



Cisco APIC Getting Started Guide, Release 1.2(x) and Release 1.3(x)

First Published: 2015-12-08

Last Modified: 2016-09-08

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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 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.

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: <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. (1110R)

© 2015-2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface xv

Audience xv

Document Conventions xv

Related Documentation xvii

Documentation Feedback xviii

Obtaining Documentation and Submitting a Service Request xviii

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Initial POD Setup and Overview 3

First-Time Access 3

Getting Started Guide Contents 3

Simplified Approach to Configuring in Cisco APIC 3

Installing the Cisco Application Centric Infrastructure Fabric Hardware 4

Changing the BIOS Default Password 4

About the APIC 4

Setting up the APIC 5

Provisioning IPv6 Management Addresses on APIC Controllers 11

Accessing the GUI 12

Accessing the REST API 13

Accessing the Object Model CLI 13

Accessing the NX-OS Style CLI 14

Overview of the GUI 15

Deployment Warning and Policy Usage Information 15

Toggling Between Basic and Advanced GUI Modes 15

Menu Bar and Submenu Bar 16

SYSTEM Tab	17
TENANTS Tab	17
FABRIC Tab	18
VM NETWORKING Tab	18
L4-L7 SERVICES Tab	18
ADMIN Tab	18
Search Icon	18
Navigation Pane	18
Work Pane	19
GUI Icons	21
Fault, Statistics, and Health Level Icons	22
API Inspector	22
Viewing an API Interchange in the GUI	22
Initializing the Fabric	24
About Fabric Initialization	24
Example Topology	24
Example Topology Connections	25
Switch Discovery with the APIC	26
About Switch Discovery with the APIC	26
Switch Registration with the APIC Cluster	26
Registering the Unregistered Switches Using the GUI	26
Switch Discovery Validation and Switch Management from the APIC	27
Validating the Registered Switches Using the GUI	27
Validating the Fabric Topology	28
Validating the Fabric Topology Using the GUI	28
Unmanaged Switch Connectivity in VM Management	28

CHAPTER 3	Using the Basic GUI	31
	Toggling Between Basic and Advanced GUI Modes	31
	About Getting Started with APIC Examples	32
	About Switch Discovery with the APIC	33
	Switch Registration with the APIC Cluster	33
	Registering the Unregistered Switches Using the GUI	33
	Switch Discovery Validation and Switch Management from the APIC	34
	Validating the Registered Switches Using the GUI	34

Validating the Fabric Topology	35
Validating the Fabric Topology Using the GUI	35
Unmanaged Switch Connectivity in VM Management	35
Configuring Network Time Protocol	36
Time Synchronization and NTP	36
In-Band and Out-of-Band Management NTP	37
NTP over IPv6	37
Configuring NTP Using the Basic GUI	37
Verifying NTP Policy Deployed to Each Node Using the CLI	38
Creating User Accounts	38
Configuring a Local User	38
Configuring a Remote User	38
Configuring a Local User Using the GUI	39
AV Pair on the External Authentication Server	40
Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs	41
Best Practice for Assigning AV Pairs	41
Configuring an AV Pair on the External Authentication Server	41
Configuring a Remote User Using the GUI	42
Adding Management Access in the GUI	43
IPv4/IPv6 Addresses and In-Band Policies	43
IPv4/IPv6 Addresses in Out-of-Band Policies	43
Configuring Management Access	44
Configuring In-Band Management Access Using the Basic GUI	44
Configuring Out-of-Band Management Access Using the Basic GUI	44
IPv6 Table Modifications to Mirror the Existing IP Tables Functionality	45
Configuring a VMM Domain	46
Configuring Virtual Machine Networking Policies	46
About the VM Manager	46
Prerequisites for Creating a VMM Domain Profile	47
Custom User Account with Minimum VMware vCenter Privileges	47
Creating a VMM Domain Profile	48
Creating a vCenter Domain Profile Using the Basic GUI	48
Creating a vCenter and a vShield Domain Profile Using the Basic GUI	50
Creating Tenants, VRF, and Bridge Domains	51
Tenants Overview	51

Tenant Creation	51
VRF and Bridge Domains	51
Creating a Tenant, VRF, and Bridge Domain Using the Basic GUI	52
Configuring Server or Service Policies	52
Configuring a DHCP Relay Policy	52
Configuring a DHCP Server Policy for the APIC Infrastructure Using the Basic GUI	53
Configuring a DNS Service Policy	54
Configuring External Destinations with an In-Band DNS Service Policy	54
Policy for Priority of IPv4 or IPv6 in a DNS Profile	55
Dual Stack IPv4 and IPv6 DNS Servers	56
Dual-Stack IPv4 and IPv6 Environment	56
Configuring a DNS Service Policy to Connect with DNS Providers Using the Basic GUI	57
Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI	58
Configuring External Connectivity for Tenants	59
Configuring an MP-BGP Route Reflector Using the Basic GUI	59
Verifying the MP-BGP Route Reflector Configuration	60
Creating OSPF External Routed Network for Management Tenant Using Basic GUI	60
Deploying an Application Policy	62
Three-Tier Application Deployment	62
Parameters to Create a Filter for http	63
Parameters to Create Filters for rmi and sql	63
Example Application Profile Database	63
Deploying an Application Policy Using the Basic GUI	64
CHAPTER 4	Using the NX-OS Style CLI
	65
Accessing the NX-OS Style CLI	65
Using the NX-OS Style CLI for APIC	66
About Getting Started with APIC Examples	69
About Switch Discovery with the APIC	69
Switch Registration with the APIC Cluster	70
Registering Unregistered Switches Using the NX-OS Style CLI	70
Switch Discovery Validation and Switch Management from the APIC	71

Configuring Network Time Protocol	71
Time Synchronization and NTP	71
In-Band and Out-of-Band Management NTP	71
NTP over IPv6	72
Configuring NTP Using the NX-OS Style CLI	72
Verifying NTP Operation Using the NX-OS Style CLI	74
Verifying NTP Policy Deployed to Each Node Using the NX-OS Style CLI	75
Creating User Accounts	75
Configuring a Local User Using the NX-OS Style CLI	75
AV Pair on the External Authentication Server	75
Changing Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs Using the NX-OS Style CLI	76
Best Practice for Assigning AV Pairs	77
Configuring an AV Pair on the External Authentication Server	77
Configuring a Remote User Using the NX-OS Style CLI	78
Configuring a Remote User With the NX-OS Style CLI	78
Adding Management Access	79
IPv4/IPv6 Addresses and In-Band Policies	79
IPv4/IPv6 Addresses in Out-of-Band Policies	79
Adding Management Access Using the NX-OS Style CLI	79
Configuring In-Band Management Access for APIC Controller, Spine, Leaf Switches Using the NX-OS CLI	80
Configuring Out-of-Band Management Access for APIC Controller, Spine, Leaf Switches Using the NX-OS CLI	82
IPv6 Table Modifications to Mirror the Existing IP Tables Functionality	84
Configuring a VLAN Domain	85
Configuring a VLAN Domain Using the NX-OS Style CLI	85
Configuring a VMM Domain	86
Configuring a VMM Domain Using the NX-OS Style CLI	86
Configuring Virtual Machine Networking Policies	86
About the VM Manager	86
Prerequisites for Creating a VMM Domain Profile	87
Custom User Account with Minimum VMware vCenter Privileges	87
Creating a VMM Domain Profile	88
Creating a vCenter Domain Profile Using the NX-OS Style CLI	88

Creating a vCenter and a vShield Domain Profile Using the NX-OS Style CLI	89
Creating Tenants, VRFs, and Bridge Domains	91
Creating a Tenant, VRF, and Bridge Domain Using the NX-OS Style CLI	91
Creating an Application Profile and EPG Using the NX-OS Style CLI	92
Mapping a VLAN on a Port to the EPG Using the NX-OS Style CLI	93
Deploying an Application Policy	94
Three-Tier Application Deployment	94
Parameters to Create an Access List for HTTP	95
Parameters to Create an Access List for RMI and SQL	95
Example Application Profile Database	96
Deploying an Application Policy Using the NX-OS Style CLI	96
Configuring External L3 Connectivity for Tenants	99
Configuring an MP-BGP Route Reflector for the ACI Fabric	99
Creating an OSPF External Routed Network for a Tenant Using the NX-OS CLI	99
Configuring Server or Service Policies	101
Configuring a DHCP Relay Policy	101
Configuring a DHCP Server Policy for the APIC Infrastructure Using the NX-OS Style CLI	102
Configuring a DNS Service Policy	102
Configuring External Destinations with an In-Band DNS Service Policy	102
Policy for Priority of IPv4 or IPv6 in a DNS Profile	104
Dual Stack IPv4 and IPv6 DNS Servers	104
Dual-Stack IPv4 and IPv6 Environment	105
Configuring a DNS Service Policy to Connect with DNS Providers Using the NX-OS Style CLI	105
Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the NX-OS Style CLI	105

CHAPTER 5**Using the Advanced GUI 107**

Toggling Between Basic and Advanced GUI Modes	107
About Getting Started with APIC Examples	108
About Switch Discovery with the APIC	109
Switch Registration with the APIC Cluster	109
Registering the Unregistered Switches Using the GUI	109
Switch Discovery Validation and Switch Management from the APIC	110

Validating the Registered Switches Using the GUI	110
Validating the Fabric Topology	111
Validating the Fabric Topology Using the GUI	111
Unmanaged Switch Connectivity in VM Management	111
Configuring Network Time Protocol	112
Time Synchronization and NTP	112
In-Band and Out-of-Band Management NTP	113
NTP over IPv6	113
Configuring NTP Using the Advanced GUI	113
Verifying NTP Operation Using the GUI	114
Verifying NTP Policy Deployed to Each Node Using the CLI	114
Creating User Accounts	115
Configuring a Local User	115
Configuring a Remote User	115
Configuring a Local User Using the GUI	116
AV Pair on the External Authentication Server	117
Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs	118
Best Practice for Assigning AV Pairs	118
Configuring an AV Pair on the External Authentication Server	118
Configuring a Remote User Using the GUI	119
Adding Management Access	120
IPv4/IPv6 Addresses and In-Band Policies	120
IPv4/IPv6 Addresses in Out-of-Band Policies	120
Configuring Management Access	121
Configuring In-Band Management Access Using the Advanced GUI	121
Configuring Out-of-Band Management Access Using the Advanced GUI	124
Modifying the IP Address of an APIC Controller Using the GUI	126
IPv6 Table Modifications to Mirror the Existing IP Tables Functionality	126
Management Connectivity Modes	127
Configuring Layer 2 Management Connectivity Using the Advanced GUI	128
Configuring Layer 3 Management Connectivity Using the Advanced GUI	129
Validating Management Connectivity	130
Configuring a VMM Domain	130
Configuring Virtual Machine Networking Policies	130
About the VM Manager	131

About Attachable Entity Profile	131
Prerequisites for Creating a VMM Domain Profile	132
Custom User Account with Minimum VMware vCenter Privileges	133
Creating a VMM Domain Profile	134
Creating a vCenter Domain Profile Using the GUI	134
Creating a vCenter and a vShield Domain Profile Using the Advanced GUI	135
Creating Tenants, VRF, and Bridge Domains	137
Tenants Overview	137
Tenant Creation	137
VRF and Bridge Domains	137
Creating a Tenant, VRF, and Bridge Domain Using the GUI	137
Configuring an Enforced Bridge Domain Using the Basic GUI	138
Configuring an Enforced Bridge Domain Using the NX-OS Style CLI	139
Configuring an Enforced Bridge Domain Using the REST API	140
Configuring Server or Service Policies	140
Configuring a DHCP Relay Policy	140
Configuring a DHCP Server Policy for the APIC Infrastructure Using the GUI	141
Configuring a DNS Service Policy	142
Configuring External Destinations with an In-Band DNS Service Policy	142
Policy for Priority of IPv4 or IPv6 in a DNS Profile	143
Dual Stack IPv4 and IPv6 DNS Servers	144
Dual-Stack IPv4 and IPv6 Environment	144
Configuring a DNS Service Policy to Connect with DNS Providers Using the Advanced GUI	145
Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI	145
Configuring External Connectivity for Tenants	147
Configuring an MP-BGP Route Reflector Using the Advanced GUI	147
Verifying the MP-BGP Route Reflector Configuration	148
Creating an OSPF External Routed Network for Management Tenant Using the Advanced GUI	148
Deploying an Application Policy	150
Three-Tier Application Deployment	150
Parameters to Create a Filter for http	151
Parameters to Create Filters for rmi and sql	151

Example Application Profile Database	152
Deploying an Application Policy Using the GUI	152
Creating a Filter Using the GUI	152
Creating a Contract Using the GUI	153
Creating an Application Profile Using the GUI	153
Creating EPGs Using the GUI	154
Consuming and Providing Contracts Using the GUI	154

CHAPTER 6**Using the REST API 157**

About Getting Started with APIC Examples	157
About Switch Discovery with the APIC	157
Switch Registration with the APIC Cluster	158
Registering the Unregistered Switches Using the REST API	158
Switch Discovery Validation and Switch Management from the APIC	158
Validating the Registered Switches Using the REST API	159
Validating the Fabric Topology	159
Validating the Fabric Topology Using the REST API	159
Unmanaged Switch Connectivity in VM Management	160
Configuring Network Time Protocol	161
Time Synchronization and NTP	161
In-Band and Out-of-Band Management NTP	162
NTP over IPv6	162
Configuring NTP Using the REST API	162
Verifying NTP Operation Using the GUI	163
Verifying NTP Policy Deployed to Each Node Using the CLI	163
Creating User Accounts	164
Configuring a Local User	164
Configuring a Remote User	164
Configuring a Local User Using the REST API	165
AV Pair on the External Authentication Server	165
Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs	166
Best Practice for Assigning AV Pairs	166
Configuring an AV Pair on the External Authentication Server	166
Configuring a Remote User Using the REST API	167
Adding Management Access	167

IPv4/IPv6 Addresses and In-Band Policies	168
IPv4/IPv6 Addresses in Out-of-Band Policies	168
Configuring Management Access	168
Configuring In-Band Management Access Using the REST API	168
Configuring Out-of-Band Management Access Using the REST API	171
Modifying the IP Address of an APIC Controller Using the REST API	173
IPv6 Table Modifications to Mirror the Existing IP Tables Functionality	174
Management Connectivity Modes	175
Configuring Layer 2 Management Connectivity Using the REST API	175
Configuring Layer 3 Management Connectivity Using the REST API	176
Validating Management Connectivity	177
Configuring a VMM Domain	178
Configuring Virtual Machine Networking Policies	178
About the VM Manager	178
About Attachable Entity Profile	179
Prerequisites for Creating a VMM Domain Profile	180
Custom User Account with Minimum VMware vCenter Privileges	180
Creating a VMM Domain Profile	181
Creating a vCenter Domain Profile Using the REST API	181
Creating a vCenter and a vShield Domain Profile Using the REST API	183
Creating Tenants, VRF, and Bridge Domains	185
Tenants Overview	185
Tenant Creation	185
VRF and Bridge Domains	185
Creating a Tenant, VRF, and Bridge Domain Using the REST API	186
Configuring Server or Service Policies	186
Configuring a DHCP Relay Policy	186
Configuring a DHCP Server Policy for the APIC Infrastructure Using the REST API	187
Configuring a DNS Service Policy	187
Configuring External Destinations with an In-Band DNS Service Policy	188
Policy for Priority of IPv4 or IPv6 in a DNS Profile	189
Dual Stack IPv4 and IPv6 DNS Servers	189
Dual-Stack IPv4 and IPv6 Environment	190

Configuring a DNS Service Policy to Connect with DNS Providers Using the REST API	190
Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI	191
Configuring External Connectivity for Tenants	192
Configuring an MP-BGP Route Reflector Using the REST API	192
Verifying the MP-BGP Route Reflector Configuration	193
Creating OSPF External Routed Network for Management Tenant Using REST API	193
Deploying an Application Policy	194
Three-Tier Application Deployment	194
Parameters to Create a Filter for http	195
Parameters to Create Filters for rmi and sql	196
Example Application Profile Database	196
Deploying an Application Profile Using the REST API	196



Preface

This preface includes the following sections:

- [Audience, page xv](#)
- [Document Conventions, page xv](#)
- [Related Documentation, page xvii](#)
- [Documentation Feedback, page xviii](#)
- [Obtaining Documentation and Submitting a Service Request, page xviii](#)

Audience

This guide is intended primarily for data center administrators with responsibilities and expertise in one or more of the following:

- Virtual machine installation and administration
- Server administration
- Switch and network administration

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).

