster.



---

**Note** Cluster synchronization stops if an existing Cisco APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.

---

- You must choose a Cisco APIC controller that is within the cluster and not the controller that is being decommissioned. For example: Log in to Cisco APIC1 or APIC2 to invoke the shutdown of APIC3 and decommission APIC3.
- Perform the replacement procedure in the following order:
  1. Make note of the configuration parameters and image of the APIC being replaced.
  2. Decommission the APIC you want to replace (see [Decommissioning a Cisco APIC in the Cluster Using the GUI, on page 80](#))
  3. Commission the replacement APIC using the same configuration and image of the APIC being replaced (see [Commissioning a Cisco APIC in the Cluster Using the GUI, on page 78](#))
- Stage the replacement Cisco APIC controller according to the instructions in its hardware installation guide. Verify in-band connectivity with a PING test.



---

**Note** Failure to decommission Cisco APIC controllers before attempting their replacement will preclude the cluster from absorbing the replacement controllers. Also, before returning a decommissioned Cisco APIC controller to service, do a wiped clean back to factory reset.

---

- Depending on the amount of data the Cisco APIC must synchronize upon the replacement of a controller, the time required to complete the replacement could be more than 10 minutes per replacement controller. Upon successful synchronization of the replacement controller with the cluster, the Cisco APIC operational size and the target size will remain unchanged.



---

**Note** Allow the Cisco APIC to complete the cluster synchronization before making additional changes to the cluster.

---

- The UUID and fabric domain name persist in a Cisco APIC controller across reboots. However, a clean back-to-factory reboot removes this information. If a Cisco APIC controller is to be moved from one fabric to another, a clean back-to-factory reboot must be done before attempting to add such a controller to a different Cisco ACI fabric.

## Expanding the APIC Cluster Using the GUI

This procedure adds one or more APICs to an existing cluster. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For expanding a cluster in release 6.0(2), you can use the **Add Node** option, as detailed in the subsequent procedure.

### Before you begin

You must first set up any Cisco APIC that you will add to the cluster. For information about setting up a Cisco APIC, see [Setting up the Cisco APIC](#), on page 5.

- 
- Step 1** On the menu bar, choose **System > Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers > apic\_name > Cluster as Seen by Node**.  
For *apic\_name*, you must choose a Cisco APIC that is within the cluster that you wish to expand.  
The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.
- Step 3** Verify that the health state of the cluster is **Fully Fit** before you proceed with contracting the cluster.
- Step 4** In the **Work** pane, click **Actions > Change Cluster Size**.
- Step 5** In the **Change Cluster Size** dialog box, in the **Target Cluster Administrative Size** field, choose the target number to which you want to expand the cluster. Click **Submit**.
- Note** You cannot have a cluster size of two Cisco APICs. You can have a cluster of one, three, or more Cisco APICs.
- Step 6** In the **Confirmation** dialog box, click **Yes**.  
In the **Work** pane, under **Properties**, the **Target Size** field must display your target cluster size.
- Step 7** Physically connect all the Cisco APICs that are being added to the cluster.  
In the **Work** pane, in the **Cluster > Controllers** area, the Cisco APICs are added one by one and displayed in the sequential order starting with  $N + 1$  and continuing until the target cluster size is achieved.
- Step 8** Verify that the Cisco APICs are in operational state, and the health state of each controller is **Fully Fit**.
- 

## Expanding the APIC Cluster Using the Add Node Option

Use this procedure on an existing Cisco Application Policy Infrastructure Controller (APIC) cluster to expand the cluster using the **Add Node** option, which was introduced in Cisco APIC release 6.0(2). To expand a cluster in Cisco APIC releases prior to 6.0(2), see the previous procedure.

The **Add Node** option is a simpler and direct method to add a Cisco APIC to a cluster.

### Before you begin

- Ensure the node *to-be-added* is a *clean* node or is in *factory-reset* state.
- Check the current **Cluster Size** in the **General** pane. If it is  $N$ , after successful node addition, the size will be  $N+1$ .

**Step 1** On the menu bar, choose **System > Controllers**. In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**.

**Step 2** In the **Active Controllers** pane, click the **Actions** button and select the **Add Node** option.  
The **Add Node** screen is displayed.

**Step 3** Enter the following details in the **Add Node** screen:  
Select the **Controller Type**. Based on your selection, proceed to the relevant substep.  
Put a check in the **Enabled** box if you need to support IPv6 addresses.

a) When the Controller Type is **Physical**:

- CIMC details pane
  - IP Address: Enter the CIMC IP address.
  - Username: Enter the username to access CIMC.
  - Password: Enter the password to access CIMC.
  - Click **Validate**. *Validation success* is displayed on successful authentication.

This pane appears only if you configured CIMC. If you did not configure CIMC, instead perform the physical APIC login step of the [Bringing up the Cisco APIC Cluster Using the GUI](#), on page 14 procedure (step 1b) on the new node to configure out-of-band management.

- General pane
  - Name: Enter a name for the controller.
  - Admin Password: Enter the admin password for the controller.
  - Controller ID: This is auto-populated based on the existing cluster size. If the current cluster size is  $N$ , the controller ID is displayed as  $N+1$ .
  - Serial Number: This is auto-populated after CIMC validation.
  - Force Add: Put a check in the **Enabled** box to add a Cisco APIC that has a release earlier than 6.0(2).
- Out of Band Network pane
  - IPv4 Address: The address is auto-populated.
  - IPv4 Gateway: The gateway address is auto-populated.

**Note** If you put a check in the **Enabled** box for IPv6 earlier, enter the IPv6 address and gateway.

- Infra Network pane
  - IPv4 Address: Enter the infra network IP address.
  - IPv4 Gateway: Enter the infra network IP address of the gateway.
  - VLAN: Enter a VLAN ID.

b) When the Controller Type is **Virtual**:

- Management IP pane
  - IP Address: Enter the management IP address.
    - Note** The management IP addresses are defined during the deployment of the virtual machines using ESXi/AWS.
  - Enter the username for the virtual APIC.
  - Enter the password for the virtual APIC.
  - Click **Validate**. *Validation success* is displayed on successful authentication.
- General pane
  - Name: User-defined name for the controller.
  - Controller ID: This is auto-populated based on the existing cluster size. If the current cluster size is  $N$ , the controller ID is displayed as  $N+1$ .
  - Serial Number: The serial number of the virtual machine is auto-populated.
  - Force Add: Put a check in the **Enabled** box to add a Cisco APIC that has a release earlier than 6.0(2).
- Out of Band Network pane
  - IPv4 Address: The IP address is auto-populated.
  - IPv4 Gateway: The gateway IP address is auto-populated.

**Note** If you put a check in the **Enabled** box for IPv6 earlier, enter the IPv6 address and gateway.
- Infra Network pane
  - IPv4 Address: Enter the infra network address.
  - IPv4 Gateway: Enter the IP address of the gateway.
  - VLAN: (Applicable only for *remotely attached* virtual APIC- ESXi) Enter the interface VLAN ID to be used.

**Note** The Infra L3 Network pane is not displayed when the virtual APIC is deployed using AWS.

**Step 4** Click **Apply**.

---

**What to do next**

The newly added controller appears in the **Unauthorized Controllers** pane. Wait for a few minutes for the latest controller to appear with the other controllers of the cluster, under the **Active Controllers** pane.

Also, check the **Current Size** and the **Target Size** in the **General** pane. The number displayed is updated with the latest node addition.

## Contracting the APIC Cluster Using the GUI

This procedure reduces the cluster size. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For contracting a cluster in release 6.0(2), you can use the **Delete Node** option as detailed in the subsequent procedure.

---

**Step 1** On the menu bar, choose **System > Controllers**. In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**.

You must choose an **apic\_name** that is within the cluster and not the controller that is being decommissioned.

The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.

**Step 2** Verify that the health state of the cluster is **Fully Fit** before you proceed with contracting the cluster.

**Step 3** In the **Work** pane, click **Actions > Change Cluster Size**.

**Step 4** In the **Change Cluster Size** dialog box, in the **Target Cluster Administrative Size** field, choose the target number to which you want to contract the cluster. Click **Submit**.

**Note** It is not acceptable to have a cluster size of two APICs. A cluster of one, three, or more APICs is acceptable.

**Step 5** From the **Active Controllers** area of the **Work** pane, choose the APIC that is last in the cluster.

**Example:**

In a cluster of three, the last in the cluster is three as identified by the controller ID.

**Step 6** Right-click on the controller you want to decommission and choose **Decommission**. When the **Confirmation** dialog box displays, click **Yes**.

The decommissioned controller displays **Unregistered** in the **Operational State** column. The controller is then taken out of service and not visible in the **Work** pane any longer.

**Step 7** Repeat the earlier step to decommission the controllers one by one for all the APICs in the cluster in the appropriate order of highest controller ID number to the lowest.

**Note** The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered, and the controller is no longer in service in the cluster.

You should be left with the remaining controllers in the APIC cluster that you desire.

---

## Contracting the APIC Cluster Using the Delete Node Option

Use this procedure to contract a cluster using the **Delete Node** option which has been introduced in Cisco APIC release 6.0(2). To contract a cluster in APIC releases prior to 6.0(2), see the previous procedure.

You can use this procedure to delete one or more than one node from an APIC cluster.

The **Delete Node** option includes two operations— reduces the cluster size, and decommissions the node.




---

**Note** A two-node cluster is not supported. You cannot delete one node from a three-node cluster. The minimum recommended cluster size is three.

