



Cisco APIC Getting Started Guide, Release 4.1(x)

First Published: 2019-03-28

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

Cisco APIC Documentation Roadmap 3

Simplified Approach to Configuring in Cisco APIC 4

Changing the BIOS Default Password 4

About the APIC 4

Setting up the Cisco APIC 5

Setup for Active and Standby APIC 8

Provisioning IPv6 Management Addresses on APIC Controllers 13

Accessing the GUI 14

Accessing the REST API 15

Accessing the NX-OS Style CLI 15

Accessing the NX-OS Style CLI from a Terminal 16

Accessing the NX-OS Style CLI from the GUI 16

Accessing the Object Model CLI 17

CHAPTER 3

APIC GUI Overview 19

Overview of the GUI 19

Menu Bar and Submenu Bar 20

Menu Bar Tabs 21

System Tab 21

Tenants Tab 21

Fabric Tab	22
Virtual Networking Tab	22
L4-L7 Services Tab	22
Admin Tab	23
Operations Tab	23
Apps Tab	23
Integrations Tab	23
Menu Bar Tools	23
Search	23
Alerts	23
User Profile and Preferences	24
System Tools	25
Navigation Pane	25
Work Pane	26
Common Pages in the Work Pane	27
Personalizing the Interface	27
Naming the APIC GUI	27
Adding a Login Banner to the CLI or GUI	28
Single-Browser Session Management	28
Deployment Warning and Policy Usage Information	28
Graphical Configuration of Ports	29
Viewing an API Interchange in the GUI	30
GUI Icons	32
Fault, Statistics, and Health Level Icons	33

CHAPTER 4 Fabric Initialization and Switch Discovery 35

Initializing the Fabric	35
About Fabric Initialization	35
Fabric Topology (Example)	35
Multi-Tier Fabric Topology (Example)	36
Switch Discovery	38
About Switch Discovery with the APIC	38
Switch Registration with the APIC Cluster	39
Registering an Unregistered Switch Using the GUI	39

Adding a Switch Before Discovery Using the GUI	41
Switch Discovery Validation and Switch Management from the APIC	42
Validating the Registered Switches Using the GUI	42
Validating the Fabric Topology	42
Validating the Fabric Topology Using the GUI	42
Unmanaged Switch Connectivity in VM Management	43
Graceful Insertion and Removal (GIR) Mode	43
Maintenance Mode	43
Removing a Switch to Maintenance Mode Using the GUI	44
Inserting a Switch to Operational Mode Using the GUI	45

CHAPTER 5

Cisco APIC Cluster Management	47
APIC Cluster Overview	47
Expanding the Cisco APIC Cluster	47
Contracting the Cisco APIC Cluster	48
Cluster Management Guidelines	48
Expanding the APIC Cluster Size	49
Reducing the APIC Cluster Size	50
Replacing Cisco APIC Controllers in the Cluster	50
Expanding the APIC Cluster Using the GUI	52
Contracting the APIC Cluster Using the GUI	52
Commissioning and Decommissioning Cisco APIC Controllers	53
Commissioning a Cisco APIC in the Cluster Using the GUI	53
Decommissioning a Cisco APIC in the Cluster Using the GUI	53
Shutting Down the APICs in a Cluster	54
Shutting Down all the APICs in a Cluster	54
Bringing Back the APICs in a Cluster	55
Cold Standby	55
About Cold Standby for a Cisco APIC Cluster	55
Verifying Cold Standby Status Using the GUI	56
Switching Over an Active APIC with a Standby APIC Using the GUI	57



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 up to this current 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 4.1(1)

Feature	Description	Where Documented
Multi-Tier Architecture (Tier-2 Leaf) Support	Cisco APIC now supports a 3-tier core-aggregation-access topology.	Multi-Tier Fabric Topology (Example), on page 36 Adding a Switch Before Discovery Using the GUI, on page 41
Bookmark button	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.	Work Pane, on page 26
Default tab	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.	Work Pane, on page 26
Integrations tab	The Integrations tab enables you to view all 3rd party integrations.	Integrations Tab, on page 23