Convention	Description
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
<i>variable</i>	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<code>boldface screen font</code>	Information you must enter is in boldface screen font.
<i><code>italic screen font</code></i>	Arguments for which you supply values are in italic screen font.
<>	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

**Warning****IMPORTANT SAFETY INSTRUCTIONS**

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Application Policy Infrastructure Controller (APIC) Documentation

The following companion guides provide documentation for APIC:

- *Cisco APIC Getting Started Guide*
- *Cisco APIC Basic Configuration Guide*
- *Cisco ACI Fundamentals*
- *Cisco APIC Layer 2 Networking Configuration Guide*
- *Cisco APIC Layer 3 Networking Configuration Guide*
- *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*
- *Cisco APIC REST API Configuration Guide*
- *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*
- *Cisco ACI Virtualization Guide*
- *Cisco Application Centric Infrastructure Best Practices Guide*

All these documents are available at the following URL: <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>

Cisco Application Centric Infrastructure (ACI) Documentation

The broader ACI documentation is available at the following URL: <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>.

Cisco Application Centric Infrastructure (ACI) Simulator Documentation

The Cisco ACI Simulator documentation is available at <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-centric-infrastructure-simulator/tsd-products-support-series-home.html>.

Cisco Nexus 9000 Series Switches Documentation

The Cisco Nexus 9000 Series Switches documentation is available at <http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html>.

Cisco Application Virtual Switch Documentation

The Cisco Application Virtual Switch (AVS) documentation is available at <http://www.cisco.com/c/en/us/support/switches/application-virtual-switch/tsd-products-support-series-home.html>.

Cisco Application Centric Infrastructure (ACI) Integration with OpenStack Documentation

Cisco ACI integration with OpenStack documentation is available at <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: <http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.



New and Changed Information

This chapter contains the following sections:

- [New and Changed Information, page 1](#)

New and Changed Information

The following table provides an overview of the significant changes to this guide up to this current release. The table does 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 Behavior in Cisco APIC and Document Reorganization

Cisco APIC Release Version	Feature	Description	Where Documented
Release 1.3(x)	No significant changes	--	--
Release 1.2(2g)	--	Support for IPv6 management address provisioning or through a policy on the APIC controller.	This content is available in the Initial POD Setup and Overview, on page 3 and with in-band and out-of-band configuration examples. Additional support for DNS and NTP policies.

Cisco APIC Release Version	Feature	Description	Where Documented
Release 1.2(1i)	Basic GUI	Introduced the Basic GUI mode in the APIC GUI	This content is available in the chapter Using the Basic GUI , on page 31.
Release 1.2(1i)	NX-OS style CLI	Introduced the NX-OS style CLI in the release.	This content is available in Using the NX-OS Style CLI , on page 65



Initial POD Setup and Overview

This chapter contains the following sections:

- [First-Time Access, page 3](#)
- [Initializing the Fabric, page 24](#)
- [Switch Discovery with the APIC, page 26](#)

First-Time Access

Getting Started Guide Contents

The *Cisco APIC Getting Started Guide* contains the following information:

- Setting up an initial pod environment
- Setting up a multipod environment

For detailed information about configuring the APIC, see the *Cisco APIC Basic Configuration Guide*. Also, see the following guides for details about installation: *Cisco ACI Fabric Hardware Installation Guide* and *Cisco APIC Installation Guide*.

Simplified Approach to Configuring in Cisco APIC

Cisco APIC supports a simplified approach to configuring the ACI with the choice of two additional user interfaces. They are the NX-OS style CLI and the Basic GUI. The existing methods of configuration using REST API and Advanced GUI are supported as well. The Advanced GUI is equivalent to the GUI of the previous releases. Cisco recommends that you use the Advanced GUI to manage any policy that you created in Release 1.2 or earlier releases.

In addition to the simple approach available for network administrators and other users of the NX-OS style CLI and the Basic GUI, there is intelligence embedded in these approaches as compared to the Advanced GUI or the REST API. In several instances, the NX-OS style CLI and the Basic GUI often 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.

Configurations using NX-OS style CLI and Basic GUI are compatible similar to the compatibility between existing methods of configuration using Advanced GUI and REST API. 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*.

Installing the Cisco Application Centric Infrastructure Fabric Hardware

For details about installing the ACI fabric hardware, see the *Application Centric Infrastructure Fabric Hardware Installation 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:

Procedure

- 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 APIC

When the APIC is launched for the first time, the 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 you are using a Cisco Integrated Management Controller (CIMC) for your setup, use only the port-side utility console port with the breakout cable. Setup the CIMC first, and then access the APIC through the CIMC KVM or continue to access the APIC locally through the port-side utility console port. Do not use the RJ-45 console port, unless access to the port side is restricted. If you choose the CIMC KVM access, you will have remote access available later which is required during operations.
- If you are using RJ-45 console port, connect to CIMC using SSH and enable the Serial over LAN port using the following parameters:
 - Scope SOL sol
 - Set Enabled to Yes
 - Commit
 - Exit

After enabling, enter the command **connect host** to access the console. If the serial port is connected, either disconnect the serial port or ensure that the connected device has the proper configuration.

- It is recommended not to modify any parameters using CIMC. If there are any issues, ensure that the default setting for CIMC management node is **Dedicated Mode** and not **Shared**. If **Dedicated Mode** is not used, it can prevent the discovery of fabric nodes.
- Do not upgrade software or firmware using the CIMC user interface, XML, or SSH interfaces unless the modified property and software or firmware version are supported with your specific 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

- Starting with APIC release 1.2(2x), 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 APIC and 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 APIC (the controller cluster) to the ACI fabric requires a 10G interface on the ACI leaf. You cannot connect the APIC directly to the N9332PQ ACI leaf switch, unless you use a 40G to 10G converter (part number CVR-QSFP-SFP10G), in which case the port on the N9332PQ switch will auto-negotiate to 10G without requiring any manual configuration.
- The fabric ID is set during the APIC controller setup and it cannot be changed unless you perform a clean reload of the fabric. To change the fabric ID, perform a clean reload on the APIC and leaf switches after changing the `sam.config` file. You must have separate fabric IDs if you want to connect two ACI fabric domains using Layer 2 configuration links. This follows the dual-fabric design.

About High Availability for APIC Cluster

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

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

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 APIC in an active-standby mode, you must have at least 3 active APICs in a cluster.
POD ID	POD ID	1
Standby controller	Setup standby controller	NO
Controller ID	Unique ID number for the active APIC instance.	Valid range: 1-19
Controller name	Active controller name	apic1

Name	Description	Default Value
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 APIC cluster is /23. If you are using Release 2.0(1) the minimum is /22.
VLAN ID for infrastructure network ¹	Infrastructure VLAN for APIC-to-switch communication including virtual switches Note Reserve this VLAN for 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.	--
IP address pool for bridge domain multicast address (GIPO)	IP addresses used for fabric multicast	225.0.0.0/15 Valid range: 225.0.0.0/15 to 231.254.0.0/15, prefixlen must be 15 (128k IPs)
IPv4/IPv6 addresses for the out-of-band management	IP address that you use to access the 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	—

Name	Description	Default Value
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
Number of active controllers	Cluster size	3 Note When setting up APIC in an active-standby mode, you must have at least 3 active 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 APIC instance .	Recommended range: >20
Controller name	Standby controller name	NA

Name	Description	Default Value
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 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 APIC-to-switch communication including virtual switches Note Reserve this VLAN for 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.	--
IPv4/IPv6 addresses for the out-of-band management	IP address that you use to access the 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	—

Name	Description	Default Value
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*.

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): 4093
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: 4093
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): 4093
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 [0:0:0:0:0: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

Procedure

Step 1 Open one of the supported browsers:

- Chrome version 35 (at minimum)
- Firefox version 26 (at minimum)
- Internet Explorer version 11 (at minimum)
- Safari version 7.0.3 (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.

Step 5 In the **Mode** field, from the drop-down list, choose the **Advanced** or the **Basic** mode as desired.

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

Procedure

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 Object Model CLI



Note From Cisco APIC Release 1.0 until 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.

Procedure

- 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. With Cisco APIC Releases 1.0 and 1.1, you are now in the object model CLI. With Cisco APIC Release 1.2, you are now in the NX-OS style CLI for APIC.
- Step 3** With Cisco APIC Release 1.2, type **bash** to enter the object model CLI. 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*.

Accessing the NX-OS Style CLI



Note

From Cisco APIC Release 1.0 until 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.

Procedure

-
- 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. From this level, you can reach these configuration modes:

- To continue in the NX-OS style CLI, you can stay in EXEC mode or you can type **configure** to enter global configuration mode.

For information about NX-OS style CLI commands, see the *Cisco APIC NX-OS Style CLI Command Reference*.

- To reach the object model CLI, type **bash**.

For information about object mode CLI commands, see the *Cisco APIC Command-Line Interface User Guide, APIC Releases 1.0 and 1.1*.

Overview of the GUI

The APIC GUI is a browser-based graphical interface to the APIC that communicates internally with the APIC engine by exchanging REST API messages. The GUI contains several areas and panes.

Deployment Warning and Policy Usage Information

When you first log in to the APIC GUI, the **Deployment Warning Settings** dialog box opens allowing you to enable and alter the scope of deployment notification that displays policy usage information. The deployment warning settings can also be accessed from the **welcome, <login_name>** drop-down list (Change Deployment Settings) and through a button on the **Policy Usage Information** dialog box.

The policy usage information allows users to identify which resources and policies are being used by the policy that the user is currently modifying or deleting. The 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 the user attempts 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.

Toggling Between Basic and Advanced GUI Modes

When logged in to the APIC GUI, you can verify the GUI mode you are in. The mode you have entered is displayed in the top right corner of the GUI. You can choose to operate in one of two modes:

Caution: Cisco recommends that you do not mix configuration modes (Advanced or Basic). When you make a configuration in either mode and change the configuration using the other mode, unintended changes can occur. For example, if you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.

- **Basic Mode**—For information about tasks that you perform in Basic Mode, see the chapter, *Getting Started with APIC Using the Basic GUI*.
- **Advanced Mode**—For information about tasks that you perform in Advanced Mode, see the chapter, *Getting Started with APIC Using the Advanced GUI*.

You can also change from one GUI mode to another or toggle between modes as follows:

- 1 In the GUI, click the **welcome, <login_name>** drop-down list and choose **Toggle GUI Mode**.
- 2 In the **Warning** dialog box, click **Yes**.
- 3 Wait for the application to complete loading and display the GUI in the changed mode.

**Caution**

Changes made through the APIC Basic GUI can be seen, but cannot be modified in the Advanced GUI, and changes made in the Advanced GUI cannot be rendered in the Basic GUI. The Basic GUI is kept synchronized with the NX-OS style CLI, so that if you make a change from the NX-OS style CLI, these changes are rendered in the Basic GUI, and changes made in the Basic GUI are rendered in the NX-OS style CLI, but the same synchronization does not occur between the Advanced GUI and the NX-OS style CLI. See the following examples:

- Do not mix Basic and Advanced GUI modes. If you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.
- Do not mix the Advanced GUI and the CLI, when doing per-interface configuration on APIC. Configurations performed in the GUI, may only partially work in the NX-OS CLI.

For example, if you configure a switch port in the GUI at **Tenants > tenant-name > Application Profiles > application-profile-name > Application EPGs > EPG-name > Static Ports > Deploy Static EPG on PC, VPC, or Interface**

Then you use the show running-config command in the NX-OS style CLI, you receive output such as:

```
leaf 102
interface ethernet 1/15
switchport trunk allowed vlan 201 tenant t1 application ap1 epg ep1
exit
exit
```

If you use these commands to configure a static port in the NX-OS style CLI, the following error occurs:

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/15
apic1(config-leaf-if)# switchport trunk allowed vlan 201 tenant t1 application ap1
epg ep1
No vlan-domain associated to node 102 interface ethernet1/15 encap vlan-201
```

This occurs because the CLI has validations that are not performed by the APIC GUI. For the commands from the show running-config command to function in the NX-OS CLI, a vlan-domain must have been previously configured. The order of configuration is not enforced in the GUI.

- Do not make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI. This may also inadvertently cause objects to be created (with names prepended with `_ui_`) which cannot be changed or deleted in the Advanced GUI.

For the steps to remove such objects, see *Troubleshooting Unwanted _ui_ Objects* in the *APIC Troubleshooting Guide*.

Menu Bar and Submenu Bar

The menu bar and the submenu bar contain the following items:

The menu bar is displayed across the top of the APIC GUI (see the following figure). It provides access to the main tabs.

Figure 1: APIC GUI Menu Bar



You can navigate to the submenu bar (see the following figure) by clicking on one of the tabs in the menu bar. When you click on a menu bar tab, the submenu bar for that tab is displayed. The submenu bar is different for each menu bar tab and might also differ depending upon your specific configurations.

Figure 2: APIC GUI Submenu Bar



Submenu Bar

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.

TENANTS Tab

Use the **Tenants** tab in the menu bar to perform tenant management. In the submenu bar, you see an **Add Tenant** link, and a drop-down list that contains all the tenants. Up to five of the most recently used tenants are also displayed on the submenu bar.

- 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.

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.

VM NETWORKING Tab

Use the **VM 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. You can view existing service nodes in the **Inventory** submenu tab.

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.

Search Icon

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

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. When you select a component in the **Navigation** pane, the object displays in the **Work** pane.

**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. See the following figure for an example view of the **Work** 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.

Figure 3: Example View of APIC Work Pane





GUI Icons

Table 4: Frequently Displayed Icons in the APIC GUI

Icons	Description
	Control arrow for Navigation pane display
	Displays online help information
	Quickstart information
	Downloads the table as an XML file
	Displays the table view
	Displays the table view of the component that you chose in the Navigation pane
	Refreshes the context of the panel. Click this icon only when there is a connection problem, because the data is updated whenever the repository changes.
	Settings
	Next view
	Previous view
	Show path
	Clear path

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.

API Inspector

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.

Procedure

-
- Step 1** Log in to the APIC GUI.
 - Step 2** In the upper right corner of the APIC window, click the "welcome, <name>" message 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.

Name	Description
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.

Name	Description
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":[]}
```

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.


Example Topology

An example 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 4: Example Fabric Topology



Example Topology Connections

An example topology with connection details 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)

Name	Connection Details
spine2	eth5/1 = leaf2 (eth1/49) eth5/2 = leaf1 (eth1/50)

Switch Discovery with the APIC

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


Note

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>.

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


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.

Registering the Unregistered Switches 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.

Procedure

- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
- Step 2** In the **Navigation** pane, click **Fabric Membership**.
In the **Work** pane, in the **Fabric Membership** table, a single leaf switch is displayed with an ID of 0. It is the leaf switch that is connected to apic1.
- Step 3** Configure the ID by double-clicking the leaf switch row, and performing the following actions:
- In the **ID** field, add the appropriate ID (leaf1 is ID 101, and leaf 2 is ID 102).
The ID must be a number that is greater than 100 because the first 100 IDs are for APIC appliance nodes.
 - In the **Switch Name** field, add the name of the switch, and click **Update**.
Note After an ID is assigned, it cannot be updated. The switch name can be updated by double-clicking the name and updating the **Switch Name** field.
An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod.
- Step 4** Monitor the **Work** pane until one or more spine switches appear.
- Step 5** Configure the ID by double-clicking the spine switch row, and perform the following actions:
- In the **ID** field, add the appropriate ID (spine1 is ID 203 and spine 2 is ID 204).
Note It is recommended that leaf nodes and spine nodes be numbered differently. For example, number spines in the 200 range and number leaves in the 100 range.
 - In the **Switch Name** field, add the name of the switch, and click **Update**.
An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod. Wait until all remaining switches appear in the **Node Configurations** table before you go to the next step.
- Step 6** For each switch listed in the **Fabric Membership** table, perform the following steps:
- Double-click the switch, enter an **ID** and a **Name**, and click **Update**.
 - Repeat for the next switch in the list.
-

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

Procedure

- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
- Step 2** In the **Navigation** pane, expand **Fabric Membership**.