---

**Step 1** On the menu bar, choose **System > Controllers**. In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**.

**Step 2** In the **Active Controllers** pane, select the controller which you want to delete by selecting the required check-box..

**Step 3** Click the **Actions** button and select the **Delete Node** option.

**Step 4** Click **OK** on the pop-up screen to confirm the deletion.

Selecting the force option has no effect. It is a *no operation* option, as it is not supported on Cisco APIC release 6.0(2).

**Note** You need to delete the nodes in decreasing order, that is, for example, you can not delete a node with ID 5 before deleting the node with ID 6.

Check the **Current Size** and **Target Size** in the General pane. The size indicated will be one lesser than it was before. If the earlier cluster size was  $N$ , now it will be  $N-1$ .

**Note** If you are deleting more than node from the cluster, the last node of the cluster is deleted first, followed by the other nodes. **Shrink In Progress** in the **General** pane is set to *Yes* until all the selected nodes are deleted.

---

#### What to do next

- After deleting an APIC from the cluster, power the controller down and disconnect it from the fabric.
- Wait for a few minutes, and confirm that the **Health State** of the remaining nodes of the cluster is displayed as *Fully fit* before further action.

# Commissioning and Decommissioning Cisco APIC Controllers

## Commissioning a Cisco APIC in the Cluster Using the GUI

Use this procedure for commissioning an APIC. This procedure is applicable for releases prior to Cisco APIC release 6.0(2). For release 6.0(2), the commissioning workflow has changed, please see the subsequent section for details.

---

**Step 1** From the menu bar, choose **System > Controllers**.

**Step 2** In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**. The **Cluster as Seen by Node** window appears in the **Work** pane with the **APIC Cluster** and **Standby APIC** tabs. In the **APIC Cluster** tab, the controller details appear. This includes the current cluster target and current sizes, the administrative, operational, and health states of each controller in the cluster.

**Step 3** From the **APIC Cluster** tab of the **Work** pane, verify in the **Active Controllers** summary table that the cluster **Health State** is **Fully Fit** before continuing.



- Step 4** From the **Work** pane, right-click the decommissioned controller that is displaying **Unregistered** in the **Operational State** column and choose **Commission**.  
The controller is highlighted.
- Step 5** In the **Confirmation** dialog box, click **Yes**.
- Step 6** Verify that the commissioned Cisco APIC is in the operational state and the health state is **Fully Fit**.
- 

## Commissioning a Cisco APIC in the Cluster

Use this procedure on an existing Cisco Application Policy Infrastructure Controller (APIC) cluster for commissioning a Cisco APIC in that cluster. This procedure is applicable for Cisco APIC release 6.0(2). From release 6.0(2), the commissioning workflow has been enhanced. It can be used to provision an existing controller and also for RMA (return material authorization).

---

- Step 1** From the menu bar, choose **System > Controllers**.
- Step 2** In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**.
- Step 3** Select a decommissioned Cisco APIC from the **Active Controllers** table.
- Step 4** In the **Active Controllers** table, click the **Actions** icon (three dots), which is displayed at the end of the row for each Cisco APIC. From the displayed options, click **Commission**.  
The **Commission** dialog box is displayed.
- Step 5** Enter the following details in the **Commission** screen:  
Choose the **Controller Type**. Based on your choice, proceed to the relevant substep.  
Put a check in the **Enabled** check box if you need to support IPv6 addresses.
- a) When the Controller Type is **Physical**:
- CIMC details pane
    - IP Address: Enter the CIMC IP address.
    - Username: Enter the username to access CIMC.
    - Password: Enter the password to access CIMC.
    - Click **Validate**. *Validation success* is displayed on successful authentication.
- This pane appears only if you configured CIMC. If you did not configure CIMC, instead perform the physical APIC login step of the [Bringing up the Cisco APIC Cluster Using the GUI](#), on page 14 procedure (step 1b) on the new node to configure out-of-band management.
- General pane
    - Name: The name of the controller. The name is entered automatically after the CIMC validation.
    - Admin Password: Enter the admin password for the controller.
    - Controller ID: This is auto-populated based on the Cisco APIC that was decommissioned. The ID of the decommissioned node is assigned.
    - Serial Number: This is auto-populated after CIMC validation.

- Pod ID: Enter the ID number of the pod for the Cisco APIC.

- Out of Band Network pane

- IPv4 Address: Enter the IPv4 address of the out-of-band network.
- IPv4 Gateway: Enter the IPv4 gateway address of the out-of-band network.

**Note** If you have selected the **Enabled** check box for IPv6 earlier, enter the IPv6 address and gateway.

b) When the Controller Type is **Virtual** :

- Virtual Instance: Enter the management IP and click **Validate**.

**Note** The management IP addresses are defined during the deployment of the VMs using ESXi/ AWS.

- General pane

- Name: A user-defined name for the controller.
- Controller ID: This is auto-populated based on the Cisco APIC that was decommissioned. The ID of the decommissioned node is assigned.
- Serial Number: The serial number of the VM is auto-populated.

- Out of Band Network pane

- IPv4 Address: The IP address is auto-populated.
- IPv4 Gateway: The gateway IP address is auto-populated.

**Note** If you have selected the **Enabled** check box for IPv6 earlier, enter the IPv6 address and gateway.

- Infra Network pane

- IPv4 Address: Enter the infra network address.
- IPv4 Gateway: Enter the IP address of the gateway.
- VLAN: (applicable only for *remotely attached* virtual APIC- ESXi) enter the interface VLAN ID to be used.

**Note** The Infra L3 Network pane is not displayed when the virtual APIC is deployed using AWS.

**Step 6** Click **Apply**.

**Step 7** Verify that the commissioned Cisco APIC is in the operational state and the health state is *Fully Fit*.

## Decommissioning a Cisco APIC in the Cluster Using the GUI

This procedure decommissions a Cisco Application Policy Infrastructure Controller (APIC) in the cluster. This procedure is applicable for APIC releases prior to Cisco APIC release 6.0(2). To decommission an APIC in release 6.0(2), see the subsequent procedure.



---

**Note** Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.

---

**Step 1** On the menu bar, choose **System > Controllers**.

**Step 2** In the **Navigation** pane, expand **Controllers > apic\_name > Cluster as Seen by Node**.

You must choose an **apic\_name** that is within the cluster and not the controller that is being decommissioned.

The **Cluster as Seen by Node** window appears in the **Work** pane with the controller details and the **APIC Cluster** and **Standby APIC** tabs.

**Step 3** In the **Work** pane, verify in the **APIC Cluster** tab that the **Health State** in the **Active Controllers** summary table indicates the cluster is **Fully Fit** before continuing.

**Step 4** In the **Active Controllers** table located in the **APIC Cluster** tab of the **Work** pane, right-click on the controller you want to decommission and choose **Decommission**.

The **Confirmation** dialog box displays.

**Step 5** Click **Yes**.

The decommissioned controller displays **Unregistered** in the **Operational State** column. The controller is then taken out of service and no longer visible in the **Work** pane.

- Note**
- After decommissioning a Cisco APIC from the cluster, power the controller down and disconnect it from the fabric. Before returning the Cisco APIC to service, perform a factory reset on the controller.
  - The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered, and the controller is no longer in service in the cluster.
  - After decommissioning the Cisco APIC, you must reboot the controller for Layer 4 to Layer 7 services. You must perform the reboot before re-commissioning the controller.

---

## Decommissioning a Cisco APIC in the Cluster

This procedure decommissions a Cisco APIC in the cluster. This procedure is applicable for Cisco APIC release 6.0(2). To decommission a Cisco APIC for releases prior to release 6.0(2), use the previous procedure.



---

**Note** Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.

---

**Step 1** On the menu bar, choose **System > Controllers**.

**Step 2** In the **Navigation** pane, expand **Controllers > apic\_name > Cluster as Seen by Node**.

You must choose an **apic\_name** that is within the cluster and not the controller that is being decommissioned.

The **Cluster as Seen by Node** window appears in the **Work** pane with the controller details.

**Step 3** In the **Work** pane, verify the **Health State** in the **Active Controllers** summary table indicates the cluster is **Fully Fit** before continuing.

**Step 4** In the **Active Controllers** table, click the **Actions** icon (three dots) displayed at the end of the row for each APIC. Select the **Decommission** option.

The **Decommission** dialog box is displayed.

**Step 5** Click **OK**.

The **Enabled** check-box for Force is a *no operation* option, as it is not supported on Cisco APIC release 6.0(2).

The decommissioned controller displays **Unregistered** in the **Operational State** column. The controller is then taken out of service and no longer visible in the **Work** pane.

- Note**
- After decommissioning a Cisco APIC from the cluster, power the controller down and disconnect it from the fabric. Before returning the Cisco APIC to service, perform a factory reset on the controller.
  - The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered, and the controller is no longer in service in the cluster.
  - After decommissioning the Cisco APIC, you must reboot the controller for Layer 4 to Layer 7 services. You must perform the reboot before re-commissioning the controller.

---

## Shutting Down the APICs in a Cluster

### Shutting Down all the APICs in a Cluster

Before you shutdown all the APICs in a cluster, ensure that the APIC cluster is in a healthy state and all the APICs are showing fully fit. Once you start this process, we recommend that no configuration changes are done during this process. Use this procedure to gracefully shut down all the APICs in a cluster.

**Step 1** Log in to Cisco APIC with appliance ID 1.

**Step 2** On the menu bar, choose **System > Controllers**.

**Step 3** In the Navigation pane, expand **Controllers > apic\_controller\_name**.

You must select the third APIC in the cluster.