CHAPTER 2

Initial Setup

This chapter contains the following sections:

- [Cisco APIC Documentation Roadmap](#), 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 4
- [Setting up the Cisco APIC](#) , on page 5
- [Accessing the GUI](#), on page 14
- [Accessing the REST API](#), on page 15
- [Accessing the NX-OS Style CLI](#), on page 15
- [Accessing the Object Model CLI](#), on page 17

Cisco APIC Documentation Roadmap

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

Document
<i>Application Centric Infrastructure Fabric Hardware Installation Guide</i>
<i>Cisco APIC Management, Installation, Upgrade, and Downgrade Guide</i>
<i>Cisco APIC Basic Configuration Guide</i>
<i>Cisco APIC Layer 2 Networking Configuration Guide</i>
<i>Cisco APIC Layer 3 Networking Configuration Guide</i>
<i>Cisco ACI Virtualization Guide</i>
<i>Cisco Application Centric Infrastructure Fundamentals</i>
<i>Cisco APIC Layer 4 to Layer 7 Services Deployment Guide</i>

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

The APIC controller 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.

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'.
- 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 mode 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

- 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.

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

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

Name	Description	Default Value
IP address pool for tunnel endpoint addresses	Tunnel endpoint address pool	<p>10.0.0.0/16</p> <p>This value is for the infrastructure virtual routing and forwarding (VRF) only.</p> <p>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.</p> <p>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.</p>
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>

Name	Description	Default Value
IPv4/IPv6 addresses for the out-of-band management	IP address that you use to access the Cisco APIC through the GUI, CLI, or API. This address must be a reserved address from the VRF of a customer	—
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	auto Valid values are as follows <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	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

Name	Description	Default Value
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.
VLAN ID for infrastructure network ²	Infrastructure VLAN for Cisco APIC-to-switch communication including virtual switches 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.	

Name	Description	Default Value
IPv4/IPv6 addresses for the out-of-band management	IP address that you use to access the Cisco APIC through the GUI, CLI, or API. This address must be a reserved address from the VRF of a customer	—
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	auto Valid values are as follows <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	Password of the system administrator This password must be at least 8 characters with one special character.	—

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 is a sample of the initial setup dialog as displayed on the console:

```
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]: sec-ifc5
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]: 172.23.142.29/21
  Enter the IPv4 address of the default gateway [None]: 172.23.136.1
  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: sec-ifc5
  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: 172.23.142.29/21
  Default gateway: 172.23.136.1
  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]:

```

Provisioning IPv6 Management Addresses on APIC Controllers

IPv6 management addresses can be provisioned on the APIC controller at setup time or through a policy once the APIC controller is operational. Pure IPv4, Pure IPv6 or dual stack (i.e both IPv6 and IPv4 addresses) are supported. The following snippet is of a typical setup screen that describes how to setup dual stack (IPv6 and IPv4) addresses for out-of-band management interfaces during the setup:

```

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:

```

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.