The switches in the fabric are displayed with their node IDs. In the **Work** pane, 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

Procedure

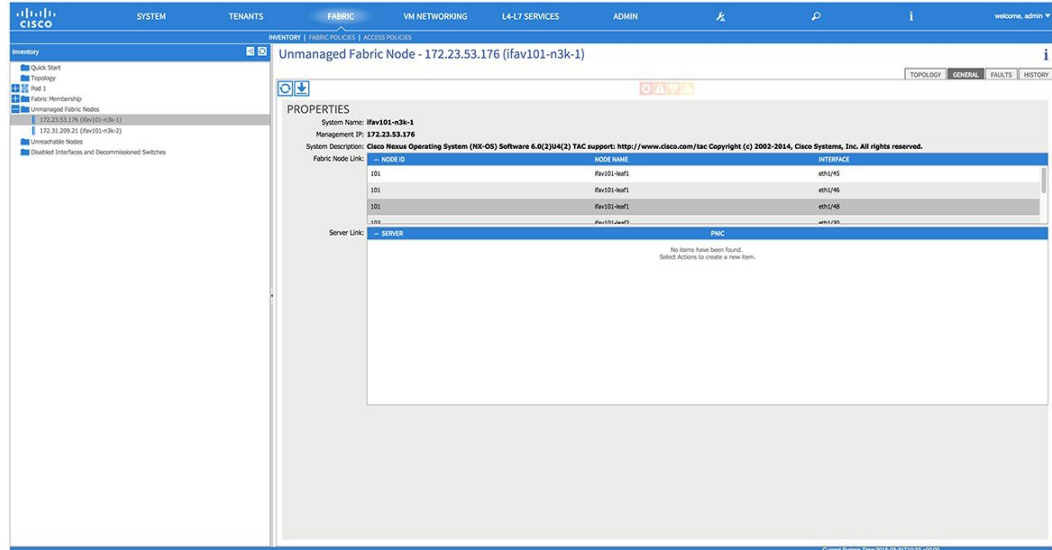
- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
 - Step 2** In the **Navigation** pane, choose the pod that you want to view.
 - Step 3** In the **Work** pane, click the **TOPOLOGY** tab.
The displayed diagram shows all attached switches, APIC instances, and links.
 - Step 4** (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
To return to the topology diagram, in the upper left corner of the **Work** pane, click the **Previous View** icon.
 - Step 5** (Optional) To refresh the topology diagram, in the upper left corner of the **Work** pane, click the **Refresh** icon.
-

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) or Cisco Discovery Protocol (CDP) 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. The following figure shows the APIC GUI displaying unmanaged switches in the **Fabric > Inventory** view.

Figure 5: Unmanaged Layer 2 Switches in the APIC Fabric Inventory



304443



Using the Basic GUI

This chapter contains the following sections:

- [Toggling Between Basic and Advanced GUI Modes, page 31](#)
- [About Getting Started with APIC Examples, page 32](#)
- [About Switch Discovery with the APIC, page 33](#)
- [Configuring Network Time Protocol, page 36](#)
- [Creating User Accounts, page 38](#)
- [Adding Management Access in the GUI, page 43](#)
- [Configuring a VMM Domain, page 46](#)
- [Creating Tenants, VRF, and Bridge Domains, page 51](#)
- [Configuring Server or Service Policies, page 52](#)
- [Configuring External Connectivity for Tenants, page 59](#)
- [Deploying an Application Policy, page 62](#)

Toggling Between Basic and Advanced GUI Modes

When logged in to the APIC GUI, you can verify the GUI mode you are in. The mode you have entered is displayed in the top right corner of the GUI. You can choose to operate in one of two modes:

Caution: Cisco recommends that you do not mix configuration modes (Advanced or Basic). When you make a configuration in either mode and change the configuration using the other mode, unintended changes can occur. For example, if you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.

- **Basic Mode**—For information about tasks that you perform in Basic Mode, see the chapter, *Getting Started with APIC Using the Basic GUI*.
- **Advanced Mode**—For information about tasks that you perform in Advanced Mode, see the chapter, *Getting Started with APIC Using the Advanced GUI*.

You can also change from one GUI mode to another or toggle between modes as follows:

- 1 In the GUI, click the **welcome**, **<login_name>** drop-down list and choose Toggle GUI Mode.
- 2 In the **Warning** dialog box, click Yes .
- 3 Wait for the application to complete loading and display the GUI in the changed mode.

**Caution**

Changes made through the APIC Basic GUI can be seen, but cannot be modified in the Advanced GUI, and changes made in the Advanced GUI cannot be rendered in the Basic GUI. The Basic GUI is kept synchronized with the NX-OS style CLI, so that if you make a change from the NX-OS style CLI, these changes are rendered in the Basic GUI, and changes made in the Basic GUI are rendered in the NX-OS style CLI, but the same synchronization does not occur between the Advanced GUI and the NX-OS style CLI. See the following examples:

- Do not mix Basic and Advanced GUI modes. If you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.
- Do not mix the Advanced GUI and the CLI, when doing per-interface configuration on APIC. Configurations performed in the GUI, may only partially work in the NX-OS CLI.

For example, if you configure a switch port in the GUI at **Tenants > tenant-name > Application Profiles > application-profile-name > Application EPGs > EPG-name > Static Ports > Deploy Static EPG on PC, VPC, or Interface**

Then you use the `show running-config` command in the NX-OS style CLI, you receive output such as:

```
leaf 102
interface ethernet 1/15
switchport trunk allowed vlan 201 tenant t1 application ap1 epg ep1
exit
exit
```

If you use these commands to configure a static port in the NX-OS style CLI, the following error occurs:

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/15
apic1(config-leaf-if)# switchport trunk allowed vlan 201 tenant t1 application ap1
epg ep1
No vlan-domain associated to node 102 interface ethernet1/15 encaps vlan-201
```

This occurs because the CLI has validations that are not performed by the APIC GUI. For the commands from the `show running-config` command to function in the NX-OS CLI, a `vlan-domain` must have been previously configured. The order of configuration is not enforced in the GUI.

- Do not make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI. This may also inadvertently cause objects to be created (with names prepended with `_ui_`) which cannot be changed or deleted in the Advanced GUI.

For the steps to remove such objects, see *Troubleshooting Unwanted _ui_ Objects* in the *APIC Troubleshooting Guide*.

About Getting Started with APIC Examples

The steps in several examples in this guide include a parameter name. These parameter names are provided as examples for convenience and ease of your understanding, and it is not required for you to use them.

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

**Note**

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>.

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

**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.

Registering the Unregistered Switches 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.

Procedure

- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
- Step 2** In the **Navigation** pane, click **Fabric Membership**.

In the **Work** pane, in the **Fabric Membership** table, a single leaf switch is displayed with an ID of 0. It is the leaf switch that is connected to apic1.

Step 3 Configure the ID by double-clicking the leaf switch row, and performing the following actions:

- a) In the **ID** field, add the appropriate ID (leaf1 is ID 101, and leaf 2 is ID 102).
The ID must be a number that is greater than 100 because the first 100 IDs are for APIC appliance nodes.
- b) In the **Switch Name** field, add the name of the switch, and click **Update**.
Note After an ID is assigned, it cannot be updated. The switch name can be updated by double-clicking the name and updating the **Switch Name** field.

An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod.

Step 4 Monitor the **Work** pane until one or more spine switches appear.

Step 5 Configure the ID by double-clicking the spine switch row, and perform the following actions:

- a) In the **ID** field, add the appropriate ID (spine1 is ID 203 and spine 2 is ID 204).
Note It is recommended that leaf nodes and spine nodes be numbered differently. For example, number spines in the 200 range and number leaves in the 100 range.
- b) In the **Switch Name** field, add the name of the switch, and click **Update**.

An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod. Wait until all remaining switches appear in the **Node Configurations** table before you go to the next step.

Step 6 For each switch listed in the **Fabric Membership** table, perform the following steps:

- a) Double-click the switch, enter an **ID** and a **Name**, and click **Update**.
- b) Repeat for the next switch in the list.

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

Procedure

Step 1 On the menu bar, choose **FABRIC > INVENTORY**.

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

The switches in the fabric are displayed with their node IDs. In the **Work** pane, 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

Procedure


- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
 - Step 2** In the **Navigation** pane, choose the pod that you want to view.
 - Step 3** In the **Work** pane, click the **TOPOLOGY** tab.
The displayed diagram shows all attached switches, APIC instances, and links.
 - Step 4** (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
To return to the topology diagram, in the upper left corner of the **Work** pane, click the **Previous View** icon.
 - Step 5** (Optional) To refresh the topology diagram, in the upper left corner of the **Work** pane, click the **Refresh** icon.
-

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) or Cisco Discovery Protocol (CDP) 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. The following figure shows the APIC GUI displaying unmanaged switches in the **Fabric > Inventory** view.

Figure 6: Unmanaged Layer 2 Switches in the APIC Fabric Inventory



Configuring Network Time Protocol

Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

In-Band and Out-of-Band Management NTP



Note

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
 - See the Adding Management Access section in this guide for information about in-band management access and out-of-band management access.
-
- Out-of-band management NTP—When an ACI fabric is deployed with out-of-band management, each node of the fabric, inclusive of spines, leaves, and all members of the APIC cluster, is managed from outside the ACI fabric. This IP reachability will be leveraged so that each node can individually query the same NTP server as a consistent clock source. To configure NTP, a Date and Time policy must be created that references an out-of-band management endpoint group. Date and Time policies are confined to a single pod and must be deployed across all pods provisioned in the ACI fabric. Currently only one pod per ACI fabric is allowed.
 - In-Band Management NTP—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication. The steps used to configure in-band management policies are identical to those used to establish an out-of-band management policy. The distinction is around how to allow the fabric to connect to the NTP server.

NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The `gai.conf` can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

Configuring NTP Using the Basic GUI

Before You Begin

Procedure

- Step 1** On the menu bar, choose **System > System Settings**.
- Step 2** In the **Navigation** pane, click **NTP**.
- Step 3** In the **Work** pane, the default NTP policy properties are displayed.
- Step 4** In the NTP Servers field, expand the + sign to display the **Create Providers** dialog box.
- Step 5** In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.
 - If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.

- In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.
-

Verifying NTP Policy Deployed to Each Node Using the CLI

Procedure

Step 1 SSH to an APIC in the fabric.

Step 2 Press the Tab key two times after entering the attach command to list all the available node names:

Example:

```
admin@apic1:~> attach <Tab> <Tab>
```

Step 3 Log in one of the nodes using the same password that you used to access the APIC.

Example:

```
admin@apic1:~> attach node_name
```

Step 4 View the NTP peer status.

Example:

```
leaf-1# show ntp peer-status
```

A reachable NTP server has its IP address prefixed by an asterisk (*), and the delay is a non-zero value.

Step 5 Repeat steps 3 and 4 to verify each node in the fabric.

Creating User Accounts

Configuring a Local User

In the initial configuration script, the admin account is configured and the admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

Configuring a Remote User

Instead of configuring local users, you can point the APIC at the centralized enterprise credential datacenter. The APIC supports Lightweight Directory Access Protocol (LDAP), active directory, RADIUS, and TACACS+.

**Note**

When an APIC is in minority (disconnected from the cluster), remote logins can fail because the ACI is a distributed system and the user information is distributed across APICS. Local logins, however, continue to work because they are local to the APIC.

To configure a remote user authenticated through an external authentication provider, you must meet the following prerequisites:

- The DNS configuration should have already been resolved with the hostname of the RADIUS server.
- You must configure the management subnet.

Configuring a Local User Using the GUI

Before You Begin

- The ACI fabric is installed, APIC controllers are online, and the APIC cluster is formed and healthy.
- As appropriate, the security domain(s) that the user will access are defined. For example, if the new user account will be restricted to accessing a tenant, the tenant domain is tagged accordingly.
- An APIC user account is available that will enable the following:
 - Creating the TACACS+ and TACACS+ provider group.
 - Creating the local user account in the target security domain(s). If the target domain is `all`, the login account used to create the new local user must be a fabric-wide administrator that has access to `all`. If the target domain is a tenant, the login account used to create the new local user must be a tenant administrator that has full read write access rights to the target tenant domain.

Procedure

- Step 1** On the menu bar, choose **ADMIN > AAA**.
- Step 2** In the **Navigation** pane, click **AAA Authentication**.
- Step 3** In the **Work** pane, verify that in the default **Authentication** field, the **Realm** field displays as Local.
- Step 4** In the **Navigation** pane, expand **Security Management > Local Users**.
The admin user is present by default.
- Step 5** In the **Navigation** pane, right-click **Create Local User**.
- Step 6** In the **User Identity** dialog box, enter a **Login ID** and **Password** for the user, and click **Next**.
- Step 7** In the **Security** dialog box, choose the desired security domain for the user, and click **Next**.
- Step 8** In the **Roles** dialog box, click the radio buttons to choose the roles for your user, and click **Next**.
You can provide read-only or read/write privileges.
- Step 9** In the **User Identity** dialog box, perform the following actions:
 - a) In the **Login ID** field, add an ID.
 - b) In the **Password** field, enter the password.
At the time a user sets their password, the APIC validates it against the following criteria:

- Minimum password length is 8 characters.
- Maximum password length is 64 characters.
- Has fewer than three consecutive repeated characters.
- Must have characters from at least three of the following characters types: lowercase, uppercase, digit, symbol.
- Does not use easily guessed passwords.
- Cannot be the username or the reverse of the username.
- Cannot be any variation of cisco, isco or any permutation of these characters or variants obtained by changing the capitalization of letters therein.

c) In the **Confirm Password** field, confirm the password.

d) Click **Finish**.

Step 10 In the **Navigation** pane, click the name of the user that you created. In the **Work** pane, expand the + sign next to your user in the **Security Domains** area.

The access privileges for your user are displayed.

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

To configure a Cisco AV Pair on an external authentication server, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair format is as follows:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(16003)
```

The first av-pair format has no UNIX user ID, while the second one does. Both are correct if all remote users have the same role and mutual file access is acceptable. If the UNIX user ID is not specified, ID 23999 is applied by the APIC system, and more than one role/read privilege is specified to any AV Pair user. This can cause users to have higher or lower permissions than configured through the group settings.



Note

The APIC Cisco AV-pair format is compatible and can co-exist with other Cisco AV-pair formats. APIC will pick up the first matching AV-pair from all the AV-pairs.

The APIC supports the following regexes:

```
shell:domains\s*[:]\s*((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0,31}) (\\ (\\d+\\)) $
shell:domains\s*[:]\s*((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0,31}) $
```

Examples:

- Example 1: A Cisco AV Pair that contains a single Login domain with only writeRoles:

```
shell:domains=domainA/writeRole1|writeRole2/
```

- Example 2: A Cisco AV Pair that contains a single Login domain with only readRoles:

```
shell:domains=domainA//readRole1|readRole2
```


Note

The "/" character is a separator between writeRoles and readRoles per Login domain and is required even if only one type of role is to be used.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

```
aaa-network-admin Cleartext-Password := "<password>"
Cisco-avpair = "shell:domains = all/aaa/read-all(16001)"
```

Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs

Procedure

- Step 1** On the menu bar, click **ADMIN > AAA**.
- Step 2** In the **Navigation** pane, click **AAA Authentication**.
- Step 3** In the **Work** pane, in the **Properties** area, from the **Remote user login policy** drop-down list, choose **Assign Default Role**.

The default value is **No Login**. The **Assign Default Role** option assigns the minimal read-only privileges to users that have missing or bad Cisco AV Pairs. Bad AV Pairs are those AV Pairs that fail the parsing rules.

Best Practice for Assigning AV Pairs

As best practice, Cisco recommends that you assign unique UNIX user ids in the range 16000-23999 for the AV Pairs that are assigned to users when in bash shell (using SSH, Telnet or Serial/KVM consoles). If a situation arises when 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 remote users with a UNIX ID of 23999.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is logged in using Secure Shell (SSH) or Telnet.

Procedure

Configure an AV pair on the external authentication server.

The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

```
* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2

* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2 (8101)
```

These are the boost regexes supported by APIC:

```
uid_regex("shell:domains\\s*[:]\\s*((\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31}) (\\(\\d+\\))$");
regex("shell:domains\\s*[:]\\s*((\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31})$");
```

The following is an example:

```
shell:domains = coke/tenant-admin/read-all,pepsi//read-all (16001)
```

Configuring a Remote User Using the GUI

Before You Begin

- The DNS configuration must have resolved the RADIUS server hostname in order for the fabric controller to reach the server.
- The APIC should have the external management subnet policy configured so that it is able to reach the RADIUS server.