**Step 4** Right-click the controller and click **Shutdown**.

**Step 5** Repeat the steps to shutdown the second APIC in the cluster.

**Step 6** Log in to Cisco IMC of the first APIC in the cluster to shutdown the APIC.

**Step 7** Choose **Server > Server Summary > Shutdown Server**.

You have now shutdown all the three APICs in a cluster.

---

## Bringing Back the APICs in a Cluster

Use this procedure to bring back the APICs in a cluster.

- 
- Step 1** Log in to Cisco IMC of the first APIC in the cluster.
- Step 2** Choose **Server > Server Summary > Power On** to power on the first APIC.
- Step 3** Repeat the steps to power on the second APIC and then the third APIC in the cluster.

After all the APICs are powered on, ensure that all the APICs are in a fully fit state. Only after verifying that the APICs are in a fully fit state, you must make any configuration changes on the APIC.

---

## Cold Standby

### About Cold Standby for a Cisco APIC Cluster

The Cold Standby functionality for a Cisco Application Policy Infrastructure Controller (APIC) cluster enables you to operate the Cisco APICs in a cluster in an Active/Standby mode. In a Cisco APIC cluster, the designated active Cisco APICs share the load and the designated standby Cisco APICs can act as a replacement for any of the Cisco APICs in the active cluster.

As an admin user, you can set up the Cold Standby functionality when the Cisco APIC is launched for the first time. We recommend that you have at least three active Cisco APICs in a cluster, and one or more standby Cisco APICs. As an admin user, you can initiate the switch over to replace an active Cisco APIC with a standby Cisco APIC.

### Guidelines and Limitations for Standby Cisco APICs

The following are guidelines and limitations for standby Cisco Application Policy Infrastructure Controllers (APICs):

- There must be three active Cisco APICs to add a standby Cisco APIC.
- The standby Cisco APIC need to run with the same firmware version of the cluster when the standby Cisco APICs join the cluster during the initial setup.
- During an upgrade process, after all the active Cisco APICs are upgraded, the standby Cisco APICs are also upgraded automatically.
- During the initial setup, IDs are assigned to the standby Cisco APICs. After a standby Cisco APIC is switched over to an active Cisco APIC, the standby Cisco APIC (new active) starts using the ID of the replaced (old active) Cisco APIC.
- The admin login is not enabled on the standby Cisco APICs. To troubleshoot a Cold Standby Cisco APIC, you must log in to the standby using SSH as *rescue-user*.
- During the switch over, the replaced active Cisco APIC needs to be powered down to prevent connectivity to the replaced Cisco APIC.

- Switch over fails under the following conditions:
  - If there is no connectivity to the standby Cisco APIC.
  - If the firmware version of the standby Cisco APIC is not the same as that of the active cluster.
- After switching over a standby Cisco APIC to be active, you can setup another standby Cisco APIC, if needed.
- If **Retain OOB IP address for Standby (new active)** is checked, the standby (new active) Cisco APIC will retain its original standby out-of-band management IP address.
- If **Retain OOB IP address for Standby (new active)** is not checked:
  - If only one active Cisco APIC is down: The standby (new active) Cisco APIC will use the old active Cisco APIC's out-of-band management IP address.
  - If more than one active Cisco APICs are down: The standby (new active) Cisco APIC will try to use the active Cisco APIC's out-of-band management IP address, but it may fail if the shard of out-of-band management IP address configuration for the active Cisco APIC is in the minority state.
- For Cisco ACI Multi-Pod, if the old active Cisco APIC and the standby Cisco APIC use different out-of-band management IP subnets, you must check the option to have the standby (new active) Cisco APIC retain its original standby out-of-band management IP address. Otherwise, you will lose out-of-band management IP connectivity to the standby (new active) Cisco APIC. This situation might happen if the old active Cisco APIC and the standby Cisco APIC are in the different pods.
 

If out-of-band management IP connectivity is lost because of this reason or if more than one active Cisco APICs are down, you must create a new Static Node Management OOB IP Address to change the new active (previously standby) Cisco APIC out-of-band management IP address. You must have the cluster out of the minority state to make the configuration change.
- The standby Cisco APIC does not participate in policy configuration or management.
- No information is replicated to the standby Cisco APICs, not even the administrator credentials.
- A standby Cisco APIC does not retain the in-band management IP address when you promote the Cisco APIC to be active. You must manually reconfigure the Cisco APIC to have the correct in-band management IP address.

## Verifying Cold Standby Status Using the GUI

1. On the menu bar, choose **System > Controllers**.
2. In the **Navigation** pane, expand **Controllers > apic\_controller\_name > Cluster as Seen by Node**.
3. In the **Work** pane, the standby controllers are displayed under **Standby Controllers**.

## Switching Over an Active APIC with a Standby APIC Using the GUI

Use this procedure to switch over an active APIC with a standby APIC.

- 
- Step 1** On the menu bar, choose **System > Controllers**.

- Step 2** In the **Navigation** pane, expand **Controllers** > **apic\_controller\_name** > **Cluster as Seen by Node**.  
The *apic\_controller\_name* should be other than the name of the controller being replaced.
- Step 3** In the **Work** pane, verify that the **Health State** in the **Active Controllers** summary table indicates the active controllers other than the one being replaced are **Fully Fit** before continuing.
- Step 4** Click an *apic\_controller\_name* that you want to switch over.
- Step 5** In the **Work** pane, click **Actions** > **Replace**.  
The **Replace** dialog box displays.
- Step 6** Choose the **Backup Controller** from the drop-down list and click **Submit**.  
It may take several minutes to switch over an active APIC with a standby APIC and for the system to be registered as active.
- Step 7** Verify the progress of the switch over in the **Failover Status** field in the **Active Controllers** summary table.
- 

## Migration of APICs

Beginning with Cisco APIC release 6.1(1), migration from a physical APIC cluster to a virtual APIC cluster deployed on an ESXi host (using VMware vCenter) is supported. Migration from a virtual APIC cluster (on an ESXi host) to a physical APIC cluster is also supported.

## Guidelines and Limitations

Following are the guidelines and limitations for migrating physical APICs to virtual APICs (and vice-versa):

- Physical APICs in layer 2 (directly attached to the fabric) can be migrated to layer 2 virtual APICs and layer 2 virtual APICs can be migrated to layer 2 physical APICs. Physical APICs in layer 3 (remotely attached to the fabric) can be migrated to layer 3 virtual APICs. Virtual APICs in layer 3 can be migrated to layer 3 physical APIC. Migration from a layer 2 APIC to a layer 3 APIC (or vice-versa) is not supported.
- No support for standby node migration. Before migration, remove standby node(s) from the cluster and then migrate.
- No support for migration for mini ACI fabric.
- If upgrade is in progress, do not initiate the migration process.
- If migration is in progress, do not initiate an upgrade.
- If NDO is configured, you must update the connection details on NDO as the migration changes the OOB IP and subnet addresses.
- If SMU is installed on the physical APIC, for Cisco APIC release 6.1(1) the migration (from physical to virtual APIC) is not recommended. You need to upgrade the cluster to an image which has the fix for the SMU before proceeding with migration.
- For app-infra, stop any running job for ELAM/FTRIAGE before migration and re-start after migration is complete.
- Any configuration which uses APIC OOB needs to be updated after the migration process is completed.

## How the migration process works

In this section, a high-level explanation of the migration process is provided. For the detailed steps, see the [Migrate a Physical APIC Cluster to a Virtual APIC Cluster \(or Virtual APIC Cluster to a Physical APIC Cluster\)](#) procedure in the subsequent section.

Consider a three-node cluster; three source nodes and correspondingly three target nodes (after migration nodes). The APIC with controller ID 1 is considered as APIC 1. Login to APIC 1 (IP address 172.16.1.1) and initiate the migration process.

| APIC   | Source Node | Target Node |
|--------|-------------|-------------|
| APIC 1 | 172.16.1.1  | 172.16.1.11 |
| APIC 2 | 172.16.1.2  | 172.16.1.12 |
| APIC 3 | 17.16.1.3   | 172.16.1.13 |

As soon as the migration is initiated, APIC 1 starts with the migration of APIC 3, followed by APIC 2. After source APIC 2 (IP address 172.16.1.2) is migrated to the target APIC 2 node (172.16.1.12), the target APIC 2 node takes control to enable the migration of APIC 1. This is called the handover process where in the control is passed on from source APIC 1 (172.16.1.1) to target APIC 2 (172.16.1.12). At this stage, a new window is displayed (URL redirect to target APIC 2). This is because, after successful migration, source APIC 1 is no longer part of the cluster (which now has the migrated target APICs).

The migration proceeds in the reverse order, that is, APIC  $N$  (APIC 3 in the example) is migrated first, followed by APIC  $N-1$  (APIC 2 in the example), so on, and then finally APIC 1.

## Migrate a Physical APIC Cluster to a Virtual APIC Cluster (or Virtual APIC Cluster to a Physical APIC Cluster)

Use this procedure to migrate the nodes of a physical APIC cluster to a virtual APIC cluster (or vice-versa).

### Before you begin

Following are the required prerequisites before you start with the migration process:

#### Cluster health

Confirm that the current APIC cluster is *Fully fit*.

#### Generic

- Ensure that the source and destination APICs' date and time are synchronized.
- Ensure that all the controllers are on Cisco APIC release 6.1(1), and all the switches are running the same version as the controller.