- 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

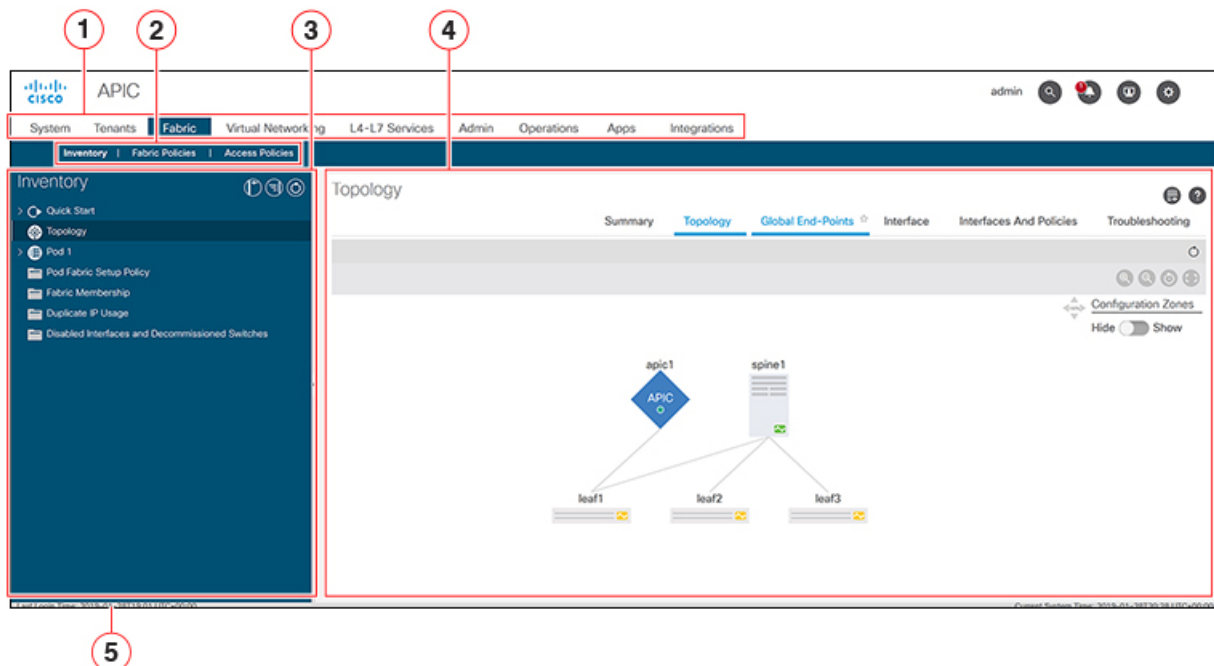
This chapter contains the following sections:

- [Overview of the GUI, on page 19](#)
- [Menu Bar and Submenu Bar, on page 20](#)
- [Navigation Pane, on page 25](#)
- [Work Pane, on page 26](#)
- [Personalizing the Interface, on page 27](#)
- [Single-Browser Session Management, on page 28](#)
- [Deployment Warning and Policy Usage Information, on page 28](#)
- [Graphical Configuration of Ports, on page 29](#)
- [Viewing an API Interchange in the GUI, on page 30](#)
- [GUI Icons, on page 32](#)

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 and tool icons—See [Menu Bar and Submenu Bar, on page 20](#)
2. Submenu bar—See [Menu Bar and Submenu Bar, on page 20](#)
3. Navigation pane—See [Navigation Pane, on page 25](#)
4. Work pane—See [Work Pane, on page 26](#)
5. Last Login—Displays the date and time of the last instance of the current user's login.
6. Integrations—See [Integrations Tab, on page 23](#)

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 30](#).





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 23
	Alerts, on page 23
	User Profile and Preferences, on page 24
	System Tools, on page 25

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).

L4-L7 Services Tab

Use the **L4-L7 Services** tab to perform services such as importing packages that define Layer 4 to Layer 7 devices such as a firewall, SSL offload, load balancer, context switch, SSL termination device, or intrusion prevention system (IPS). In the **Inventory** submenu tab, you can view existing Layer 4 to Layer 7 devices registered with the controller. The **Packages** submenu tab allows you to import L4-L7 device packages, which are used to define, configure, and monitor a network service device.

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.

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 (*).

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 3: 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.


User Profile and Preferences

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 4: User Profile and Preferences



The following selections are available:

- **Favorites**—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 28](#).
 - **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 28](#).
- **Change Deployment Settings**—Enable and set the scope of the deployment notification. See [Deployment Warning and Policy Usage Information, on page 28](#).
- **Logout**—Exit the APIC configuration GUI.

System Tools

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

Figure 5: System Tools



The following selections are available:

- **Help**—Display the online help.
- **Documentation**—Display links to API documentation and to the APIC documentation home page.
- **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 30](#).
- **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**—
- **About**—Display the APIC version.