Procedure

-
- Step 1** On the menu bar, choose **ADMIN > AAA**. In the **Navigation** pane, expand **RADIUS Management**.
 - Step 2** Right-click **RADIUS Providers**, and click **Create RADIUS Provider**.
 - Step 3** In the **Create RADIUS Provider** dialog box, and perform the following actions:
 - a) In the **Host Name (or IP Address)** field, add the hostname.
 - b) In the **Authorization Port** field, add the port number required for authorization. This number depends on the RADIUS server configured.
 - c) Click the required **Authorization Protocol** radio button.
 - d) In the **Key** and **Confirm Key** fields, enter the preshared key. This key is the same information that is shared with the server key configured on the RADIUS server.
 - Step 4** In the **Navigation** pane, under **RADIUS Providers**, click the RADIUS provider that you created. Details about the configurations for the RADIUS provider are displayed in the **Work** pane.
 - Step 5** In the **Navigation** pane, right-click **RADIUS Provider Groups**, and click **Create RADIUS Provider Group**.
 - Step 6** In the **Create RADIUS Provider Group** dialog box, perform the following actions:
 - a) In the **Name** field, enter a name.
 - b) Expand the **Providers** field, and from the **Name** field drop-down list, choose the provider created earlier.
 - c) In the **Priority** field, assign a priority. Click **Update**, and click **Submit**.

The radius provider group is created.

Step 7 In the **Navigation** pane, expand **AAA Authentication**, and right-click **Login Domain** to click **Create Login Domain**.

Step 8 In the **Create Login Domain** dialog box, perform the following actions:

- a) In the **Name** field, enter a domain name.
- b) In the **Realm** field drop-down list, choose the RADIUS realm.
- c) In the **RADIUS Provider Group** field drop-down list, choose the provider group that was created earlier. Click **Submit**.

The login domain is created and is now available for remote user login and configuration.

Adding Management Access in the GUI

An APIC controller has two routes to reach the management network, one is by using the in-band management interface and the other is by using the out-of-band management interface.

The in-band management network allows APIC to communicate with the leaf switches and with the outside using the ACI fabric, and it makes it possible for external management devices to communicate with the APIC or the leaf switches and spine switches using the fabric itself.

The out-of-band management network configuration defines the configuration of the management port on the controllers, the leaf switches and the spine switches.

The APIC controller always selects the in-band management interface over the out-of-band management interface, if the in-band management interface is configured. The out-of-band management interface is used only when the in-band management interface is not configured or if the destination address is on the same subnet as the out-of-band management subnet of the APIC. This behavior cannot be changed or reconfigured. The APIC management interface does not support an IPv6 address and cannot connect to an external IPv6 server through this interface.

The APIC out-of-band management connection link must be 1 Gbps.

IPv4/IPv6 Addresses and In-Band Policies

In-band management addresses can be provisioned on the APIC controller only through a policy (Postman REST API, NX-OS Style CLI, or GUI). Additionally, the in-band management addresses must be configured statically on each node.

IPv4/IPv6 Addresses in Out-of-Band Policies

Out-of-band management addresses can be provisioned on the APIC controller either at the time of bootstrap or by using a policy (Postman REST API, NX-OS Style CLI, GUI). Additionally, the out-of-band management addresses must be configured statically on each node or by specifying a range of addresses (IPv4/IPv6) to the entire cluster. IP addresses are randomly assigned from a range to the nodes in the cluster.

Configuring Management Access

Configuring In-Band Management Access Using the Basic GUI



Note IPv4 and IPv6 addresses are supported for in-band management access. IPv6 configurations are supported using static configurations (for both in-band and out-of-band). IPv4 and IPv6 dual in-band and out-of-band configurations are supported only through static configuration. For more information, see the KB article, *Configuring Static Management Access in Cisco APIC*.

Procedure

-
- Step 1** Login to the **Basic Mode** in the APIC GUI, and on the menu bar, click **System > In Band & Out Of Band**.
- Step 2** In the **Navigation** pane, choose **InBand Management Configuration**.
- Step 3** (Optional) In the **Encap** field, enter a new value to change the default VLAN that is used for in-band management, if desired.
- Step 4** Expand **Nodes** and perform the following actions:
- In the **Nodes** field, choose the appropriate node to associate the in-band address.
 - In the **IP address** field, enter the desired IPv4 or IPv6 address.
 - In the **Gateway** field, enter the desired IPv4 or IPv6 gateway address. Click **Submit**.

Note The default gateway IP address will be the pervasive gateway of the ACI fabric on the VRF for the inband management.
- Step 5** Click the **L2 Connectivity** tab, expand **Ports**, and perform the following actions:
- In the **Path** field, from the drop-down list, choose the port that is connected to a server for management or to the outside.
 - In the **Encap** field, specify a VLAN to use on this port.
- Step 6** Expand **Gateway IP Address for External Connectivity** and in the **IP address** fields, list the desired gateway IPv4 and Pv6 address for external connectivity.
- Step 7** Expand **ACLs**, and add the desired ports that you want to connect to the inband management network. Click **Submit**.
-

The in-band management access is now established.

Configuring Out-of-Band Management Access Using the Basic GUI



Note IPv4 and IPv6 addresses are supported for out-of-band management access.

Procedure

-
- Step 1** Log in to the **Basic Mode** of the APIC GUI, and on the menu bar, choose **System > In Band & Out Of Band**.
 - Step 2** In the **Navigation** pane, click **Out-of-Band EPG - default**.
 - Step 3** In the **Work** pane, under **Properties**, expand **Nodes** and associate the appropriate nodes with an IPv4 or IPv6 address and a default gateway. Click **Update**.
 - Step 4** Expand **Access Restrictions**, and add the list of desired external subnets that will communicate with the out-of-band management address of the ACI fabric nodes.
 - Step 5** Expand **ACLs** for external subnets and enter the appropriate information for the L4 ports that you want to allow for management of the ACI fabric nodes.
The out-of-band management access is now configured.
-

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

All IPv6 tables mirror the existing IP tables functionality, except for Network Address Translation (NAT).

Existing IP Tables

- 1 Earlier, every rule in the IPv6 tables were executed one at a time and a system call was made for every rule addition or deletion.
- 2 Whenever a new policy was added, rules were appended to the existing IP tables file and no extra modifications were done to the file.
- 3 When a new source port was configured in the out-of-band policy, it added source and destination rules with the same port number.

Modifications to IP Tables

- 1 When IP tables are created, they are first written into hash maps that are then written into intermediate file IP tables-new which are restored. When saved, a new IP tables file is created in the /etc/sysconfig/ folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.
- 2 When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (/etc/sysconfig/iptables-new) and saved.
- 3 It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.
- 4 When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.

```
-A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
```
- 5 When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.

**Note**

- If only IPv4 is enabled, do not configure an IPv6 policy.
- If only IPv6 is enabled, do not configure an IPv4 policy.
- If both IPv4 and IPv6 are enabled and a policy is added, it will be configured to both the versions . So when you add an IPv4 subnet, it will be added to IP tables and similarly an IPv6 subnet is added to IPv6 tables.

Configuring a VMM Domain

Configuring Virtual Machine Networking Policies

The APIC integrates with third-party VM manager (VMM) (for example, VMware vCenter and SCVMM) to extend the benefits of ACI to the virtualized infrastructure. The APIC enables the ACI policies inside the VMM system to be used by its administrator.

This section provides examples of VMM integration using VMware vCenter and vShield. For details about the different modes of Cisco ACI and VMM integration, see the *ACI Virtualization Guide*.

About the VM Manager

**Note**

Information about the necessary configuration of the APIC for integration with the vCenter is described here. For instructions about configuring the VMware components, see the VMware documentation.

The following are details of some VM manager terms:

- A VM controller is an external virtual machine management entity such as VMware vCenter, and the VMware vShield. The APIC communicates with the controller to publish network policies that are applied to virtual workloads. A VM controller administrator provides an APIC administrator with a VM controller authentication credential; multiple controllers of the same type can use the same credential.
- Credentials represent the authentication credentials to communicate with VM controllers. Multiple controllers can use the same credentials.
- A virtual machine mobility domain (vCenter mobility domain) is a grouping of VM controllers with similar networking policy requirements. This mandatory container holds one or more VM controllers with policies such as for a VLAN pool, server to network MTU policy, or server to network access LACP policy. When an endpoint group gets associated with a vCenter domain, network policies get pushed to all the VM controllers in the vCenter domain.
- A pool represents a range of traffic encapsulation identifiers (for example, VLAN IDs, VNIDs, and multicast addresses). A pool is a shared resource and can be consumed by multiple domains such as VMM and Layer 4 to Layer 7 services. A leaf switch does not support overlapping VLAN pools. You must not associate different overlapping VLAN pools with the VMM domain. The two types of VLAN-based pools are as follows:

- Dynamic pools—Managed internally by the APIC to allocate VLANs for endpoint groups (EPGs). A vCenter Domain can associate only to a dynamic pool.
- Static pools—The EPG has a relation to the domain, and the domain has a relation to the pool. The pool contains a range of encapsulated VLANs and VXLANs. For static EPG deployment, the user defines the interface and the encapsulation. The encapsulation must be within the range of a pool that is associated with a domain with which the EPG is associated.
- For a VMware vCenter to be deployed, it must operate in VLAN mode or VXLAN mode. A VMM domain must be associated with a VLAN pool and a vShield must be associated with the vCenter.

Prerequisites for Creating a VMM Domain Profile

To configure a VMM domain profile, you must meet the following prerequisites:

- All fabric nodes are discovered and configured.
- Inband (inb) or out-of-band (oob) management has been configured on the APIC.
- A Virtual Machine Manager (VMM) is installed, configured, and reachable through the inb/oob management network (for example, a vCenter).
- You have the administrator/root credentials to the VMM (for example vCenter).

**Note**

If you prefer not to use the vCenter admin/root credentials, you can create a custom user account with minimum required permissions. See [Custom User Account with Minimum VMware vCenter Privileges](#), on page 47 for a list of the required user privileges.

- A DNS policy for the APIC must be configured if you plan to reference the VMM by hostname rather than an IP address.

Custom User Account with Minimum VMware vCenter Privileges

This allows the APIC to send VMware API commands to vCenter to allow the creation of the DVS/AVS, creation of the VMK interface (AVS), publish port groups and relay all necessary alerts.

To configure the vCenter from Cisco APIC, your credentials must allow the following minimum set of privileges within the vCenter:

- **Alarms**

APIC creates two alarms on the folder. One for DVS and another for port-group. The alarm is raised when the EPG or Domain policy is deleted on APIC, but for port-group or DVS it cannot be deleted due to the VMs are attached.

- **Distributed Switch**

- **dvPort Group**

- **Folder**

- **Network**

APIC manages the network settings such as add or delete port-groups, setting host/DVS MTU, LLDP/CDP, LACP etc.

- **Host**

If you use AVS in addition to above, you need the Host privilege on the data center where APIC will create DVS.

- **Host.Configuration.Advanced settings**
- **Host.Local operations.Reconfigure virtual machine**
- **Host.Configuration.Network configuration**

This is needed for AVS and the auto-placement feature for virtual Layer 4 to Layer 7 Service VMs. For AVS, APIC creates VMK interface and places it in 'vtep' port-group which is used for OpFlex.

- **Virtual machine**

If you use Service Graph in addition to above, you need the Virtual machine privilege for the virtual appliances which will be used for Service Graph.

- **Virtual machine.Configuration.Modify device settings**
- **Virtual machine.Configuration.Settings**

Creating a VMM Domain Profile

In this section, examples of a VMM domain are vCenter domain.

Creating a vCenter Domain Profile Using the Basic GUI

Before You Begin

Before you create a VMM domain profile, you must establish connectivity to external network using in-band management network on the APIC.

Procedure

-
- Step 1** Login to the **Basic Mode** in the APIC GUI.
 - Step 2** On the menu bar, choose **VM NETWORKING > Inventory**.
 - Step 3** In the **Navigation** pane, right-click **VMware** and click **Create vCenter Domain**.
 - Step 4** In the **Create vCenter Domain** dialog box, in the **Virtual Switch Name** field, enter a **Name**.
 - Step 5** In the **Virtual Switch** field, verify that **VMware vSphere Distributed Switch** is selected.
 - Step 6** In the **VLAN Pool** drop-down list, choose **Create VLAN Pool**. In the **Create VLAN Pool** dialog box, perform the following actions:
 - Note** This step provides the VLAN range for all port groups and EPGs that will be created under this server.
 - a) Enter a **Name**.
 - b) In the **Allocation Mode** field, verify that **Dynamic Allocation** is selected.

- c) Expand **Encap Blocks** to add a VLAN block. In the **Create Ranges** dialog box, enter a VLAN range.
Note We recommend that you use a range of at least 200 VLAN numbers.
- d) Click **OK**, and click **Submit**.

Step 7 In the **Create vCenter Domain** dialog box, expand **vCenter** and perform the following tasks:

- a) In the **Add vCenter** dialog box, in the **Type** field, click the **vCenter** radio button.
- b) In the vCenter Controller **Host Name (or IP Address)** field, enter the name or IP address of your vCenter.
- c) In the **Datacenter** field, enter the data center as appropriate.
- d) In the **vCenter Credential Name** field, enter a name.
- e) In the **Username** field, enter a username.
The username must be a credential to log in as an administrator of the vCenter.
- f) In the **Password** field, enter the password and repeat the password in the **Confirm Password** field. Click **OK**, and click **Submit**.
The password must be a credential to log in as an administrator of the vCenter.

Step 8 On the menu bar, choose **FABRIC > Inventory**.

Step 9 In the **Navigation** pane, expand **Pod**, click on the **Configure** tab and perform the following actions:

- a) In the **Configure** pane, click on **Add Switches** and select the switch/switches to configure. Click **Add Selected**.
Note Use the **Command** button to select more than one switch.
- b) Click on the port numbers to associate them to the VMware and click on **Configure Interface**.
- c) In the **Configure Interface** pane, click on the **VLAN** tab.
- d) In the **VLAN** pane, expand **ESX And SCVMM**.
- e) In the **Name** field, choose the VMware that you have just created from the drop-down list. Click **Update** and **Apply Changes** to complete VMware configuration.

Step 10 Verify the new domain and profiles by performing the following actions:

- Note** To ensure that the controllers are operational after the policy has been submitted, the administrator of the vCenter must add the hosts to the distributed switch.
- a) On the menu bar, choose **VM Networking > Inventory**.
 - b) In the **Navigation** pane, expand **VMware**, and expand the vCenter domain name.
 - c) In the **Navigation** pane, click the controller names to verify that the controllers are online.
In the **Work** pane, the properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the server is established, and the inventory is available.
-

Creating a vCenter and a vShield Domain Profile Using the Basic GUI

Procedure

-
- Step 1** Login to the **Basic Mode** in the APIC GUI.
- Step 2** On the menu bar, choose **VM NETWORKING > Inventory**.
- Step 3** In the **Navigation** pane, right-click **VMware** and click **Create vCenter Domain**.
- Step 4** In the **Create vCenter Domain** dialog box, in the **Virtual Switch Name** field, enter a **Name**.
- Step 5** In the **Virtual Switch** field, verify that **VMware vSphere Distributed Switch** is selected.
- Step 6** This step provides the VLAN range for all port groups and EPGs that will be created under this server. In the **VLAN Pool** drop-down list, choose **Create VLAN Pool**. In the **Create VLAN Pool** dialog box, perform the following actions:
- Enter a **Name**.
 - In the **Allocation Mode** field, verify that **Dynamic Allocation** is selected.
 - Expand **Encap Blocks** to add a VLAN block. In the **Create Ranges** dialog box, enter a VLAN range.

Note We recommend that you use a range of at least 200 VLAN numbers.
 - Click **OK**, and click **Submit**.
- Step 7** Expand **vCenter/vShield** and perform the following tasks:
- In the **Create vCenter/vShield Controller** dialog box, in the **Type** field, click the **vCenter + vShield** radio button.
 - In the vCenter Controller **Host Name (or IP Address)** field, enter the name or IP address of your vCenter.
 - In the **Datacenter** field, enter the data center as appropriate.
 - In the **vCenter Credential Name** field, enter a name.
 - In the **Username** field, enter a username.
The username must be a credential to log in as an administrator of the vCenter.
 - In the **Password** field, enter the password and repeat the password in the **Confirm Password** field.
The password must be a credential to log in as an administrator of the vCenter.
 - In the vShield Controller **Host Name (or IP Address)** field, enter the name or IP address of your vShield.
 - In the **vCenter Credential Name** field, enter a name.
 - In the **Datacenter** field, enter the data center as appropriate.
 - In the **Username** field, enter a username.
The username must be a credential to log in as an administrator of the vShield.
 - In the **Password** field, enter the password and repeat the password in the **Confirm Password** field.
The password must be a credential to log in as an administrator of the vShield.
 - Click **OK**, and click **Submit**.
- Step 8** On the menu bar, choose **FABRIC > Inventory**.
- Step 9** In the **Navigation** pane, expand **Pod**, click on the **Configure** tab and perform the following actions:
- In the **Configure** pane, click on **Add Switches** and select the switch/switches to configure. Click **Add Selected**.