#### Source and Target Nodes



- For directly connected APIC migration, ensure both source and target nodes are on the same layer 2 network.
- For remotely connected APIC migration, ensure both source and target nodes have infra network connectivity between them. This means the new target APIC should have the correct IPN configuration such that it can interact with the infra network of the fabric.
- Target nodes have the same admin password as the source cluster.
- Target nodes' OOB IP address should be different while all other fields can be same or different from the source node. Infra addresses will remain same for layer 2 (directly attached); for layer 3 (remotely attached) cluster, they can be same or different based on deployments.
- Source cluster and target cluster OOB networking stacks should match. For example, if source cluster is using dual stack (IPv4 and IPv6) for OOB, dual stack (IPv4 and IPv6) address details should be provided for target nodes too.
- Ensure OOB connectivity between the source and destination APICs.
- Ensure the OOB contracts and reachability for the new APIC are configured correctly; the migration process uses the OOB IP address to communicate between the APICs.

#### For virtual APIC to physical APIC migration

- Ensure the physical APIC nodes are factory reset; use the **acdiag touch setup** and **acdiag reboot** commands.
- For migration with/ without CIMC (applicable for physical APICs):

| If...          | Then....                                                                                                                                |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Using CIMC     | ensure that the physical APIC CIMC addresses are reachable from the OOB network of the virtual APIC.                                    |
| Not using CIMC | ensure that the OOB IP address is configured manually on the physical APIC after factory-reset and use the OOB option for connectivity. |

#### For Physical to virtual migration

- Ensure that you have deployed the virtual APIC nodes as per the procedure in the [Deploying Cisco Virtual APIC Using VMware vCenter](#) guide.
- If virtual APICs are deployed on a vCenter that is part of a VMM domain, ensure that Infrastructure VLAN is enabled on the AEP configured on the interfaces connected to the ESXi host(s) where the virtual APIC is deployed.

---

**Step 1** On the **Cluster as Seen by Node** screen, click **Migrate** (displayed in the **Cluster Overview** area).

All the available controllers in the cluster are displayed.

**Note** The **Migrate** button is displayed only on APIC 1 (of the cluster).

**Step 2** Click the pencil icon next to the Validate column, to start the migration process of the selected controller.

The **Migrate Node** screen is displayed.

**Step 3** Enter the following details in the **Migrate Node** screen:

- a) For the **Controller Type**, select Virtual or Physical, as the case may be (migration from physical APIC to virtual APIC and vice-versa is supported).
- b) For the **Connectivity Type**, select OOB, if you are migrating a physical APIC to a virtual APIC. If you are migrating a virtual APIC to a physical APIC, you can either select the OOB option or the CIMC option.

It is recommended to select the CIMC option for the virtual to physical migration. To use the OOB option, connect to the CIMC address of physical APICs and configure the OOB IP addresses manually before starting the migration process.

**Controller Type** and **Connectivity Type** are auto selected based on the source controller type. If required, you can modify them.

- c) In the Management IP pane, enter the following (target APIC details) — Management IP address, Username, Password.  
or

(applicable only for virtual to physical APIC migration) In the CIMC Details pane, enter the following details of the physical APIC— CIMC IP address, username and password of the node.

- d) Click **Validate**.

After you click Validate, the details displayed in the General and Out of Band management panes change to match the details of the controller. The only editable fields are the Name and Pod ID (applicable only for layer 2), the other fields cannot be modified. For virtual to physical APIC migration, confirm the Admin Password too.

**Note** If dual stack is supported, fill in the IPv4 and IPv6 addresses.

- e) In the Infra Network pane (applicable only for Layer 3, APIC is remotely attached to the fabric), enter the following:
  - IPv4 Address: the infra network address.
  - IPv4 Gateway: the IP address of the gateway.
  - VLAN: the interface VLAN ID to be used.

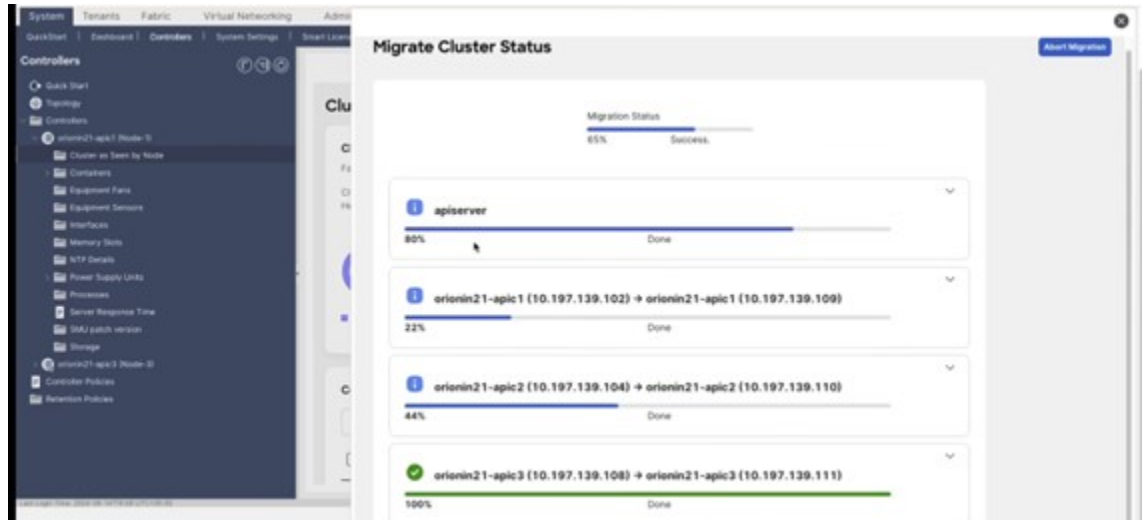
The OOB gateway and IP addresses are auto-populated in the table (based on the validation); click **Apply**. The validation status is displayed as Complete (on the Migrate Nodes screen).

Repeat the same process for the other APICs in the cluster by clicking the pencil icon (next to the Validation column). After providing all the controller details, click the **Migrate** button at the bottom of the Migrate screen.

## Checking the Migration Status

Click **Check Migration Status** to check the status of the migration. The **Check Type** section has details of the completed steps. The migration process involves a series of activities, and this is displayed in stages and each stage is indicated with a color-coded bar.


Figure 13: Migration Status



As displayed in the image, the overall migration status is displayed, followed by the status of the `apiserver`. The `apiserver` is the process that orchestrates the whole migration process. Below the `apiserver` status, the controller-wise migration status is displayed. The source IP address and the target IP address of the nodes are also indicated.

The `apiserver` status is indicated as 100% done (green bar) after the handover to APIC 2 is completed. At this stage, a new window is displayed (URL redirected to target APIC 2). Login to target APIC 2. A banner indicating that *Migration is in progress* is displayed at the top of the GUI until the migration is complete. After the handover process, the banner which was displayed on source APIC 1 is displayed on the target APIC 2. Click the **View Status** link on the banner to check the migration status.

Figure 14: Migration Status Banner



You can also abort the migration process from source APIC 1 by clicking the **Abort** button available on the Migrate Cluster Status screen. The **Abort** button is displayed only after a certain period of time after initiating the migration.

After successful migration:

- the migration status is no longer displayed. If the migration has failed, then a failure message is explicitly displayed.
- to confirm if the target cluster is healthy and fully fit, navigate to **System > Controllers > Expand Controllers**. Expand **Controller 1 > Cluster as seen by Node** page.
- verify if all the fabric nodes are in active state; navigate to **Fabric > Fabric Membership**.
- if the Pod ID of the target APIC has changed, inband address for the node needs to be reconfigured on the Tenant Management screen. Navigate to **Tenants > Mgmt > Node Management Addresses** page.

## Operations in case of Migration Failure

A migration process may be interrupted due to a failure, or you may choose to abort the migration. In case the migration is not successful, it is recommended to revert or resume migration to either the source or target controller type. It is not recommended to have a APIC cluster in a migration failed state with a mix of physical and virtual controllers. Before attempting revert or resume, follow the steps in the next section, [Basic Troubleshooting](#), to get the cluster to a healthy state.

If you choose to resume— migration process is continued. On the **Migrate Node** screen (source APIC 1):

1. Enter the details of all the target nodes based on the controller type you want to migrate to.
2. Click **Migrate**.

If you choose to revert— migration process is restarted. Migration process needs to be restarted after getting the controllers of the cluster to the initial (source) IP addresses.

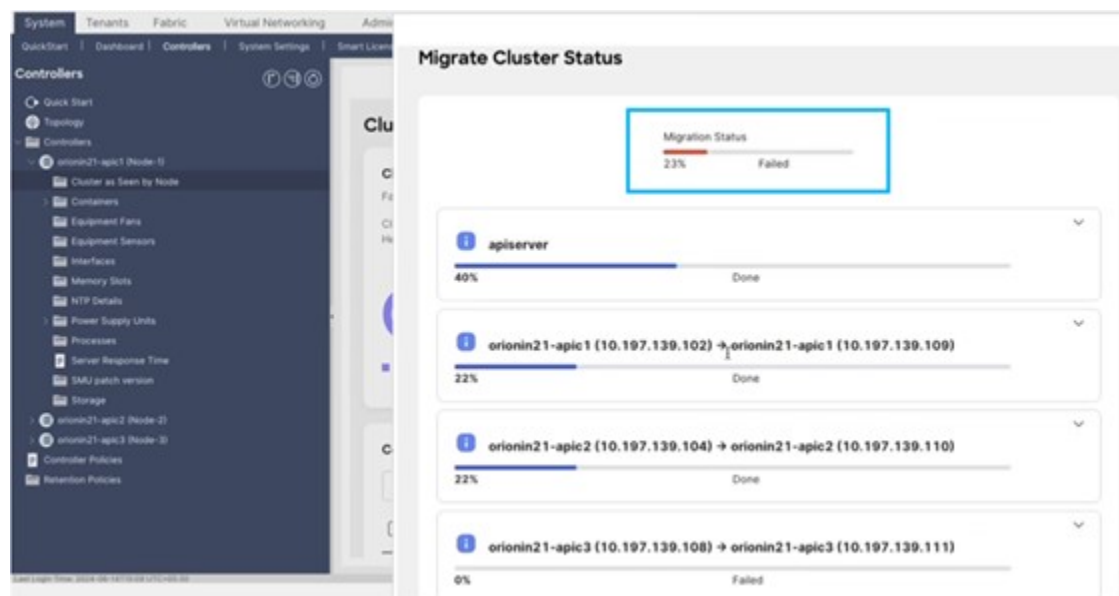
1. Factory reset each of the source APIC nodes which are being migrated using **acdiag touch setup** and **acdiag reboot** commands.
2. On the **Migrate Node** screen, enter the source APIC details for all the nodes, as the migration process reverts the previously migrated APICs to the source controller type.
3. Click **Migrate**.