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

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.
- A link to context-sensitive online help that is represented by a question mark icon in the upper right

corner.



- 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 23
	Alerts, on page 23
	User Profile and Preferences, on page 24
	System Tools, on page 25
	Bookmark this page
	Displays online help information for the current menu 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

Icons	Description
	Pod
	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.



CHAPTER 4

Fabric Initialization and Switch Discovery

This chapter contains the following sections:

- [Initializing the Fabric, on page 35](#)
- [Switch Discovery, on page 38](#)
- [Graceful Insertion and Removal \(GIR\) Mode, on page 43](#)

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 6: 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 7: 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 6: 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

* 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).

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.

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, navigate to **Fabric > Inventory > Fabric Membership**.

Step 2 In the **Fabric Membership** 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 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 3 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 4 Right-click the leaf switch row, select **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
Pod ID	Identifier of the pod where the node is located.
Node ID	<p>A number greater than 100. The first 100 IDs are reserved for APIC appliance nodes.</p> <p>Note 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>Note After the node ID is assigned, it cannot be updated. After the node has been added to the Registered Nodes tab table, you can update the node name by right-clicking the table row and choosing Edit Node and Rack Name.</p>
RL TEP Pool	Tunnel endpoint (TEP) pool identifier for the node.
Role	<p>The assigned node role. The options are:</p> <ul style="list-style-type: none"> • tier-2 leaf • leaf • remote leaf • spine • unspecified
Node Name	The node name, such as leaf1 or spine3.
Rack Name	The name of the rack in which the node is installed. Select Default or select Create Rack to add a name and description.

- Click **Register**.

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 5 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, navigate to **Fabric > Inventory > Fabric Membership**.

Step 2 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 3 Configure the following settings:

Field	Setting
Pod ID	Identify the pod where the node is located.
Serial Number	Required: Enter the serial number of the switch.
Node ID	<p>Required: Enter a number greater than 100. The first 100 IDs are reserved for APIC appliance nodes.</p> <p>Note We recommend that leaf nodes and spine nodes be numbered differently. For example, number leaves in the 100 range (such as 101, 102) and number spines in the 200 range (such as 201, 202).</p> <p>Note After the node ID is assigned, it cannot be updated. After the node has been added to the Registered Nodes tab table, you can update the node name by right-clicking the table row and choosing Edit Node and Rack Name.</p>
Switch Name	The node name, such as leaf1 or spine3.
Role	<p>Choose the assigned node role. The options are:</p> <ul style="list-style-type: none"> • leaf <p>Check one of the following boxes if applicable:</p> <ul style="list-style-type: none"> • Is Remote • Is Virtual • Is Tier-2 Leaf • spine <p>Check the following box if applicable:</p> <ul style="list-style-type: none"> • Is Virtual • unspecified

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, 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 39](#) 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.

Graceful Insertion and Removal (GIR) Mode

Maintenance Mode

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

- **Graceful Insertion and Removal (GIR)**: The operation used to isolate a switch from user traffic.
- **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. Therefore, performing a firmware upgrade for a switch in this state is not supported, since it may fail or may get stuck in an incomplete status indefinitely if you attempt to perform a firmware upgrade on it while it is in this state.

The maintenance mode allows you to isolate a switch from the network with minimum service disruption. In the maintenance mode you can perform real-time debugging without affecting traffic.

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
- Upgrading or downgrading a switch in maintenance mode is not supported.
 - 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, **IS-IS metric for redistributed routes** should be set to less than 63. 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.

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**.
-



CHAPTER 5

Cisco APIC Cluster Management

- [APIC Cluster Overview, on page 47](#)
- [Expanding the Cisco APIC Cluster, on page 47](#)
- [Contracting the Cisco APIC Cluster, on page 48](#)
- [Cluster Management Guidelines, on page 48](#)
- [Expanding the APIC Cluster Using the GUI, on page 52](#)
- [Contracting the APIC Cluster Using the GUI, on page 52](#)
- [Commissioning and Decommissioning Cisco APIC Controllers, on page 53](#)
- [Shutting Down the APICs in a Cluster, on page 54](#)
- [Cold Standby, on page 55](#)

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 is comprised of 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, follow the guidelines below for making changes to the Cisco APIC cluster.