Note Use the **Command** button to select more than one switch.
 - Click on the port numbers to associate them to the VMware and click on **Configure Interface**.
 - In the **Configure Interface** pane, click on the **VLAN** tab.

- d) In the **VLAN** pane, expand **ESX And SCVMM**.
- e) In the **Name** field, choose the VMware that you have just created from the drop-down list. Click **Update** and **Apply Changes** to complete VMware configuration.

Step 10 Verify the new domain and profiles by performing the following actions:

To ensure that the controllers are operational after the policy has been submitted, the administrator of the vCenter and vShield must add the hosts to the distributed switch.

- a) On the menu bar, choose **VM Networking > Inventory**.
- b) In the **Navigation** pane, expand **VMware**, and expand the vCenter domain name. Both the vCenter and vShield should be displayed in the VMware **Work** pane.
- c) In the **Navigation** pane, click the controller names to verify that the controllers are online. In the **Work** pane, the properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the server is established, and the inventory is available.

Creating Tenants, VRF, and Bridge Domains

Tenants Overview

- 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.

Tenant Creation

A tenant contains primary elements such as filters, contracts, bridge domains, and application profiles that you can create after you first create a tenant.

VRF and Bridge Domains

You can create and specify a VRF and a bridge domain for the tenant. The defined bridge domain element subnets reference a corresponding Layer 3 context.

For details about enabling IPv6 Neighbor Discovery see *IPv6 and Neighbor Discovery in Cisco APIC Layer 3 Networking Guide*.

Creating a Tenant, VRF, and Bridge Domain Using the Basic GUI

Procedure

-
- Step 1** Log in to the **Basic Mode** in the APIC GUI, and on the menu bar, click **TENANT > Add Tenant**.
- Step 2** In the **Create Tenant** dialog box, perform the following tasks:
- In the **Name** field, enter a name.
 - Click the **Security Domains +** icon to open the **Create Security Domain** dialog box.
 - In the **Name** field, enter a name for the security domain. Click **Submit**.
 - In the **Create Tenant** dialog box, check the check box for the security domain that you created, and click **Submit**.
- Step 3** In the **Navigation** pane, expand **Tenant-name > Networking**, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:
- In the **Name** field, enter a name.
 - Click **Submit** to complete the VRF configuration.
- Step 4** In the **Networking** pane, drag the **BD** icon to the canvas while connecting it to the **VRF** icon. In the **Create Bridge Domain** dialog box that displays, perform the following tasks:
- In the **Name** field, enter a name.
 - Expand **Subnets** to open the **Create Subnet** dialog box, enter the subnet mask in the **Gateway IP** field and click **OK**.
 - Click **Submit** to complete bridge domain configuration.
- Step 5** In the **Networking** pane, drag the **L3** icon down to the canvas while connecting it to the **VRF** icon. In the **Create Routed Outside** dialog box that displays, perform the following tasks:
- In the **Node ID** field, enter a node ID.
 - In the **Router ID** field, enter the router ID.
 - Expand **Static Routes** and enter the IPv4 or IPv6 addresses in the **IP Address** and the **Next Hop IP** fields and click **Update**.

Note The gateway IPv6 address must be a global unicast IPv6 address.
 - Click the **Protocols** box and select BGP, OSPF, and EIGRP for configuration as desired.
 - Click **OK** and then click **Submit** to complete Layer 3 configuration.
- To confirm L3 configuration, in the **Navigation** pane, expand **VRFs > VRF name > Deployed VRFs**.
-

Configuring Server or Service Policies

Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets. If the client is on an ESX hypervisor with a deployed vShield Domain profile, then the use of a DHCP relay policy configuration is mandatory.

When a vShield controller deploys a Virtual Extensible Local Area Network (VXLAN), the hypervisor hosts create a kernel (vxn, virtual tunnel end-point [VTEP]) interface. These interfaces need an IP address in the infrastructure tenant that uses DHCP. Therefore, you must configure a DHCP relay policy so that the APIC can act as the DHCP server and provide these IP addresses.

When an ACI fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the ACI fabric must support Option 82.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the Basic GUI



Note Configuring a DHCP Server Policy is only available in the **Advanced Mode**. Toggle to the **Advanced Mode** in the APIC GUI to perform this task.

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

- Step 1** On the menu bar, choose **TENANTS > infra**. In the **Navigation** pane, under **Tenant infra**, expand **Networking > Protocol Policies > DHCP > Relay Policies**.
- Step 2** Right-click **Relay Policies** and click **Create DHCP Relay Policy**.
- Step 3** In the **Create DHCP Relay Policy** dialog box, perform the following actions:
 - a) In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).
 - b) Expand **Providers**. In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.
 - c) In the **Application EPG** area, in the **Tenant** field, from the drop-down list, choose the tenant. (infra)
 - d) In the **Application Profile** field, from the drop-down list, choose the application. (access)
 - e) In the **EPG** field, from the drop-down list, choose the EPG. (default)
 - f) In the **DHCP Server Address** field, enter the IP address for the infra DHCP server. Click **Update**.

Note The infra DHCP IP address is the infra IP address of APIC1. You must enter the default IP address of 10.0.0.1 if deploying for vShield controller configuration.
 - g) Click **Submit**.

The DHCP relay policy is created.

Step 4 In the **Navigation** pane, expand **Networking > Bridge Domains > default > DHCP Relay Labels**.

Step 5 Right-click **DHCP Relay Labels**, and click **Create DHCP Relay Label**.

Step 6 In the **Create DHCP Relay Label** dialog box, perform the following actions:

- a) In the **Scope** field, click the tenant radio button.
This action displays, in the **Name** field drop-down list, the DHCP relay policy created earlier.
- b) In the **Name** field, from the drop-down list, choose the name of the DHCP policy created (DhcpRelayP).
- c) Click **Submit**.

The DHCP server is associated with the bridge domain.

Step 7 In the **Navigation** pane, expand **Networking > Bridge Domains > default > DHCP Relay Labels** to view the DHCP server created.

Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.
- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.
- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking > VRF** policy configuration in the tenants configuration.

Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere

Source	In-Band Management	Out-of-Band Management	External Server Location
Leaf switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be attached to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/fabric/dnsp-default.xml -->
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The `gai.conf` settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a `gai.conf` to prioritize IPv4 over IPv6:

```
# Generated by APIC
label ::1/128      0
label ::/0        1
label 2002::/16   2
label ::/96       3
label ::ffff:0:0/96 4
precedence ::1/128      50
precedence ::/0        40
precedence 2002::/16   30
precedence ::/96       20
# For APICs preferring IPv4 connections, change the value to 100.
precedence ::ffff:0:0/96 10
```

Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPv4) or AAAA records (IPv6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to `/etc/resolv.conf`.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in `/etc/resolv.conf`. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in `/etc/resolv.conf`). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add “options single-request-reopen” to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of `resolv.conf` in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (`glibc`) will use the IPv6 network by default because `getaddrinfo()` will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (`::ffff/96`). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the `glibc` IPv6 selection preference for `getaddrinfo()` in `/etc/gai.conf`.

In order to allow `glibc` to return multiple addresses when using `/etc/hosts`, “multi on” should be added to the `/etc/hosts` file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

Configuring a DNS Service Policy to Connect with DNS Providers Using the Basic GUI

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

- Step 1** On the menu bar, choose **System > System Settings**. In the **Navigation** pane, expand **System Settings > DNS**, and click the default DNS profile.
- Step 2** In the **Work** pane, in the **Management EPG** field, from the drop-down list, choose the appropriate management EPG (default (Out-of-Band)).
- Step 3** Expand **DNS Providers**, and perform the following actions:
- In the **Address** field, enter the provider address.
 - In the **Preferred** column, check the check box if you want to have this address as the preferred provider. You can have only one preferred provider.
 - Click **Update**.
 - (Optional) To add a secondary DNS provider, expand **DNS Providers**, and in the **Address** field, type the provider address. Click **Update**.
- Step 4** Expand **DNS Domains**, and perform the following actions:
- In the **Name** field, enter the domain name (cisco.com).
 - In the **Default** column, check the check box to make this domain the default domain. You can have only one domain name as the default.
 - Click **Update**.
 - (Optional) To add a secondary DNS domain, expand **DNS Domains**. In the **Address** field, enter the secondary domain name. Click **Update**.
- Step 5** Click **Submit**.
The DNS server is configured.
- Step 6** On the menu bar, click **Tenants > mgmt**.
- Step 7** In the **Navigation** pane, expand **Networking > VRFs > oob**.
- Step 8** In the **Work** pane, under **Properties**, in the **DNS labels** field, enter the appropriate DNS label (default). Click **Submit**.
The DNS profile label is now configured on the tenant and VRF.
-

Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI

Procedure

Step 1 Verify the configuration for the default DNS profile.

Example:

```
admin@apic1:~> cd /aci/fabric/fabric-policies/global-policies/dns-profiles/default
admin@apic1:default> cat summary
# dns-profile
name : default
description : added via CLI by tdeleon@cisco.com
ownerkey :
ownertag :

dns-providers:
address preferred
-----
10.44.124.122 no
10.70.168.183 no
10.37.87.157 no
10.102.6.247 yes
dns-domains:
name default description
-----
cisco.com yes
management-epg : tenants/mgmt/node-management-epgs/default/out-of-band/default
```

Step 2 Verify the configurations for the DNS labels.

Example:

```
admin@apic1:default> cd
/aci/tenants/mgmt/networking/private-networks/oob/dns-profile-labels/default
admin@apic1:default> cat summary
# dns-lbl
name : default
description :
ownerkey :
ownertag :
tag : yellow-green
```

Step 3 Verify that the applied configuration is operating on the fabric controllers.

Example:

```
admin@apic1:~> cat /etc/resolv.conf
# Generated by IFC
search cisco.com
nameserver 10.102.6.247
nameserver 10.44.124.122
nameserver 10.37.87.157
nameserver 10.70.168.183
admin@apic1:~> ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=238 time=35.4 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=238 time=29.0 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=238 time=29.2 ms
```

Step 4 Verify that the applied configuration is operating on the leaf and spine switches.

Example:

```
leaf1# cat /etc/resolv.conf
search cisco.com
nameserver 10.102.6.247
nameserver 10.70.168.183
nameserver 10.44.124.122
nameserver 10.37.87.157
leaf1# cat /etc/dcos_resolv.conf
# DNS enabled
leaf1# ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161): 56 data bytes
64 bytes from 72.163.4.161: icmp_seq=0 ttl=238 time=29.255 ms
64 bytes from 72.163.4.161: icmp_seq=1 ttl=238 time=29.212 ms
64 bytes from 72.163.4.161: icmp_seq=2 ttl=238 time=29.343 ms
```

Configuring External Connectivity for Tenants

Before you can distribute the static route to the other leaf switches on the Application Centric Infrastructure (ACI) fabric, a multiprotocol BGP (MP-BGP) process must first be operating, and the spine switches must be configured as BGP route reflectors.

To integrate the ACI fabric into an external routed network, you can configure Open Shortest Path First (OSPF) for management tenant Layer 3 connectivity.

Configuring an MP-BGP Route Reflector Using the Basic GUI

Procedure

- Step 1** On the menu bar, choose **System > System Settings**.
 - Step 2** In the **Navigation** pane, expand **System Settings > BGP Route Reflector**, right-click **BGP Route Reflector**, and click **Create Route Reflector Node Policy EP**.
 - Step 3** In the **Create Route Reflector Node Policy EP** dialog box, from the **Spine Node** drop-down list, choose the appropriate spine node. Click **Submit**.
Note Repeat the above steps to add additional spine nodes as required.
The spine switch is marked as the route reflector node.
 - Step 4** In the **Autonomous System Number** field, choose the appropriate number. Click **Submit**.
Note The autonomous system number must match the leaf connected router configuration if Border Gateway Protocol (BGP) is configured on the router. If you are using routes learned using static or Open Shortest Path First (OSPF), the autonomous system number value can be any valid value.
-

Verifying the MP-BGP Route Reflector Configuration

Procedure

- Step 1** Verify the configuration by performing the following actions:
- Use secure shell (SSH) to log in as an administrator to each leaf switch as required.
 - Enter the **show processes | grep bgp** command to verify the state is S.
If the state is NR (not running), the configuration was not successful.
- Step 2** Verify that the autonomous system number is configured in the spine switches by performing the following actions:
- Use the SSH to log in as an administrator to each spine switch as required.
 - Execute the following commands from the shell window

Example:
`cd /mit/sys/bgp/inst`

Example:
`grep asn summary`

The configured autonomous system number must be displayed. If the autonomous system number value displays as 0, the configuration was not successful.

Creating OSPF External Routed Network for Management Tenant Using Basic GUI

Before You Begin

- You must verify that the router ID and the logical interface profile IP address are different and do not overlap.
- The following steps are for creating an OSPF external routed network for a management tenant. To create an OSPF external routed network for a tenant, you must choose a tenant and create a VRF for the tenant.
- For more details, see *Cisco APIC and Transit Routing*.

Procedure

- Step 1** On the menu bar, click **Fabric > Inventory**. In the **Navigation** pane, click the leaf switch where you want to deploy the VRF.
- Step 2** Right-click and click **Configure VRF**.
- Step 3** In the **Configure VRF** dialog box, perform the following actions:

- a) From the **Tenant** field drop-down list, choose a tenant.
In this case it is the mgmt tenant.
- b) From the **VRF** field drop-down list, choose the VRF.
- c) In the **Router ID** field, enter the router ID.
- d) In the **Protocols** area, check the check box for OSPF.
- e) In the **Route Maps** area that gets displayed, click the + sign.
- f) In the **Create Route Map** dialog box, in the **Name** field, enter a name for the route map.

Step 4 Expand the **Prefix List** area.

- a) In the **Create Prefix List** dialog box, in the **Prefix Name** field, enter an area ID.
- b) In the **IP Prefixes** field, enter at least one prefix.
- c) Enter additional details in the dialog box as required.
- d) In the **Prefix List** dialog box, click **OK**.

Step 5 In the **Configure VRF with Leaf** dialog box, expand the **OSPF Configuration** area.

- a) In the **Configure OSPF** dialog box, in the **Area ID** field, enter an area ID.
- b) In the **Area Type** field, choose the desired type.
- c) From the **Route Map** field drop-down list, choose the appropriate route map. Click **OK**.
- d) In the **Configure VRF** dialog box, click **Submit**.

In the **Navigation** pane under VRFs, the VRF is deployed.

Step 6 In the **Navigation** pane, expand **Interfaces > Physical Interfaces**.

Step 7 Choose the desired interface where you want to configure Layer 3 and perform the following actions:

- a) Click the Convert to L3 button. Click **Yes** in the **Warning** dialog box for "**L2 configuration will be deleted. Do you want to continue?**"
- b) In the **Subinterface** field click the **Off** button.
- c) From the **Tenant** field drop-down list, choose the appropriate tenant (mgmt).
- d) From the **VRF** field drop-down list, choose the appropriate VRF.
- e) In the **IPv4 Address** field, enter an IP address for the interface.
- f) In the **Protocols** field, check the check box for OSPF.
In the top right corner of the dialog box, the OSPF tab is displayed.
- g) Click the **OSPF** tab.
- h) In the **Work** pane, in the **IP V4 OSPF Area ID** field, from the drop-down list, choose the desired area ID.
- i) In the **Policy Name** field, choose the default policy.
Alternatively, choose the authentication type, authentication key, policy name in this area. Click **Submit**.
Note This creates the management interface. The OSPF is deployed in the VRF.
- j) In the **Navigation** pane, when you click on the OSPF deployed in the VRF, the **Work** pane displays the statistics.

The OSPF external routed network for the management tenant is created.

Deploying an Application Policy

Three-Tier Application Deployment

A filter specifies the data protocols to be allowed or denied by a contract that contains the filter. A contract can contain multiple subjects. A subject can be used to realize uni- or bidirectional filters. A unidirectional filter is a filter that is used in one direction, either from consumer-to-provider (IN) or from provider-to-consumer (OUT) filter. A bidirectional filter is the same filter that is used in both directions. It is not reflexive.