**Note** If the migration process fails after the handover process (control is passed on to target APIC 2 from source APIC 1), the migration cannot be resumed or reverted.

As mentioned earlier, the various sub-stages of migration with their completion progress are indicated as bars, for each controller. In case of failure at any stage, collect the relevant tech-support details and contact Cisco TAC for further assistance. To collect the logs for tech support, navigate to **Admin > Import/Export > Export Policies > On-demand Tech Support > migration\_techsupport**.

**Figure 15: Migration Status - Failed**



## Basic Troubleshooting

Consider a three-node cluster where in two nodes have migrated successfully, and a failure is detected during the migration of the third node. Check the status of the failed node. If the controller is not in the *Fully fit* state, the migration could fail.

The procedure to get the cluster to a healthy state:

- 
- Step 1** (for migration failures with APIC 1) Check the cluster health from target APIC 2 by navigating to, **System > Controllers**. Select **Controller 2 > Cluster as seen by Node** screen.
- or
- (for migration failures with APIC 2 to N) Check the cluster health from source APIC 1 by navigating to, **System > Controllers**. Select **Controller 1 > Cluster as seen by Node** screen.
- Step 2** If APIC 1 (or any other node of the cluster) is not *Fully fit*, click the three dots adjacent to the serial number of the controller. Select **Maintenance > Decommission**. Click **Force Decommission** as the node is not in a *Fully fit* state. Connect to the source APIC node N using SSH and factory-reset the node using the following commands - **acdiag touch setup, acdiag reboot**.
- Step 3** From source APIC 1 navigate to, **System > Controllers**. Click **Controllers > Controller 1 > Cluster as seen by Node** screen.
- or
- From target APIC 2 navigate to, **System > Controllers**. Click **Controllers > Controller 2 > Cluster as seen by Node** screen.
- Step 4** To commission a controller, click the three dots adjacent to the serial number of the controller. Select **Maintenance > Commission**. Enter the details, as required. Refer to the *Commissioning a Node* procedure (described earlier in this chapter). The only difference here is that the controller ID is pre-populated with the number corresponding to the ID of the controller in the cluster.
- After the controller is commissioned, the cluster is indicated as *Fully fit*.
- Step 5** Check the status of the cluster after commissioning the failed node. If the cluster is in a healthy state, resume the migration by clicking **Migrate** on the **Cluster As Seen By Node** screen. If the migration fails again, contact Cisco TAC for further assistance.
-



## APPENDIX **A**

# Configuring the Cisco APIC Using the CLI

- [Cluster Management Guidelines](#), on page 93
- [Replacing a Cisco APIC in a Cluster Using the CLI](#), on page 94
- [Reducing the APIC Cluster Size](#), on page 95
- [Contracting the Cisco APIC Cluster](#), on page 96
- [Switching Over Active APIC with Standby APIC Using CLI](#), on page 97
- [Verifying Cold Standby Status Using the CLI](#), on page 97
- [Registering an Unregistered Switch Using the CLI](#), on page 98
- [Adding a Switch Before Discovery Using the CLI](#), on page 98
- [Removing a Switch to Maintenance Mode Using the CLI](#), on page 98
- [Inserting a Switch to Operation Mode Using the CLI](#), on page 99
- [Configuring a Remote Location Using the NX-OS Style CLI](#), on page 99
- [Finding Your Switch Inventory Using the NX-OS CLI](#), on page 100
- [Verifying the Cisco APIC Cluster Using the CLI](#), on page 102

## Cluster Management Guidelines

The Cisco Application Policy Infrastructure Controller (APIC) cluster comprises multiple Cisco APICs that provide operators a unified real time monitoring, diagnostic, and configuration management capability for the Cisco Application Centric Infrastructure (ACI) fabric. To assure optimal system performance, use the following guidelines when making changes to the Cisco APIC cluster:

- Prior to initiating a change to the cluster, always verify its health. When performing planned changes to the cluster, all controllers in the cluster should be healthy. If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding. Also, assure that cluster controllers added to the Cisco APIC are running the same version of firmware as the other controllers in the Cisco APIC cluster.
- We recommend that you have at least 3 active Cisco APICs in a cluster, along with additional standby Cisco APICs. Cisco APIC clusters can have from 3 to 7 active Cisco APICs. Refer to the [Verified Scalability Guide](#) to determine how many active Cisco APICs are required for your deployment.
- Disregard cluster information from Cisco APICs that are not currently in the cluster; they do not provide accurate cluster information.
- Cluster slots contain a Cisco APIC `ChassisID`. Once you configure a slot, it remains unavailable until you decommission the Cisco APIC with the assigned `ChassisID`.

- If a Cisco APIC firmware upgrade is in progress, wait for it to complete and the cluster to be fully fit before proceeding with any other changes to the cluster.
- When moving a Cisco APIC, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to shut down. After the Cisco APIC has shut down, move the Cisco APIC, re-connect it, and then turn it back on. From the GUI, verify that the all controllers in the cluster return to a fully fit state.




---

**Note** Only move one Cisco APIC at a time.

---

- When moving a Cisco APIC that is connected to a set of leaf switches to another set of leaf switches or when moving a Cisco APIC to different port within the same leaf switch, first ensure that you have a healthy cluster. After verifying the health of the Cisco APIC cluster, choose the Cisco APIC that you intend to move and decommission it from the cluster. After the Cisco APIC is decommissioned, move the Cisco APIC and then commission it.
- Before configuring the Cisco APIC cluster, ensure that all of the Cisco APICs are running the same firmware version. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.
- Unlike other objects, log record objects are stored only in one shard of a database on one of the Cisco APICs. These objects get lost forever if you decommission or replace that Cisco APIC.
- When you decommission a Cisco APIC, the Cisco APIC loses all fault, event, and audit log history that was stored in it. If you replace all Cisco APICs, you lose all log history. Before you migrate a Cisco APIC, we recommend that you manually backup the log history.

## Replacing a Cisco APIC in a Cluster Using the CLI



- 
- Note**
- For more information about managing clusters, see [Cluster Management Guidelines, on page 70](#).
  - When you replace an Cisco APIC, the password will always be synced from the cluster. When replacing APIC 1, you will be asked for a password but it will be ignored in favor of the existing password in the cluster. When replacing Cisco APIC 2 or 3, you will not be asked for a password.
- 

### Before you begin

Before replacing a Cisco Application Policy Infrastructure Controller (APIC), ensure that the replacement Cisco APIC is running the same firmware version as the Cisco APIC to be replaced. If the versions are not the same, you must update the firmware of the replacement Cisco APIC before you begin. Initial clustering of Cisco APICs running differing versions is an unsupported operation and may cause problems within the cluster.

- 
- Step 1** Identify the Cisco APIC that you want to replace.
- Step 2** Note the configuration details of the Cisco APIC to be replaced by using the **acdiag avread** command.



**Step 3** Decommission the Cisco APIC using the **decommission controller** *controller-id* command in config mode.

Decommissioning the Cisco APIC removes the mapping between the APIC ID and Chassis ID. The new Cisco APIC typically has a different APIC ID, so you must remove this mapping in order to add a new Cisco APIC to the cluster.

Beginning with Cisco APIC release 6.0(2), an optional argument (*force*) is added to the **decommission** command to allow forcing the decommission operation. The revised command is **decommission controller** *controller-id* [**force**], with the following behaviors:

- When *force* is not declared, the decommission proceeds only if the cluster is not in an unhealthy or upgrade state, where a decommission may not be proper.
- When *force* is declared, the decommission proceeds regardless of the cluster state.

For example, `decommission controller 3 force` decommissions APIC3 regardless of the cluster state.

**Step 4** To commission the new Cisco APIC, follow these steps:

- a) Disconnect the old Cisco APIC from the fabric.
- b) Connect the replacement Cisco APIC to the fabric.

The new Cisco APIC appears in the Cisco APIC GUI menu **System > Controllers > *apic\_controller\_name* > Cluster as Seen by Node** in the **Unauthorized Controllers** list.

- c) Commission the new Cisco APIC using the **controller** *controller-id* **commission** command.
- d) Boot the new Cisco APIC.
- e) Allow several minutes for the new Cisco APIC information to propagate to the rest of the cluster.

The new Cisco APIC appears in the Cisco APIC GUI menu **System > Controllers > *apic\_controller\_name* > Cluster as Seen by Node** in the **Active Controllers** list.