Note 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.

Follow these general guidelines when managing clusters:

- Cisco recommends that you have at least 3 active Cisco APICs in a cluster, along with additional standby Cisco APICs. A cluster size of 3, 5, or 7 Cisco APICs is recommended. A cluster size of 4 or 6 Cisco APICs is not recommended.
- 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 you intend to shut down. After the Cisco APIC has shutdown, 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 a Cisco APIC cluster is split into two or more groups, the ID of a node is changed and the changes are not synchronized across all Cisco APICs. This can cause inconsistency in the node IDs between Cisco

APICs and also the affected leaf nodes may not appear in the inventory in the Cisco APIC GUI. When you split a Cisco APIC cluster, decommission the affected leaf nodes from Cisco APIC and register them again, so that the inconsistency in the node IDs is resolved and the health status of the Cisco APICs in a cluster are in a fully fit state.

- Before configuring the Cisco APIC cluster, ensure that all 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 APIC cluster size and decommission the APIC controllers that are removed from the cluster:



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

- Reducing the cluster size increases the load on the remaining APIC controllers. Schedule the APIC controller 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 APIC controllers' 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 controller one by one until the cluster reaches the new lower target size.

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



Note After decommissioning an APIC controller from the cluster, power it down and disconnect it from fabric. Before returning it to service, do a wiped clean back to factory reset.

- Cluster synchronization stops if an existing APIC controller becomes unavailable. Resolve this issue before attempting to proceed with the cluster synchronization.
- Depending on the amount of data the 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 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 53](#))
 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 53](#))
- 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 Cisco Application Policy Infrastructure Controllers (APICs) to an existing cluster.

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**.
-

Contracting the APIC Cluster Using the GUI

-
- 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_controller_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 three tabs: **APIC Cluster**, **APIC-X**, and **Standby APIC**. 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 APIC controllers. A cluster of one, three, or more APIC controllers 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.

Commissioning and Decommissioning Cisco APIC Controllers

Commissioning a Cisco APIC in the Cluster Using the GUI

- 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**.
-

Decommissioning a Cisco APIC in the Cluster Using the GUI

This procedure decommissions a Cisco Application Policy Infrastructure Controller (APIC) in the cluster.



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.
-

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.

Important Notes

- The standby Cisco APICs are automatically updated with firmware updates to keep the backup Cisco APIC at same firmware version as the active cluster.
- During an upgrade process, after all the active Cisco APICs are upgraded, the standby Cisco APICs are also upgraded automatically.
- Temporary IDs are assigned to the standby Cisco APICs. After a standby Cisco APIC is switched over to an active Cisco APIC, a new ID is assigned.

- 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 is 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, if it was the only standby, you must configure a new standby.
- The following limitations are observed for retaining out of band address for the standby Cisco APIC after a fail over:
 - The standby (new active) Cisco APIC may not retain its out of band address if more than 1 active Cisco APICs are down or unavailable.
 - The standby (new active) Cisco APIC may not retain its out of band address if it is in a different subnet than the active Cisco APIC. This limitation is only applicable for Cisco APIC release 2.x.
 - The standby (new active) Cisco APIC may not retain its IPv6 out of band address. This limitation is not applicable starting from Cisco APIC release 3.1.x.
 - The standby (new active) Cisco APIC may not retain its out of band address if you have configured a non-static OOB management IP address policy for the replacement (old active) Cisco APIC.
 - The standby (new active) Cisco APIC may not retain its out of band address if it is not in a pod that has an active Cisco APIC.



Note If you want to retain the standby Cisco APIC's out of band address despite the limitations, you must manually change the OOB policy for the replaced Cisco APIC after the replace operation had completed successfully.

- There must be three active Cisco APICs to add a standby Cisco APIC.
- 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.

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.
-