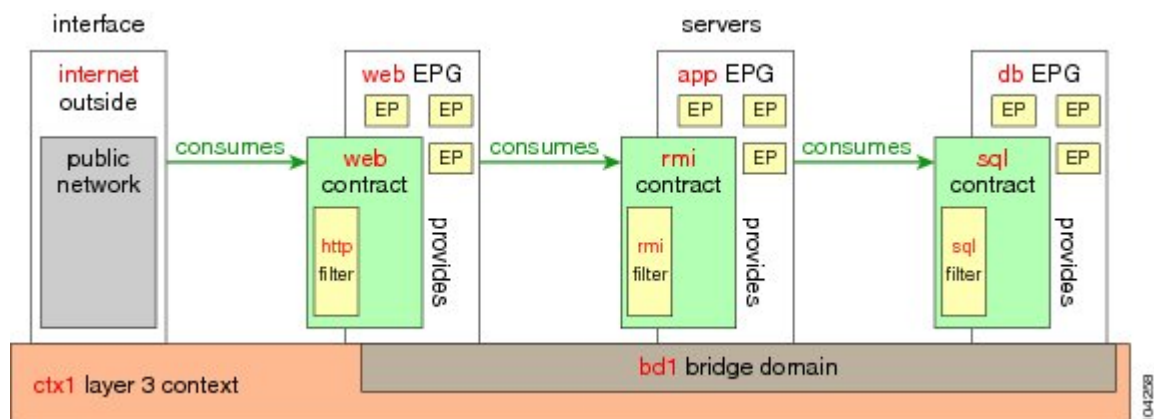
Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

Application profiles enable you to model application requirements that the APIC then automatically renders in the network and data center infrastructure. The application profiles enable administrators to approach the resource pool in terms of applications rather than infrastructure building blocks. The application profile is a container that holds EPGs that are logically related to one another. EPGs can communicate with other EPGs in the same application profile and with EPGs in other application profiles.

To deploy an application policy, you must create the required application profiles, filters, and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

The web server has the HTTP filter, the application server has the Remote Method Invocation (RMI) filter, and the database server has the Structured Query Language (SQL) filter. The application server consumes the SQL contract to communicate with the database server. The web server consumes the RMI contract to communicate with the application server. The traffic enters from the web server and communicates with the application server. The application server then communicates with the database server, and the traffic can also communicate externally.

Figure 7: Three-Tier Application Diagram



Parameters to Create a Filter for http

The parameters to create a filter for http in this example is as follows:

Parameter Name	Filter for http
Name	http
Number of Entries	2
Entry Name	Dport-80 Dport-443
Ethertype	IP
Protocol	tcp tcp
Destination Port	http https

Parameters to Create Filters for rmi and sql

The parameters to create filters for rmi and sql in this example are as follows:

Parameter Name	Filter for rmi	Filter for sql
Name	rmi	sql
Number of Entries	1	1
Entry Name	Dport-1099	Dport-1521
Ethertype	IP	IP
Protocol	tcp	tcp
Destination Port	1099	1521

Example Application Profile Database

The application profile database in this example is as follows:

EPG	Provided Contracts	Consumed Contracts
web	web	rmi
app	rmi	sql
db	sql	--

Deploying an Application Policy Using the Basic GUI

Before You Begin

Verify that the tenant, network, and bridge domain have been created.

Procedure

-
- Step 1 Note** Log in to the **Basic Mode** of the APIC GUI.
- On the menu bar, click **Tenants > Tenant-name**.
- Step 2** In the **Navigation** pane, right-click **Application Profiles** and click **Create Application Profile**.
- Step 3** In the **Create Application Profile** dialog box, enter a name for the profile. Click **Submit**.
- Step 4** In the **Navigation** pane, click and choose the new application profile.
- Step 5** In the **Work** pane, from the **Drag and drop to configure toolbar**, drag and drop the first **EPG** to the blank screen below.
- Step 6** In the **Create Application EPG** dialog box that is displayed, perform the following actions:
- Enter the name for the application EPG.
 - In the **Bridge Domain** field, from the drop-down list, choose the desired bridge domain. Click **OK**. Repeat this step to create additional EPGs as desired in different bridge domains.
- Step 7** From the **Drag and drop to configure toolbar**, drag and drop **Contract**, and it auto connects as the provider EPG the consumer EPG as the user desires and drags. The relationship is displayed with arrows. The **Config Contract With L4-L7 Service Graph** dialog box is displayed with the selected details auto populated. and the provider and consumer contracts associated.
- In the **Contract Name** field, enter a contract name. Click **OK**.
 - In the **No Filter** field, uncheck the check box to create a customized filter.

Note A default filter will be auto created if you do not uncheck the check box.
 - (Optional) To create a customized filter, enter the appropriate information in the **Filter Entries** fields as desired. Click **OK**.
- Step 8** In the **Application Profile** Work pane, click **Submit**. This completes the steps for deploying an application profile.
-



Using the NX-OS Style CLI

This chapter contains the following sections:

- [Accessing the NX-OS Style CLI, page 65](#)
- [Using the NX-OS Style CLI for APIC, page 66](#)
- [About Getting Started with APIC Examples, page 69](#)
- [About Switch Discovery with the APIC, page 69](#)
- [Configuring Network Time Protocol, page 71](#)
- [Creating User Accounts, page 75](#)
- [Adding Management Access, page 79](#)
- [Configuring a VLAN Domain, page 85](#)
- [Configuring a VMM Domain, page 86](#)
- [Creating Tenants, VRFs, and Bridge Domains, page 91](#)
- [Deploying an Application Policy, page 94](#)
- [Configuring External L3 Connectivity for Tenants, page 99](#)
- [Configuring Server or Service Policies, page 101](#)

Accessing the NX-OS Style CLI



Note

From Cisco APIC Release 1.0 until 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.

Procedure

-
- 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. From this level, you can reach these configuration modes:

- To continue in the NX-OS style CLI, you can stay in EXEC mode or you can type **configure** to enter global configuration mode.
For information about NX-OS style CLI commands, see the *Cisco APIC NX-OS Style CLI Command Reference*.
- To reach the object model CLI, type **bash**.
For information about object mode CLI commands, see the *Cisco APIC Command-Line Interface User Guide, APIC Releases 1.0 and 1.1*.

Using the NX-OS Style CLI for APIC

Using CLI Command Modes

The NX-OS style CLI is organized in a hierarchy of command modes with EXEC mode as the root, containing a tree of configuration submodes beginning with global configuration mode. The commands available to you depend on the mode you are in. To obtain a list of available commands in any mode, type a question mark (?) at the system prompt.

This table lists and describes the two most commonly used modes (EXEC and global configuration) along with an example submode (DNS). The table shows how to enter and exit the modes, and the resulting system prompts. The system prompt helps to identify which mode you are in and the commands that are available to you in that mode.

Mode	Access Method	Prompt	Exit Method
EXEC	From the APIC prompt, enter <code>execsh</code> .	<code>apic#</code>	To exit to the login prompt, use the exit command.
Global configuration	From EXEC mode, enter the <code>configure</code> command.	<code>apic(config)#</code>	To exit from a configuration submode to its parent mode, use the exit command.
DNS configuration	From global configuration mode, enter the <code>dns</code> command.	<code>apic(config-dns)#</code>	To exit from any configuration mode or submode to EXEC mode, use the end command.

CLI Command Hierarchy

Configuration mode has several submodes, with commands that perform similar functions grouped under the same level. For example, all commands that display information about the system, configuration, or hardware are grouped under the **show** command, and all commands that allow you to configure the switch are grouped under the **configure** command.

To execute a command that is not available in EXEC mode, you navigate to its submode starting at the top level of the hierarchy. For example, to configure DNS settings, use the **configure** command to enter the global configuration mode, then enter the **dns** command. When you are in the DNS configuration submode, you can query the available commands, as in this example:

```
apic1# configure
apic1(config)# dns
apic1(config-dns)# ?
  address  Configure the ip address for dns servers
  domain   Configure the domains for dns servers
  exit     Exit from current mode
  fabric   Show fabric related information
  no       Negate a command or set its defaults
  show     Show running system information
  use-vrf  Configure the management vrf for dns servers
  where    Show the current mode

apic1(config-dns)# end
apic1#
```

Each submode places you further down in the prompt hierarchy. To view the hierarchy for the current mode, use the **configure** command, as shown in this example:

```
apic1# configure
apic1(config)# bgp-fabric
apic1(config-bgp-fabric)# where
configure t; bgp-fabric
apic1(config-bgp-fabric)#
```

To leave the current level and return to the previous level, type **exit**. To return directly to the EXEC level, type **end**.

EXEC Mode Commands

When you start a CLI session, you begin in EXEC mode. From EXEC mode, you can enter configuration mode. Most EXEC commands are one-time commands, such as show commands, which display the current configuration status.

Configuration Mode Commands

Configuration mode allows you to make changes to the existing configuration. When you save the configuration, these commands are saved across switch reboots. Once you are in configuration mode, you can enter a variety of protocol-specific modes. Configuration mode is the starting point for all configuration commands.

Listing Commands and Syntax

In any command mode, you can obtain a list of available commands by entering a question mark (?).

```
apicl(config-dns)# ?
address    Configure the ip address for dns servers
domain     Configure the domains for dns servers
exit       Exit from current mode
fabric     Show fabric related information
no         Negate a command or set its defaults
show       Show running system information
use-vrf    Configure the management vrf for dns servers
where      Show the current mode

apicl(config-dns)# end
apicl#
```

To see a list of commands that begin with a particular character sequence, type those characters followed by a question mark (?). Do not include a space before the question mark.

```
apicl(config)# sh ?
aaa        Show AAA information
access-list Show Access-list Information
accounting Show accounting information
acllog     Show acllog information
. . .
```

To complete a command after you begin typing, type a tab.

```
apicl# qu<TAB>
apicl# quota
```

To list keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help because it reminds you which keywords or arguments are applicable based on the commands, keywords, and arguments you have already entered.

```
apicl(config-dns)# use-vrf ?
inband-mgmt Configure dns on inband
oob-mgmt    Configure dns on out-of-band

apicl(config-dns)#
```

You can also abbreviate a command if the abbreviation is unambiguous. In this example, the **configure** command is abbreviated.

```
apicl# conf
apicl(config)#
```

Undoing or Reverting to Default Values or Conditions Using the 'no' Prefix

For many configuration commands, you can precede the command with the **no** keyword to remove a setting or to restore a setting to the default value. This example shows how to remove a previously-configured DNS address from the configuration.

```
apic1(config-dns)# address 192.0.20.123 preferred
apic1(config-dns)# show dns-address
Address                Preferred
-----
192.0.20.123          yes

apic1(config-dns)# no address 192.0.20.123
apic1(config-dns)# show dns-address
Address                Preferred
-----
```

Executing BASH Commands From the NX-OS Style CLI

To execute a single command in the bash shell, type **bash -c 'path/command'** as shown in this example.

```
apic1# bash -c '/controller/sbin/acidiag avread'
```

You can execute a bash command from any mode or submode in the NX-OS style CLI.

Entering Configuration Text with Spaces or Special Characters

When a configuration field consists of user-defined text, special characters such as '\$' should be escaped ('\\$\$') or the entire word or string should be wrapped in single quotes to avoid misinterpretation by Bash.

About Getting Started with APIC Examples

The steps in several examples in this guide include a parameter name. These parameter names are provided as examples for convenience and ease of your understanding, and it is not required for you to use them.

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



Note 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>.

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



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.

Registering Unregistered Switches Using the NX-OS Style CLI



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.

Procedure

Step 1 Start in Configuration mode, shown as follows:

Example:

```
apicl# configure
apicl(config)#
```

Step 2 Register the switches, as shown in the following example:

Note To obtain the serial number, find the serial number that is physically printed on the node itself, or use the command **acidiag fmvread** for a list of discovered node serial numbers.

Example:

```
apicl(config)# system switch-id FGE173900ZD 101 leaf1
```

Step 3 Repeat the previous step for the remaining switches.

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.

Configuring Network Time Protocol

Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

In-Band and Out-of-Band Management NTP

**Note**

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
 - See the Adding Management Access section in this guide for information about in-band management access and out-of-band management access.
-
- Out-of-band management NTP—When an ACI fabric is deployed with out-of-band management, each node of the fabric, inclusive of spines, leaves, and all members of the APIC cluster, is managed from outside the ACI fabric. This IP reachability will be leveraged so that each node can individually query the same NTP server as a consistent clock source. To configure NTP, a Date and Time policy must be created that references an out-of-band management endpoint group. Date and Time policies are confined

to a single pod and must be deployed across all pods provisioned in the ACI fabric. Currently only one pod per ACI fabric is allowed.

- **In-Band Management NTP**—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication. The steps used to configure in-band management policies are identical to those used to establish an out-of-band management policy. The distinction is around how to allow the fabric to connect to the NTP server.

NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The `gai.conf` can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

Configuring NTP Using the NX-OS Style CLI

When an ACI fabric is deployed with out-of-band management, each node of the fabric is managed from outside the ACI fabric. You can configure an out-of-band management NTP server so that each node can individually query the same NTP server as a consistent clock source.

Procedure

Step 1 **configure**

Enters configuration mode.

Example:

```
apicl# configure
```

Step 2 **template ntp-fabric** *ntp-fabric-template-name*

Specifies the NTP template (policy) for the fabric.

Example:

```
apicl(config)# template ntp-fabric poll
```

Step 3 **[no] server** *dns-name-or-ipaddress* **[prefer]** **[use-vrf {inband-mgmt | oob-default}]** **[key** *key-value* **]**

Configures an NTP server for the active NTP policy. To make this server the preferred server for the active NTP policy, include the **prefer** keyword. If NTP authentication is enabled, specify a reference key ID. To specify the in-band or out-of-band management access VRF, include the **use-vrf** keyword with the **inb-default** or **oob-default** keyword.

Example:

```
apicl(config-template-ntp-fabric)# server 192.0.20.123 prefer use-vrf oob-mgmt
```

Step 4 **[no] authenticate**

Enables (or disables) NTP authentication.

Example:

```
apic1(config-template-ntp-fabric)# no authenticate
```

Step 5 [no] **authentication-key** *key-value*

Configures an authentication NTP authentication. The range is 1 to 65535.

Example:

```
apic1(config-template-ntp-fabric)# authentication-key 12345
```

Step 6 [no] **trusted-key** *key-value*

Configures a trusted NTP authentication. The range is 1 to 65535.

Example:

```
apic1(config-template-ntp-fabric)# trusted-key 54321
```

Step 7 **exit**

Returns to global configuration mode

Example:

```
apic1(config-template-ntp-fabric)# exit
```

Step 8 **template pod-group** *pod-group-template-name*

Configures a pod-group template (policy).

Example:

```
apic1(config)# template pod-group allPods
```

Step 9 **inherit ntp-fabric** *ntp-fabric-template-name*

Configures the NTP fabric pod-group to use the previously configured NTP fabric template (policy).

Example:

```
apic1(config-pod-group)# inherit ntp-fabric poll
```

Step 10 **exit**

Returns to global configuration mode

Example:

```
apic1(config-template-pod-group)# exit
```

Step 11 **pod-profile** *pod-profile-name*

Configures a pod profile.

Example:

```
apic1(config)# pod-profile all
```

Step 12 **pods** {*pod-range-1-255* | **all**}

Configures a set of pods.

Example:

```
apic1(config-pod-profile)# pods all
```

Step 13 **inherit pod-group** *pod-group-name*

Associates the pod-profile with the previously configured pod group.

Example:

```
apicl(config-pod-profile-pods)# inherit pod-group allPods
```

Step 14 end

Returns to EXEC mode.