---

#### What to do next

For each decommissioned controller, verify that the operational state of the controller is unregistered and that the controller is no longer in service in the cluster.



---

**Note** If a decommissioned Cisco APIC is not promptly removed from the fabric, it might be rediscovered, which could cause problems. In that case, follow the instructions in [Reducing the APIC Cluster Size, on page 71](#) to remove the controller.

---

## Reducing the APIC Cluster Size

Follow these guidelines to reduce the Cisco Application Policy Infrastructure Controller (APIC) cluster size and decommission the Cisco APICs that are removed from the cluster:




---

**Note** Failure to follow an orderly process to decommission and power down Cisco APICs from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized Cisco APICs to remain connected to the fabric.

---

- Reducing the cluster size increases the load on the remaining Cisco APICs. Schedule the Cisco APIC size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC one by one until the cluster reaches the new lower target size.

Upon the decommissioning and removal of each controller, the Cisco APIC synchronizes the cluster.




---

**Note** After decommissioning a Cisco APIC from the cluster, promptly power it down and disconnect it from the fabric to prevent its rediscovery. Before returning it to service, do a wiped clean back to factory reset.

If the disconnection is delayed and a decommissioned controller is rediscovered, follow these steps to remove it:

1. Power down the Cisco APIC and disconnect it from the fabric.
  2. In the list of Unauthorized Controllers, reject the controller.
  3. Erase the controller from the GUI.
- 

- Cluster synchronization stops if an existing Cisco APIC becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the Cisco APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.




---

**Note** Complete the entire necessary decommissioning steps, allowing the Cisco APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

---

## Contracting the Cisco APIC Cluster

Contracting the Cisco APIC cluster is the operation to decrease any size mismatches, from a cluster size of N to size N -1, within legal boundaries. As the contraction results in increased computational and memory

load for the remaining APICs in the cluster, the decommissioned APIC cluster slot becomes unavailable by operator input only.

During cluster contraction, you must begin decommissioning the last APIC in the cluster first and work your way sequentially in reverse order. For example, APIC4 must be decommissioned before APIC3, and APIC3 must be decommissioned before APIC2.

## Switching Over Active APIC with Standby APIC Using CLI

Use this procedure to switch over an active APIC with a standby APIC.

### Step 1 `replace-controller replace ID number Backup serial number`

Replaces an active APIC with an standby APIC.

#### Example:

```
apic1#replace-controller replace 2 FCH1804V27L
Do you want to replace APIC 2 with a backup? (Y/n): Y
```

### Step 2 `replace-controller reset ID number`

Resets fail over status of the active controller.

#### Example:

```
apic1# replace-controller reset 2
Do you want to reset failover status of APIC 2? (Y/n): Y
```

## Verifying Cold Standby Status Using the CLI

To verify the Cold Standby status of APIC, log in to the APIC as admin and enter the command **show controller**.

```
apic1# show controller
Fabric Name : vegas
Operational Size : 3
Cluster Size : 3
Time Difference : 496
Fabric Security Mode : strict
```

| ID  | Pod | Address<br>Version                    | In-Band IPv4<br>Flags | In-Band IPv6<br>Serial Number | Health | OOB IPv4                  | OOB IPv6 |
|-----|-----|---------------------------------------|-----------------------|-------------------------------|--------|---------------------------|----------|
| 1*  | 1   | 10.0.0.1<br>fe80::26e9:b3ff:fe91:c4e0 | 0.0.0.0<br>2.2(0.172) | fc00::1<br>crva- FCH1748V0DF  |        | 172.23.142.4<br>fully-fit |          |
| 2   | 1   | 10.0.0.2<br>fe80::26e9:bf8f:fe91:f37c | 0.0.0.0<br>2.2(0.172) | fc00::1<br>crva- FCH1747V0YF  |        | 172.23.142.6<br>fully-fit |          |
| 3   | 1   | 10.0.0.3<br>fe80::4e00:82ff:fead:bc66 | 0.0.0.0<br>2.2(0.172) | fc00::1<br>crva- FCH1725V2DK  |        | 172.23.142.8<br>fully-fit |          |
| 21~ |     | 10.0.0.21                             |                       | ----- FCH1734V2DG             |        |                           |          |

Flags - c:Commissioned | r:Registered | v:Valid Certificate | a:Approved | f/s:Failover fail/success  
 (\*)Current (~)Standby

## Registering an Unregistered Switch Using the CLI

Use this procedure to register a switch from the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.



**Note** This procedure is identical to "Adding a Switch Before Discovery Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node exists, the system registers it.

### Procedure

|               | Command or Action                                                                                      | Purpose                                           |
|---------------|--------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| <b>Step 1</b> | <code>[no] system switch-id serial-number switch-id name pod id role leaf node-type tier-2-leaf</code> | Adds the switch to the pending registration list. |

## Adding a Switch Before Discovery Using the CLI

Use this procedure to add a switch to the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the CLI.



**Note** This procedure is identical to "Registering an Unregistered Switch Using the CLI". When you execute the command, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

`[no] system switch-id serial-number switch-id name pod id role leaf node-type tier-2-leaf`

Adds the switch to the pending registration list.

## Removing a Switch to Maintenance Mode Using the CLI

Use this procedure to remove a switch to maintenance mode using the CLI.



**Note** While the switch is in maintenance mode, CLI 'show' commands on the switch show the front panel ports as being in the up state and the BGP protocol as up and running. The interfaces are actually shut and all other adjacencies for BGP are brought down, but the displayed active states allow for debugging.

**[no]debug-switch** *node\_id* or *node\_name*

Removes the switch to maintenance mode.

## Inserting a Switch to Operation Mode Using the CLI

Use this procedure to insert a switch to operational mode using the CLI.

**[no]no debug-switch** *node\_id* or *node\_name*

Inserts the switch to operational mode.

## Configuring a Remote Location Using the NX-OS Style CLI

In the ACI fabric, you can configure one or more remote destinations for exporting techsupport or configuration files.

### SUMMARY STEPS

1. **configure**
2. **[no] remote path** *remote-path-name*
3. **user** *username*
4. **path** {**ftp** | **scp** | **sftp**} *host[:port]* [**remote-directory** ]

### DETAILED STEPS

|        | Command or Action                                                                                                     | Purpose                                      |
|--------|-----------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| Step 1 | <b>configure</b><br><b>Example:</b><br><code>apic1# configure</code>                                                  | Enters global configuration mode.            |
| Step 2 | <b>[no] remote path</b> <i>remote-path-name</i><br><b>Example:</b><br><code>apic1(config)# remote path myFiles</code> | Enters configuration mode for a remote path. |

|               | Command or Action                                                                                                                                                                                                     | Purpose                                                                                  |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <b>Step 3</b> | <b>user</b> <i>username</i><br><b>Example:</b><br>apicl(config-remote)# <b>user admin5</b>                                                                                                                            | Sets the user name for logging in to the remote server. You are prompted for a password. |
| <b>Step 4</b> | <b>path</b> {ftp   scp   sftp} <i>host[:port]</i> [ <b>remote-directory</b> ]<br><b>Example:</b><br>apicl(config-remote)# <b>path sftp</b><br><b>filehost.example.com:21 remote-directory</b><br><b>/reports/apic</b> | Sets the path and protocol to the remote server. You are prompted for a password.        |

### Examples

This example shows how to configure a remote path for exporting files.

```

apicl# configure
apicl(config)# remote path myFiles
apicl(config-remote)# user admin5
You must reset the password when modifying the path:
Password:
Retype password:
apicl(config-remote)# path sftp filehost.example.com:21 remote-directory /reports/apic
You must reset the password when modifying the path:
Password:
Retype password:

```

## Finding Your Switch Inventory Using the NX-OS CLI

This section explains how to find your switch model and serial numbers using the NX-OS CLI.

Find your switch inventory as follows:

### Example:

```

switch# show hardware
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

Software
 BIOS: version 07.56
 kickstart: version 12.1(1h) [build 12.1(1h)]

```

```
system: version 12.1(1h) [build 12.1(1h)]
PE: version 2.1(1h)
BIOS compile time: 06/08/2016
kickstart image file is: /bootflash/aci-n9000-dk9.12.1.1h.bin
kickstart compile time: 10/01/2016 20:10:40 [10/01/2016 20:10:40]
system image file is: /bootflash/auto-s
system compile time: 10/01/2016 20:10:40 [10/01/2016 20:10:40]
```

#### Hardware

```
cisco N9K-C93180YC-EX ("supervisor")
 Intel(R) Xeon(R) CPU @ 1.80GHz with 16400384 kB of memory.
 Processor Board ID FDO20101H1W
```

```
Device name: ifav41-leaf204
bootflash: 62522368 kB
```

Kernel uptime is 02 day(s), 21 hour(s), 42 minute(s), 31 second(s)

Last reset at 241000 usecs after Sun Oct 02 01:27:25 2016

```
Reason: reset-by-installer
System version: 12.1(1e)
Service: Upgrade
```

#### plugin

```
Core Plugin, Ethernet Plugin
```

#### Switch hardware ID information

```
Switch is booted up
Switch type is : Nexus C93180YC-EX Chassis
Model number is N9K-C93180YC-EX
H/W version is 0.2010
Part Number is 73-15298-01
Part Revision is 1
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
CLEI code is 73-15298-01
```

```
Chassis has one slot
```

#### Module1 ok

```
Module type is : 48x10/25G
1 submodules are present
Model number is N9K-C93180YC-EX
H/W version is 0.2110
Part Number is 73-17776-02
Part Revision is 11
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
CLEI code is 73-17776-02
```