Example:

```
apicl(config-pod-profile-pods)# end
```

Examples

This example shows how to configure a preferred out-of-band NTP server and how to verify the configuration and deployment.

```
apicl# configure t
apicl(config)# template ntp-fabric poll
apicl(config-template-ntp-fabric)# server 192.0.20.123 use-vrf oob-default
apicl(config-template-ntp-fabric)# no authenticate
apicl(config-template-ntp-fabric)# authentication-key 12345
apicl(config-template-ntp-fabric)# trusted-key 12345
apicl(config-template-ntp-fabric)# exit
apicl(config)# template pod-group allPods
apicl(config-pod-group)# inherit ntp-fabric poll
apicl(config-pod-group)# exit
apicl(config)# pod-profile all
apicl(config-pod-profile)# pods all
apicl(config-pod-profile-pods)# inherit pod-group allPods
apicl(config-pod-profile-pods)# end
apicl#
```

```
apicl# show ntpq
nodeid  remote          refid  st  t  when  poll  reach  delay  offset  jitter
-----  -  -----  ----  --  -----  -----  -----  -----  -----  -----
1         *  192.0.20.123  .GPS.  u  27    64    377    76.427  0.087  0.067
2         *  192.0.20.123  .GPS.  u  3     64    377    75.932  0.001  0.021
3         *  192.0.20.123  .GPS.  u  3     64    377    75.932  0.001  0.021
```

Verifying NTP Operation Using the NX-OS Style CLI

Procedure

Verify that the NTP policy is deployed to APIC using the NX-OS CLI using **show ntp**:

Example:

```
apicl# show ntp
```

Verifying NTP Policy Deployed to Each Node Using the NX-OS Style CLI

Procedure

-
- Step 1** Log onto an APIC controller in the fabric using the SSH protocol.
 - Step 2** Attach to a node and check the NTP peer status, shown as follows:

```
apic1# fabric node_name show ntp peer-status
```
 - Step 3** Repeat step 2 for different nodes in the fabric.
-

Creating User Accounts

Configuring a Local User Using the NX-OS Style CLI

In the initial configuration script, the admin account is configured and the admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

To configure a Cisco AV Pair on an external authentication server, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair format is as follows:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(16003)
```

The first av-pair format has no UNIX user ID, while the second one does. Both are correct if all remote users have the same role and mutual file access is acceptable. If the UNIX user ID is not specified, ID 23999 is applied by the APIC system, and more than one role/read privilege is specified to any AV Pair user. This can cause users to have higher or lower permissions than configured through the group settings.



Note

The APIC Cisco AV-pair format is compatible and can co-exist with other Cisco AV-pair formats. APIC will pick up the first matching AV-pair from all the AV-pairs.

The APIC supports the following regexes:

```
shell:domains\\s* [=:]\\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31}) (\\ (\\d+\\))$
shell:domains\\s* [=:]\\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31})$
```

Examples:

- Example 1: A Cisco AV Pair that contains a single Login domain with only writeRoles:

```
shell:domains=domainA/writeRole1|writeRole2/
```

- Example 2: A Cisco AV Pair that contains a single Login domain with only readRoles:

```
shell:domains=domainA//readRole1|readRole2
```

**Note**

The "/" character is a separator between writeRoles and readRoles per Login domain and is required even if only one type of role is to be used.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

```
aaa-network-admin Cleartext-Password := "<password>"
Cisco-avpair = "shell:domains = all/aaa/read-all(16001)"
```

Changing Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs Using the NX-OS Style CLI

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. To do so, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+. One AV pair format contains a Cisco UNIX user ID and one does not. Both are correct if all remote users have the same role and mutual file access is acceptable. If the UNIX user ID is not specified, ID 23999 is applied by the APIC system, and more than one role/read privilege is specified to any AV Pair user. This can cause users to have higher or lower permissions than configured through the group settings. This topic explains how to change the behavior if that is not acceptable.

To change the default behavior for remote users with missing or bad Cisco AV pairs using the NX-OS CLI:

Procedure

Step 1 In the NX-OS CLI, start in Configuration mode.

Example:

```
apic1#
apic1# configure
```

Step 2 Configure the aaa user default role.

Example:

```
apic1(config)# aaa user default-role
assign-default-role assign-default-role
no-login no-login
```

Step 3 Configure the aaa authentication login methods.

Example:

```

apic1(config)# aaa authentication
login Configure methods for login

apic1(config)# aaa authentication login
console Configure console methods
default Configure default methods
domain Configure domain methods

apic1(config)# aaa authentication login console
<CR>

apic1(config)# aaa authentication login domain
WORD Login domain name
fallback

```

Best Practice for Assigning AV Pairs

As best practice, Cisco recommends that you assign unique UNIX user ids in the range 16000-23999 for the AV Pairs that are assigned to users when in bash shell (using SSH, Telnet or Serial/KVM consoles). If a situation arises when 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 remote users with a UNIX ID of 23999.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is logged in using Secure Shell (SSH) or Telnet.

Procedure

Configure an AV pair on the external authentication server.

The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

```

* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2

* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(8101)

These are the boost regexes supported by APIC:
uid_regex("shell:domains\\s*[:]\\s*(\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31}) (\\d+\\S*)$");
regex("shell:domains\\s*[:]\\s*(\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31})$");

```

The following is an example:

```
shell:domains = coke/tenant-admin/read-all,pepsi//read-all(16001)
```

Configuring a Remote User Using the NX-OS Style CLI

Instead of configuring local users, you can point the APIC at the centralized enterprise credential datacenter. The APIC supports Lightweight Directory Access Protocol (LDAP), active directory, RADIUS, and TACACS+.

To configure a remote user authenticated through an external authentication provider, you must meet the following prerequisites:

- The DNS configuration should have already been resolved with the hostname of the RADIUS server.
- You must configure the management subnet.

Configuring a Remote User With the NX-OS Style CLI

Procedure

Step 1 In the NX-OS CLI, start in Configuration mode, shown as follows:

Example:

```
apicl# configure
apicl(config)#
```

Step 2 Create a RADIUS provider, shown in the following example:

Example:

```
apicl(config)# radius-server
  host      RADIUS server's DNS name or its IP address
  retries   Global RADIUS server retransmit count
  timeout   Global RADIUS server timeout period in seconds

apicl(config)# radius-server host 1.1.1.1
apicl(config-host)#
  descr     RADIUS server descr for authentication
  exit      Exit from current mode
  fabric    show fabric related information
  key       RADIUS server key for authentication
  no        Negate a command or set its defaults
  port      RADIUS server port for authentication
  protocol  RADIUS server protocol for authentication
  retries   RADIUS server retries for authentication
  show      Show running system information
  timeout   RADIUS server timeout for authentication
  where     show the current mode

apicl(config-host)# exit
```

Step 3 Create a TACACS+ provider, shown in the following example:

Example:

```
apicl(config)# tacacs-server
  host      TACACS+ server's DNS name or its IP address
  retries   Global TACACS+ server retries period in seconds
  timeout   Global TACACS+ server timeout period in seconds
```



```
apicl(config)# tacacs-server host 1.1.1.1
apicl(config-host)# exit
```

Step 4 Create an LDAP provider, shown in the following example:

Example:

```
apicl(config)# ldap-server
attribute  An LDAP endpoint attribute to be used as the CiscoAVPair
basedn    The LDAP base DN for user lookup in the LDAP directory tree
filter    LDAP search filter for the LDAP endpoint
host      LDAP server DNS name or IP address
retries   Global LDAP server retransmit count
timeout   Global LDAP server timeout period in seconds

apicl(config)# ldap-server host 1.1.1.1
apicl(config-host)#
enable-ssl  enabling an SSL connection with the LDAP provider
exit       Exit from current mode
fabric     show fabric related information
filter     Set the LDAP filter to be used in a user search
key        LDAP server key for authentication
no         Negate a command or set its defaults
port      LDAP server port for authentication
retries    LDAP server retries for authentication
show      Show running system information
ssl-validation-level  Set the LDAP Server SSL Certificate validation level
timeout   LDAP server timeout for authentication
where     show the current mode

apicl(config-host)# exit
apicl(config)#
```

Adding Management Access

IPv4/IPv6 Addresses and In-Band Policies

In-band management addresses can be provisioned on the APIC controller only through a policy (Postman REST API, NX-OS Style CLI, or GUI). Additionally, the in-band management addresses must be configured statically on each node.

IPv4/IPv6 Addresses in Out-of-Band Policies

Out-of-band management addresses can be provisioned on the APIC controller either at the time of bootstrap or by using a policy (Postman REST API, NX-OS Style CLI, GUI). Additionally, the out-of-band management addresses must be configured statically on each node or by specifying a range of addresses (IPv4/IPv6) to the entire cluster. IP addresses are randomly assigned from a range to the nodes in the cluster.

Adding Management Access Using the NX-OS Style CLI

An APIC controller has two routes to reach the management network, one is by using the in-band management interface and the other is by using the out-of-band management interface.

- In-band management access—You can configure in-band management connectivity to the APIC and the ACI fabric. You first configure the VLANs that will be used by APIC when the APIC is communicating with the leaf switches, and then you configure the VLANs that the VMM servers will use to communicate with the leaf switches.
- Out-of-band management access—You can configure out-of-band management connectivity to the APIC and the ACI fabric. You configure an out-of-band contract that is associated with an out-of-band endpoint group (EPG), and attach the contract to the external network profile.



Note The APIC out-of-band management connection link must be 1 Gbps.

The APIC controller always selects the in-band management interface over the out-of-band management interface, if the in-band management interface is configured. The out-of-band management interface is used only when the in-band management interface is not configured, or if the destination address is on the same subnet as the out-of-band management subnet of the APIC. This behavior cannot be changed or reconfigured.

The APIC management interface does not support an IPv6 address and cannot connect to an external IPv6 server through this interface.

Configuring the external management instance profile under the management tenant for in-band or out-of-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet.

Configuring In-Band Management Access for APIC Controller, Spine, Leaf Switches Using the NX-OS CLI



Note IPv4 and IPv6 addresses are supported for in-band management access. IPv6 configurations are supported using static configurations (for both in-band and out-of-band). IPv4 and IPv6 dual in-band and out-of-band configurations are supported only through static configuration. For more information, see the KB article, *Configuring Static Management Access in Cisco APIC*.

Procedure

-
- Step 1** Start in the configuration mode to modify one or more of the IP addresses of the in-band management interface of the APIC controllers, as shown in the following example:

Example:

```
apic1# configure
apic1(config)# controller 1
apic1(config-controller)# interface inband-mgmt0
apic1(config-controller-if)# ip address 10.13.1.1/24 gateway 10.13.1.254
apic1(config-controller-if)# exit
```

Note You wouldn't have this for the inband management interface for the switches.

- Step 2** You can configure the in-band management interface of spine and leaf switches by entering into the switch configuration mode. Enter the switch followed by the ID of the switch, shown as follows:

Example:

```
apic1(config)# switch 101
apic1(config-switch)# interface inband-mgmt0
apic1(config-switch-if)# ip address 10.13.1.101/24 gateway 10.13.1.254
```

Note Switch 101 in the above example can be a leaf or spine switch. There are two types of switches (spines and leaves), but for the purposes of management configuration, it does not matter if the switches are spines or leaves, you can use the same configuration for both.

Example:

To configure a range of switches using successive addresses from a IP address pool, the **ip address-range** command can be used. The following example shows how you can configure IP addresses of the inband management ports for a range of switches at the same time.

```
apic1(config)# switch 101-104
apic1(config-switch)# interface mgmt0
apic1(config-switch-if)# ip address-range 172.23.48.21/21 gateway 172.23.48.1
apic1(config-switch-if)# exit
apic1(config-switch)# exit
```

Step 3 To establish connectivity to in-band management ports from the outside network, perform the following configuration steps:

- a) Create a VLAN domain for the VLAN used for external in-band connectivity.

Example:

Note In the following example, the Management station that is used to connect to the in-band network is connected to Leaf 102, port 2, on VLAN 11, and is in the subnet 179.10.1.0/24.

```
apic1(config)# vlan-domain external-inband
apic1(config-vlan)# vlan 11
apic1(config-vlan)# exit
```

- b) Add the port connected to the external Management station to the VLAN domain and open up the VLAN on the port for in-band connectivity, as shown in the following example:

Example:

Note The address 179.10.1.254/24 is the gateway address used by the external management station and the gateway functionality is provided by the ACI fabric.

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/2
apic1(config-leaf-if)# switchport trunk allowed vlan 11 inband-mgmt 179.10.1.254/24
```

With the previous configuration, the external management station can connect to the in-band port on the spine and leaf switches.

For connectivity to the APIC controller inband port, additional configuration is required in order to open up the VLAN on the port connected to the controller, as explained in the following example. Controller 1 is connected to port on Ethernet 1/1 on leaf 110 and VLAN 10 is used for the APIC controller in-band connectivity.

To configure the in-band VLAN on the controller:

```
apic1(config)# controller 1
apic1(config-controller)# interface inband-mgmt0
apic1(config-controller-if)# ip address x.x.x.x gateway x.x.x.y
apic1(config-controller-if)# vlan 10
apic1(config-controller-if)# inband-mgmt epg inb-default
```

To create a VLAN domain for the APIC in-band VLAN:

```
apic1(config)# vlan-domain apic-inband
apic1(config-vlan)# vlan 10
apic1(config-vlan)# exit
```

To allow the VLAN on the port connected to the controller:

```
apicl(config)# leaf 101
apicl(config-leaf)# interface ethernet 1/1
apicl(config-leaf-if)# vlan-domain member apic-inband
```

Note With the previous configuration, the external management station can connect to the in-band port on the controller. Note that the inband VLAN (VLAN 10 in this example) must be the same on all the controllers.

Step 4 To provide access control for specific protocols on APIC in-band ports from the external network:

Example:

```
apicl(config)# tenant mgmt
apicl(config-tenant)# access-list inband-default
apicl(config-tenant-acl)# no match raw inband-default
apicl(config-tenant-acl)# match tcp dest 443
apicl(config-tenant-acl)# match tcp dest 22
```

In the previous example, "no match raw inband-default" deletes the allow all entry in the default access-list filter. The following match tcp dest 443 and 22 allow access to only these tcp ports on the in-band ports.

What to Do Next

- You must use the new IP address to reconnect to the APIC controller.
- You must delete the old IP address of the controller once a new IP address is assigned to it.

Configuring Out-of-Band Management Access for APIC Controller, Spine, Leaf Switches Using the NX-OS CLI



Note IPv4 and IPv6 addresses are supported for out-of-band management access.

Procedure

Step 1 Start in configuration mode to modify one or more of the IP addresses of the out-of-band management interface of the APIC controllers, as shown in the following example:

Example:

```
apicl# configure
apicl(config)# controller 1
apicl(config-controller)# interface mgmt0
apicl(config-controller-if)# ip address 172.23.48.16/21 gateway 172.23.48.1
apicl(config-controller-if)# exit
apicl(config-controller)# exit
```

Example:

The following example shows how to enter a range of IP addresses when configuring multiple controllers:

Note In this example, controller 1 is assigned 172.23.48.16/21, controller 2 is assigned 172.23.48.17/21, and controller 3 is assigned 172.23.48.18/21 address.

```
apic1(config)# controller 1-3
apic1(config-controller)# interface mgmt0
apic1(config-controller-if)# ip address-range 172.23.48.16/21 gateway 172.23.48.1
```

Step 2 You can configure the out-of-band management interface of spine and leaf switches by entering into the switch configuration mode. Enter the switch followed by the ID of the switch, shown as follows:

Example:

```
apic1(config)# switch 101
apic1(config-switch)# interface mgmt0
apic1(config-switch-if)# ip address 172.23.48.101/21 gateway 172.23.48.1
```

Example:

Note Switch 101 in the above example can be a leaf or spine switch. There are two types of switches (spines and leafs), but for the purposes of management configuration, it does not matter if the switches are spines or leafs, you can use the same configuration for both.

Example:

To configure a range of switches using successive addresses from a IP address pool, the **ip address-range** command can be used. The following example shows how you can configure IP addresses of four switches at the same time.

```
apic1(config)# switch 101-104
apic1(config-switch)# interface mgmt0
apic1(config-switch-if)# ip address-range 172.23.48.21/21 gateway 172.23.48.1
apic1(config-switch-if)# exit
apic1(config-switch)# exit
```

Step 3 To establish connectivity to out-of-band management ports from the outside network, perform the following configuration steps:

a) Provide access control for the out-of-band management interface to specific external subnets.

Example:

Note In this example, except for 179.10.1.0/24 network, none of the other external networks have connectivity to the out-of-band management interfaces in the APIC controller or leaf / spine switches. System Management policies are configured under a special tenant called mgmt.

```
apic1(config)# tenant mgmt
apic1(config-tenant)# external-l3 epg default oob-mgmt
apic1(config-tenant-l3ext-epg)# match ip 179.10.1.0/24
apic1(config-tenant-l3ext-epg)# exit
apic1(config-tenant)# exit
apic1(config)#
```

b) To provide access-control to specific protocols on the out-of-band management ports from the external network:

Example:

Note In this example, "no match raw oob-default" deletes the allow all entry in the default access-list filter. The following match tcp dest 443 and 22 allow access on the management interface only on these specified ports.

```
apic1(config)# tenant mgmt
apic1(config-tenant)# access-list oob-default
apic1(config-tenant-acl)# no match raw oob-default
apic1(config-tenant-acl)# match tcp dest 443
apic1(config-tenant-acl)# match tcp dest 22
```

What to Do Next

- You must use the new IP address to reconnect to the APIC controller.
- You must delete the old IP address of the controller once a new IP address is assigned to it.