#### GEM ok

```
Module type is : 6x40/100G Switch
1 submodules are present
Model number is N9K-C93180YC-EX
H/W version is 0.2110
Part Number is 73-17776-02
Part Revision is 11
Manufacture Date is Year 20 Week 10
Serial number is FDO20101H1W
```

```

CLEI code is 73-17776-02

Chassis has 2 PowerSupply Slots

PS1 shut
Power supply type is : 54.000000W 220v AC
Model number is NXA-PAC-650W-PE
H/W version is 0.0
Part Number is 341-0729-01
Part Revision is A0
Manufacture Date is Year 19 Week 50
Serial number is LIT19500ZEK
CLEI code is 341-0729-01

PS2 ok
Power supply type is : 54.000000W 220v AC
Model number is NXA-PAC-650W-PE
H/W version is 0.0
Part Number is 341-0729-01
Part Revision is A0
Manufacture Date is Year 19 Week 50
Serial number is LIT19500ZEA
CLEI code is 341-0729-01

Chassis has 4 Fans

FT1 ok

Fan1(sys_fan1) (fan_model:NXA-FAN-30CFM-F) is inserted but info
is not available

FT2 ok

Fan2(sys_fan2) (fan_model:NXA-FAN-30CFM-F) is inserted but info
is not available

FT3 ok

Fan3(sys_fan3) (fan_model:NXA-FAN-30CFM-F) is inserted but info
is not available

FT4 ok

Fan4(sys_fan4) (fan_model:NXA-FAN-30CFM-F) is inserted but info
is not available

=====

```

## Verifying the Cisco APIC Cluster Using the CLI

Cisco Application Policy Infrastructure Controller (APIC) release 4.2.(1) introduces the **cluster\_health** command, which enables you to verify the Cisco APIC cluster status step-by-step. The following output example demonstrates a scenario where everything is fine except for one node (ID 1002), which is inactive.





**Note** To use the `cluster_health` command, you must be logged in as admin.

To verify the cluster status:

```
F1-APIC1# cluster_health
Password:

Running...

Checking Wiring and UUID: OK
Checking AD Processes: Running
Checking All Apics in Commission State: OK
Checking All Apics in Active State: OK
Checking Fabric Nodes: Inactive switches: ID=1002(IP=10.1.176.66/32)
Checking Apic Fully-Fit: OK
Checking Shard Convergence: OK
Checking Leadership Degration: Optimal leader for all shards
Ping OOB IPs:
APIC-1: 172.31.184.12 - OK
APIC-2: 172.31.184.13 - OK
APIC-3: 172.31.184.14 - OK
Ping Infra IPs:
APIC-1: 10.1.0.1 - OK
APIC-2: 10.1.0.2 - OK
APIC-3: 10.1.0.3 - OK
Checking APIC Versions: Same (4.2(0.261a))
Checking SSL: OK

Done!
```

**Table 9: Cluster\_Health Verification Steps**

| Step                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Checking Wiring and UUID               | <p>Leaf switches provide infra connectivity between each Cisco APIC by detecting the Cisco APICs using LLDP. This step checks wiring issues between a leaf and a Cisco APIC that is detected during LLDP discovery.</p> <p>Any issues in here implies a leaf switch cannot provide infra connectivity for a Cisco APIC as it doesn't have a valid information. For example, a Cisco APIC UUID mismatch means the new APIC2 has a different UUID than the previously known APIC2.</p> <p>UUID – Universally Unique ID, or chassis ID in some outputs</p> |
| Checking AD Processes                  | Cisco APIC clustering is handled by the Appliance Director process on each Cisco APIC. This step checks if the process is running correctly.                                                                                                                                                                                                                                                                                                                                                                                                            |
| Checking All APICs in Commission State | To complete the Cisco APIC clustering, all Cisco APICs need to be commissioned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Step                                     | Description                                                                                                                                                                                                                                                                                            |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Checking All APICs in Active State       | To complete the Cisco APIC clustering, all commissioned Cisco APICs need to be active. If it is not active, the Cisco APIC may not be up yet.                                                                                                                                                          |
| Checking Fabric Nodes: Inactive switches | The Cisco APIC's communication are through infra connectivity provided by leaf and spine switches. This step checks inactive switches to ensure switches are providing infra connectivity.                                                                                                             |
| Checking APIC Fully-Fit                  | When Cisco APICs have established IP reachability to each other through infra network, it will synchronize its database to each other. When the synchronization completes, the status of all Cisco APICs become "Fully-Fit." Otherwise, the status will be "Data Layer Partially Diverged," and so on. |
| Checking Shard Convergence               | When Cisco APICs are not fully-fit, database shards need to be checked to see which service is not fully synchronized. If there is any service that has problems in synchronization, you may reach out to Cisco TAC for further troubleshooting.                                                       |
| Checking Leadership Degration            | In ACI, each database shard has one leader shard distributed to each Cisco APIC in the cluster. This step shows if all shards have an optimal leader. If there is an issue in here when all Cisco APICs are up, you may reach out to Cisco TAC for further troubleshooting.                            |
| Ping OOB IPs                             | This step is to check if all Cisco APICs are up and operational by pinging the OOB IP which is configured separately from clustering.                                                                                                                                                                  |
| Ping Infra IPs                           | This step is to check if there is infra connectivity between each Cisco APIC. Cisco APIC clustering is performed through infra connectivity instead of OOB.                                                                                                                                            |
| Checking APIC Versions                   | All Cisco APICs should be on a same version to complete clustering.                                                                                                                                                                                                                                    |
| Checking SSL                             | All Cisco APICs need to have a valid SSL that should be built-in when a Cisco APIC is shipped as an appliance. Without a valid SSL, the server cannot operate the Cisco APIC OS correctly.                                                                                                             |



## APPENDIX **B**

# Configuring the Cisco APIC Using the REST API

---

- [Expanding the APIC Cluster Using the REST API](#), on page 105
- [Contracting the APIC Cluster Using the REST API](#), on page 105
- [Reducing the APIC Cluster Size](#), on page 107
- [Switching Over Active APIC with Standby APIC Using REST API](#), on page 108
- [Registering an Unregistered Switch Using the REST API](#), on page 109
- [Adding a Switch Before Discovery Using the REST API](#), on page 109
- [Removing a Switch to Maintenance Mode Using the REST API](#), on page 110
- [Inserting a Switch to Operational Mode Using the REST API](#), on page 110
- [Configuring a Remote Location Using the REST API](#), on page 111
- [Sending an On-Demand Tech Support File Using the REST API](#), on page 111
- [Finding Your Switch Inventory Using the REST API](#), on page 112

## Expanding the APIC Cluster Using the REST API

The cluster drives its actual size to the target size. If the target size is higher than the actual size, the cluster size expands.

---

**Step 1** Set the target cluster size to expand the APIC cluster size.

**Example:**

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=3/>
```

**Step 2** Physically connect the APIC controllers that you want to add to the cluster.

---

## Contracting the APIC Cluster Using the REST API

Use this procedure to shrink the cluster size by removing controllers. For more information about contracting the cluster size, see [Contracting the Cisco APIC Cluster](#), on page 70.



**Note** Beginning with Cisco APIC Release 6.0(2), two additional properties are added to the API command to allow forcing the decommission operation. The new object properties are:

- `infraClusterPol:shrink`
  - `false`: (default) If the target cluster size (`infraClusterPol:size`) is less than the current operational cluster size, the APICs to be removed must be decommissioned manually, as in earlier releases.
  - `true`: If the target cluster size is less than the current operational cluster size, a cluster shrink decommission is triggered. The APICs to be removed are decommissioned automatically, beginning with the APIC with the highest controller ID number.
- `infraWiNode:force`
  - `false`: (default) The decommission proceeds only if the cluster is not in an unhealthy or upgrade state, where a decommission may not be proper.
  - `true`: The decommission proceeds regardless of the cluster state.

This example shows how to contract the cluster from three APIC controllers to a single controller. To achieve a target size of one, APIC3 and APIC2 must be decommissioned, in that order.

**Step 1** Set the target cluster size so as to contract the APIC cluster size.

When the cluster size is reduced with `shrink='true'`, the APICs to be removed are decommissioned automatically. Otherwise, they must be decommissioned manually.

**Example:**

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 />
```

You must now manually decommission the APICs to be removed, as shown in the next steps.

**Example:**

Cisco APIC Release 6.0(2) and later, using 'shrink' property:

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 shrink='true' />
```

With `shrink='true'`, the following steps can be skipped. The APICs to be removed are decommissioned automatically.

```
POST
https://<IP address>/api/node/mo/uni/controller.xml
<infraClusterPol name='default' size=1 shrink='false' />
```

With `shrink='false'`, you must now manually decommission the APICs to be removed, as shown in the next steps.

**Step 2** Decommission APIC3 on APIC1 for cluster contraction.

**Example:**

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=3 adminSt='out-of-service' />
```

The decommission proceeds only if the cluster is in a healthy state.

**Example:**

Cisco APIC Release 6.0(2) and later, using 'force' property:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=3 adminSt='out-of-service' force='true' />
```

With `force='true'`, the decommission proceeds regardless of the cluster state.

**Step 3** Decommission APIC2 on APIC1 for cluster contraction.

**Example:**

Cisco APIC Release 6.0(1) and earlier:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 adminSt='out-of-service' />
```

The decommission proceeds only if the cluster is in a healthy state.

**Example:**

Cisco APIC Release 6.0(2) and later, using 'force' property:

```
POST
https://<IP address>/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 adminSt='out-of-service' force='false' />
```

With `force='false'`, the decommission proceeds only if the cluster is in a healthy state.

---

The operation cluster size shrinks only after the last appliance is decommissioned, and not after the administrative size is changed. Verify after each controller is decommissioned that the operational state of the controller is unregistered and that the controller is no longer in service in the cluster.



---

**Note** If a decommissioned APIC controller is not promptly removed from the fabric, it might be rediscovered, which could cause problems. In that case, follow the instructions in [Reducing the APIC Cluster Size, on page 71](#) to remove the controller.

---

## Reducing the APIC Cluster Size

Follow these guidelines to reduce the Cisco Application Policy Infrastructure Controller (APIC) cluster size and decommission the Cisco APICs that are removed from the cluster:



---

**Note** Failure to follow an orderly process to decommission and power down Cisco APICs from a reduced cluster can lead to unpredictable outcomes. Do not allow unrecognized Cisco APICs to remain connected to the fabric.

---

- Reducing the cluster size increases the load on the remaining Cisco APICs. Schedule the Cisco APIC size reduction at a time when the demands of the fabric workload will not be impacted by the cluster synchronization.
- If one or more of the Cisco APICs' health status in the cluster is not "fully fit," remedy that situation before proceeding.
- Reduce the cluster target size to the new lower value. For example if the existing cluster size is 6 and you will remove 3 controllers, reduce the cluster target size to 3.
- Starting with the highest numbered controller ID in the existing cluster, decommission, power down, and disconnect the APIC one by one until the cluster reaches the new lower target size.

Upon the decommissioning and removal of each controller, the Cisco APIC synchronizes the cluster.



**Note** After decommissioning a Cisco APIC from the cluster, promptly power it down and disconnect it from the fabric to prevent its rediscovery. Before returning it to service, do a wiped clean back to factory reset.

If the disconnection is delayed and a decommissioned controller is rediscovered, follow these steps to remove it:

1. Power down the Cisco APIC and disconnect it from the fabric.
2. In the list of Unauthorized Controllers, reject the controller.
3. Erase the controller from the GUI.

- Cluster synchronization stops if an existing Cisco APIC becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the Cisco APIC must synchronize upon the removal of a controller, the time required to decommission and complete cluster synchronization for each controller could be more than 10 minutes per controller.



**Note** Complete the entire necessary decommissioning steps, allowing the Cisco APIC to complete the cluster synchronization accordingly before making additional changes to the cluster.

## Switching Over Active APIC with Standby APIC Using REST API

Use this procedure to switch over an active APIC with standby APIC using REST API.

Switch over active APIC with standby APIC.

```
URL for POST: https://ip address/api/node/mo/topology/pod-initiator_pod_id/node-initiator_id/av.xml
Body: <infraWInode id=outgoing_apic_id targetMbSn=backup-serial-number/>
where initiator_id = id of an active APIC other than the APIC being replaced.
pod-initiator_pod_id = pod ID of the active APIC
backup-serial-number = serial number of standby APIC
```

**Example:**

```
https://ip address/api/node/mo/topology/pod-1/node-1/av.xml
<infraWiNode id=2 targetMbSn=FCH1750V00Q/>
```

## Registering an Unregistered Switch Using the REST API

Use this procedure to register a switch from the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the REST API.



**Note** This procedure is identical to "Adding a Switch Before Discovery Using the REST API". When you apply the code, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

Add a switch description.

**Example:**

```
POST
https://<IP address>/api/policymgr/mo/uni.xml

<!-- /api/policymgr/mo/uni.xml -->
<polUni>
<ctrlrInst>
 <fabricNodeIdentPol>
 <fabricNodeIdentP nodeType="tier-2-leaf" podId="1" serial="XXXXXXXXXX"
 name="tier-2-leaf-leaf1" nodeId="101"/>
 </fabricNodeIdentPol>
</ctrlrInst>
</polUni>
```

## Adding a Switch Before Discovery Using the REST API

Use this procedure to add a switch to the **Nodes Pending Registration** tab on the **Fabric Membership** work pane using the REST API.



**Note** This procedure is identical to "Registering an Unregistered Switch Using the REST API". When you apply the code, the system determines if the node exists and, if not, adds it. If the node does exist, the system registers it.

Add a switch description.

**Example:**

```

POST
https://<IP address>/api/policymgr/mo/uni.xml

<!-- /api/policymgr/mo/uni.xml -->
<polUni>
<ctrlrInst>
 <fabricNodeIdentPol>
 <fabricNodeIdentP nodeType="tier-2-leaf" podId="1" serial="XXXXXXXX"
 name="tier-2-leaf1" nodeId="101"/>
 </fabricNodeIdentPol>
</ctrlrInst>
</polUni>

```

---

## Removing a Switch to Maintenance Mode Using the REST API

Use this procedure to remove a switch to maintenance mode using the REST API.

---

Remove a switch to maintenance mode.

**Example:**

```

POST
https://<IP address>/api/node/mo/uni/fabric/outofsvc.xml

<fabricOOServicePol
 descr=""
 dn=""
 name="default"
 nameAlias=""
 ownerKey=""
 ownerTag="">
 <fabricRsDecommissionNode
 debug="yes"
 dn=""
 removeFromController="no"
 tDn="topology/pod-1/node-102"/>
</fabricOOServicePol>

```

---

## Inserting a Switch to Operational Mode Using the REST API

Use this procedure to insert a switch to operational mode using the REST API.

---

Insert a switch to operational mode.

**Example:**



```

POST
https://<IP address>/api/node/mo/uni/fabric/outofsvc.xml

<fabricOOServicePol
 descr=""
 dn=""
 name="default"
 nameAlias=""
 ownerKey=""
 ownerTag="">
 <fabricRsDecommissionNode
 debug="yes"
 dn=""
 removeFromController="no"
 tDn="topology/pod-1/node-102"
 status="deleted"/>
</fabricOOServicePol>

```

## Configuring a Remote Location Using the REST API

This procedure explains how to create a remote location using the REST API.

```

<fileRemotePath name="local" host="host or ip" protocol="ftp|scp|sftp" remotePath="path to
folder" userName="uname" userPasswd="pwd" />

```

## Sending an On-Demand Tech Support File Using the REST API

**Step 1** Set the remote destination for a technical support file using the REST API, by sending a POST with XML such as the following example:

**Example:**

```

<fileRemotePath userName="" remotePort="22" remotePath="" protocol="sftp" name="ToSupport"
host="192.168.200.2"
dn="uni/fabric/path-ToSupport" descr="">

<fileRsARemoteHostToEpg tDn="uni/tn-mgmt/mgmt-default/oob-default"/>

</fileRemotePath>

```

**Step 2** Generate an on-demand technical support file using the REST API by sending a POST with XML such as the following:

**Example:**

```

<dbgexpTechSupOnD upgradeLogs="no" startTime="unspecified" name="Tech_Support_9-20-16"
exportToController="no" endTime="unspecified" dn="uni/fabric/tsod-Tech_Support_9-20-16" descr=""
compression="gzip" category="forwarding" adminSt="untriggered">
 <dbgexpRsExportDest tDn="uni/fabric/path-ToSupport"/>
 <dbgexpRsTsSrc tDn="topology/pod-1/node-102/sys"/>
 <dbgexpRsTsSrc tDn="topology/pod-1/node-103/sys"/>
 <dbgexpRsTsSrc tDn="topology/pod-1/node-101/sys"/>
 <dbgexpRsData tDn="uni/fabric/tscont"/>
</dbgexpTechSupOnD>
<fabricFuncP>

```

```
<fabricCtrlrPGrp name="default">
 <fabricRsApplTechSupOnDemand tnDbgexpTechSupOnDName=" Tech_Support_9-20-16"/>
</fabricCtrlrPGrp>
</fabricFuncP>
```

## Finding Your Switch Inventory Using the REST API

This section explains how to find your switch model and serial numbers using the REST API

Find your switch inventory as follows:

### Example:

```
GET
https://192.0.20.123/api/node/mo/topology/pod-1.json?query-target=children&target-subtree-class=fabricNode
```

The following response is returned:

```
response:
{
 "totalCount":"8",
 "imdata":
 [{
 "fabricNode":{
 "attributes":{
 "adSt":"on",
 "childAction":"",
 "delayedHeartbeat":"no",
 "dn":"topology/pod-1/node-103",
 "fabricSt":"active",
 "id":"103",
 "lcOwn":"local",
 "modTs":"2016-10-08T14:49:35.665+00:00",
 "model":"N9K-C9396PX",
 "monPolDn":"uni/fabric/monfab-default",
 "name":"leaf3",
 "nameAlias":"",
 "role":"leaf",
 "serial":"TEP-1-103",
 "status":"","uid":"0",
 "vendor":"Cisco Systems, Inc",
 "version":""}
 },{
 "fabricNode":{
 "attributes":{
 "adSt":"on",
 "childAction":"",
 "delayedHeartbeat":"no",
 "dn":"topology/pod-1/node-105",
 "fabricSt":"active",
 "id":"105",
 "lcOwn":"local",
 "modTs":"2016-10-08T14:47:52.011+00:00",
 "model":"N9K-C9508",
 "monPolDn":"uni/fabric/monfab-default",
 "name":"spine2",
```

```
 "nameAlias": "",
 "role": "spine",
 "serial": "TEP-1-105", "status": "",
 "uid": "0",
 "vendor": "Cisco Systems, Inc",
 "version": ""
 ...
 [TRUNCATED]
 ...
}
```

---