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

All IPv6 tables mirror the existing IP tables functionality, except for Network Address Translation (NAT).

Existing IP Tables

- 1 Earlier, every rule in the IPv6 tables were executed one at a time and a system call was made for every rule addition or deletion.
- 2 Whenever a new policy was added, rules were appended to the existing IP tables file and no extra modifications were done to the file.
- 3 When a new source port was configured in the out-of-band policy, it added source and destination rules with the same port number.

Modifications to IP Tables

- 1 When IP tables are created, they are first written into hash maps that are then written into intermediate file IP tables-new which are restored. When saved, a new IP tables file is created in the /etc/sysconfig/ folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.
- 2 When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (/etc/sysconfig/iptables-new) and saved.
- 3 It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.
- 4 When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.


```
-A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
```
- 5 When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.

**Note**

- If only IPv4 is enabled, do not configure an IPv6 policy.
- If only IPv6 is enabled, do not configure an IPv4 policy.
- If both IPv4 and IPv6 are enabled and a policy is added, it will be configured to both the versions . So when you add an IPv4 subnet, it will be added to IP tables and similarly an IPv6 subnet is added to IPv6 tables.

Configuring a VLAN Domain

Configuring a VLAN Domain Using the NX-OS Style CLI

The ACI fabric can be partitioned into groups of 4K VLANs to allow a large number of Layer 2 (L2) domains across the fabric, which can be used by multiple tenants.

A VLAN domain represents a set of VLANs that can be configured on group of nodes and ports. VLAN domains allow multiple tenants to share the common fabric resources, such as nodes, ports, and VLANs, without conflicting with each other and independently managing them. A tenant can be provided access to one or more VLAN domains.

VLAN domains can be static or dynamic. Static VLAN domains support static VLAN pools, while dynamic VLAN domains can support both static and dynamic VLAN pools. VLANs in the static VLAN pools are managed by the user and are used for applications such as connectivity to bare metal hosts. VLANs in the dynamic VLAN pool are allocated and managed by the APIC without user intervention and are used for applications such as VMM. The default type for VLAN domains and VLAN pools within the domain is static.

You must perform this procedure before tenants can start using the fabric resources for their L2/L3 configurations. For detailed steps with examples on how to use the NX-OS CLI for this procedure, see [Creating a Tenant, VRF, and Bridge Domain Using the NX-OS Style CLI](#), on page 91.

Procedure

Step 1 Create a VLAN domain and assign VLANs in each VLAN domain.

Example:

```
apic1# configure
apic1(config)# vlan-domain dom1
apic1(config-vlan)# vlan 5-100
apic1(config-vlan)# exit
apic1(config)# vlan-domain dom2 dynamic
apic1(config-vlan)# vlan 101-200
apic1(config-vlan)# vlan 301-400 dynamic
apic1(config-vlan)# exit
apic1(config)# vlan-domain dom3
apic1(config-vlan)# vlan 401-500
```

Step 2 Configure the VLAN domain membership on the ports on the leaf switches.

Example:

```
apic1(config)# leaf 101,102
apic1(config-leaf)# interface ethernet 1/10-20
```

```

apicl(config-leaf-if)# vlan-domain member dom1
apicl(config-leaf-if)# vlan-domain member dom2
apicl(config-leaf-if)#exit
apicl(config-leaf)# interface ethernet 1/21
apicl(config-leaf-if)# vlan-domain member dom3
apicl(config-leaf-if)#exit

```

Step 3 Convert some ports to be used as L3 Port for External-L3 connectivity through L3Port or through its sub-interfaces.

Example:

```

apicl(config)# leaf 101
apicl(config-leaf)# interface ethernet 1/21
apicl(config-leaf-if)# no switchport

```

In this example, sub-interface encapsulations on ethernet1/21 come from vlans allowed in dom3.

Step 4 Verify the configuration.

Related Topics

[Creating a Tenant, VRF, and Bridge Domain Using the NX-OS Style CLI, on page 91](#)

Configuring a VMM Domain

Configuring a VMM Domain Using the NX-OS Style CLI

Configuring Virtual Machine Networking Policies

The APIC integrates with third-party VM manager (VMM) (for example, VMware vCenter and SCVMM) to extend the benefits of ACI to the virtualized infrastructure. The APIC enables the ACI policies inside the VMM system to be used by its administrator.

This section provides examples of VMM integration using VMware vCenter and vShield. For details about the different modes of Cisco ACI and VMM integration, see the ACI Virtualization Guide.

About the VM Manager



Note

Information about the necessary configuration of the APIC for integration with the vCenter is described here. For instructions about configuring the VMware components, see the VMware documentation.

The following are details of some VM manager terms:

- A VM controller is an external virtual machine management entity such as VMware vCenter, and the VMware vShield. The APIC communicates with the controller to publish network policies that are applied to virtual workloads. A VM controller administrator provides an APIC administrator with a VM controller authentication credential; multiple controllers of the same type can use the same credential.

- A virtual machine mobility domain (vCenter mobility domain) is a grouping of VM controllers with similar networking policy requirements. This mandatory container holds one or more VM controllers with policies such as for a VLAN pool, server to network MTU policy, or server to network access LACP policy. When an endpoint group gets associated with a vCenter domain, network policies get pushed to all the VM controllers in the vCenter domain.
- See [Configuring a VLAN Domain Using the NX-OS Style CLI](#), on page 85 for information about VLAN domains.
- For a VMware vCenter to be deployed, it must operate in VLAN mode or VXLAN mode. A VMM domain must be associated with a VLAN pool and a vShield must be associated with the vCenter.

Prerequisites for Creating a VMM Domain Profile

To configure a VMM domain profile, you must meet the following prerequisites:

- All fabric nodes are discovered and configured.
- Inband (inb) or out-of-band (oob) management has been configured on the APIC.
- A Virtual Machine Manager (VMM) is installed, configured, and reachable through the inb/oob
- You have the administrator/root credential to the VMM (for example vCenter).



Note

If you prefer not to use the vCenter admin/root credentials, you can create a custom user account with minimum required permissions. See [Custom User Account with Minimum VMware vCenter Privileges](#), on page 87 for a list of the required user privileges.

- A DNS policy for the APIC must be configured if you plan to reference the VMM by hostname rather than an IP address.
- A DHCP server and relay policy must be configured if you are creating a domain profile for VMware vShield.

Custom User Account with Minimum VMware vCenter Privileges

To configure the vCenter from Cisco APIC, your credentials must allow the following minimum set of privileges within the vCenter:

- Alarms
- Datacenter
- Folder
- Distributed Switch
- dvPortgroup
- Network
- VM

- Host

Creating a VMM Domain Profile

This section describes how to create a VMM domain profile using the NX-OS CLI and provides examples for a vCenter domain or vCenter and vShield domains.

Creating a vCenter Domain Profile Using the NX-OS Style CLI

Before You Begin

This section describes how to create a vCenter domain profile using the NX-OS style CLI:

Procedure

Step 1 In the CLI, enter configuration mode:

Example:

```
apicl# configure
apicl(config)#
```

Step 2 Configure a VLAN domain:

Example:

```
apicl(config)# vlan-domain dom1 dynamic
apicl(config-vlan)# vlan 150-200 dynamic
apicl(config-vlan)# exit
apicl(config)#
```

Step 3 Add interfaces to this VLAN domain. These are the interfaces to be connected to VMware hypervisor uplink ports:

Example:

```
apicl(config)# leaf 101-102
apicl(config-leaf)# interface ethernet 1/2-3
apicl(config-leaf-if)# vlan-domain member dom1
apicl(config-leaf-if)# exit
apicl(config-leaf)# exit
```

Step 4 Create a VMware domain and add VLAN domain membership:

Example:

```
apicl(config)# vmware-domain vmmdom1
apicl(config-vmware)# vlan-domain member dom1
apicl(config-vmware)#
```

Create the domain with a specific delimiter:

Example:

```
apic1(config)# vmware-domain vmmdom1 delimiter @
```

Step 5 Configure the domain type to DVS:**Example:**

```
apic1(config-vmware)# configure-dvs
apic1(config-vmware-dvs)# exit
apic1(config-vmware)#
```

Step 6 Configure a controller in the domain:**Example:**

```
apic1(config-vmware)# vcenter 192.168.66.2 datacenter prodDC
apic1(config-vmware-vc)# username administrator
Password:
Retype password:
apic1(config-vmware-vc)# exit
apic1(config-vmware)# exit
apic1(config)# exit
```

Note When configuring the password, you must precede special characters such as '\$' or '!' with a backslash ('\') to avoid misinterpretation by the Bash shell. The escape backslash is necessary only when configuring the password; the backslash does not appear in the actual password.

Step 7 Verify configuration:**Example:**

```
apic1# show running-config vmware-domain vmmdom1
# Command: show running-config vmware-domain vmmdom1
# Time: Wed Sep  2 22:14:33 2015
vmware-domain vmmdom1
  vlan-domain member dom1
  vcenter 192.168.66.2 datacenter prodDC
  username administrator password *****
  configure-dvs
  exit
exit
```

Creating a vCenter and a vShield Domain Profile Using the NX-OS Style CLI

Before You Begin

This section describes how to create a vCenter and vShield domain profile using the NX-OS CLI:

Procedure**Step 1** In the NX-OS CLI, enter configuration mode as follows:

Example:

```
apicl# configure
apicl(config)# exit
```

Step 2 Configure a VLAN domain as follows:

Example:

```
apicl(config)# vlan-domain dom1 dynamic
apicl(config-vlan)# vlan 150-200 dynamic
apicl(config-vlan)# exit
apicl(config)#
```

Step 3 Add interfaces to this VLAN domain. These are the interfaces to be connected to VMware hypervisor uplink ports as follows:

Example:

```
apicl(config)# leaf 101-102
apicl(config-leaf)# interface ethernet 1/2-3
apicl(config-leaf-if)# vlan-domain member dom1
apicl(config-leaf-if)# exit
apicl(config-leaf)# exit
apicl(config)#
```

Step 4 Create a VMware domain and add VLAN domain membership as follows:

Example:

```
apicl(config)# vmware-domain vmmdom1
apicl(config-vmware)# vlan-domain member dom1
apicl(config-vmware)#
```

Step 5 Configure the domain type to DVS as follows:

Example:

```
apicl(config-vmware)# configure-dvs
apicl(config-vmware-dvs)# exit
apicl(config-vmware)#
```

Step 6 Configure a vCenter controller in the domain as follows:

Example:

```
apicl(config-vmware)# vcenter 192.168.66.2 datacenter prodDC
apicl(config-vmware-vc)# username administrator password "password"
apicl(config-vmware-vc)#
```

Step 7 Configure a VShield controller attached to this VCenter, and configure vxlan and multicast address pools for this VShield as follows:

Example:

```
apicl(config-vmware-vc)# vshield 123.4.5.6
apicl(config-vmware-vc-vs)# username administrator password "password"
apicl(config-vmware-vc-vs)# vxlan pool 10000-12000
apicl(config-vmware-vc-vs)# vxlan multicast-pool 224.3.4.5-224.5.6.7
apicl(config-vmware-vc-vs)# exit
apicl(config-vmware-vc)#
```

Step 8 Verify the configuration as follows:

Example:

```
apicl# show running-config vmware-domain vmmdom1
# Command: show running-config vmware-domain vmmdom1
# Time: Wed Sep 2 22:14:33 2015
vmware-domain vmmdom1
```

```

vlan-domain member dom1
vcenter 192.168.66.2 datacenter prodDC
  username administrator password *****
  vshield 123.4.5.6
    username administrator password *****
    vxlan pool 10000-12000
    vxlan multicast-pool 224.3.4.5-224.5.6.7
  exit
exit
configure-dvs
exit
exit

```

Creating Tenants, VRFs, and Bridge Domains

Creating a Tenant, VRF, and Bridge Domain Using the NX-OS Style CLI

This section provides information on how to create tenants, VRFs, and bridge domains.



Note Before creating the tenant configuration, you must create a VLAN domain using the **vlan-domain** command and assign the ports to it.

Procedure

Step 1 Create a VLAN domain (which contains a set of VLANs that are allowable in a set of ports) and allocate VLAN inputs, as follows:

Example:

In the following example ("exampleCorp"), note that VLANs 50 - 500 are allocated.

```

apic1# configure
apic1(config)# vlan-domain dom_exampleCorp
apic1(config-vlan)# vlan 50-500
apic1(config-vlan)# exit

```

Step 2 Once the VLANs have been allocated, specify the leaf (switch) and interface for which these VLANs can be used. Then, enter "vlan-domain member" and then the name of the domain you just created.

Example:

In the following example, these VLANs (50 - 500) have been enabled on leaf 101 on interface ethernet 1/2-4 (three ports including 1/2, 1/3, and 1/4). This means that if you are using this interface, you can use VLANS 50-500 on this port for any application that the VLAN can be used for.

```

apic1(config-vlan)# leaf 101
apic1(config-vlan)# interface ethernet 1/2-4
apic1(config-leaf-if)# vlan-domain member dom_exampleCorp
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit

```

Step 3 Create a tenant in global configuration mode, as shown in the following example:

Example:

```
apicl(config)# tenant exampleCorp
```

Step 4 Create a private network (also called VRF) in tenant configuration mode as shown in the following example:

Example:

```
apicl(config)# tenant exampleCorp
apicl(config-tenant)# vrf context exampleCorp_v1
apicl(config-tenant-vrf)# exit
```

Step 5 Create a bridge domain (BD) under the tenant, as shown in the following example:

Example:

```
apicl(config-tenant)# bridge-domain exampleCorp_b1
apicl(config-tenant-bd)# vrf member exampleCorp_v1
apicl(config-tenant-bd)# exit
```

Note In this case, the VRF is "exampleCorp_v1".

Step 6 Allocate IP addresses for the BD (ip and ipv6), as shown in the following example.

Example:

```
apicl(config-tenant)# interface bridge-domain exampleCorp_b1
apicl(config-tenant-interface)# ip address 172.1.1.1/24
apicl(config-tenant-interface)# ipv6 address 2001:1:1::1/64
apicl(config-tenant-interface)# exit
```

What to Do Next

The next section describes how to add an application profile, create an application endpoint group (EPG), and associate the EPG to the bridge domain.

Related Topics

[Configuring a VLAN Domain Using the NX-OS Style CLI, on page 85](#)

Creating an Application Profile and EPG Using the NX-OS Style CLI

Before You Begin

Before you can create an application profile and an application endpoint group (EPG), you must create a VLAN domain, tenant, VRF, and BD (as described in the previous section).

Procedure

Step 1 Create an application profile, as shown in the following example ("exampleCorp_web1"):

Example:

```
apic1(config)# tenant exampleCorp
apic1(config-tenant)# application exampleCorp_web1
```

Step 2 Create an EPG under the application, as shown in the following example ("exampleCorp_webepg1"):

Example:

```
apic1(config-tenant-app)# epg exampleCorp_webepg1
```

Step 3 Associate the EPG to the bridge domain, shown as follows:

Example:

```
apic1(config-tenant-app-epg)# bridge-domain member exampleCorp_b1
apic1(config-tenant-app-epg)# exit
apic1(config-tenant-app)# exit
apic1(config-tenant)# exit
```

Note Every EPG belongs to a BD. An EPG can belong to a BD from the same tenant (or) from tenant Common. If you look at the chain, the lowest end is the EPG, and above that is the BD. The BD belongs to a VRF, and the VRF belongs to the tenant.

What to Do Next

These examples have shown how to configure an application EPG on a tenant. The next section discusses how to map a VLAN on a port to the EPG.

Mapping a VLAN on a Port to the EPG Using the NX-OS Style CLI

This step discusses how to open or enable a VLAN on a port on the leaf switch and associate it with an application EPG. The pre-requisite for this step is that the interface should be a member of a VLAN domain (vlan-domain), which contains this VLAN. Creation of VLAN domain is discussed in [Configuring a VLAN Domain Using the NX-OS Style CLI, on page 85](#).

Procedure

Step 1 Enter into leaf configuration mode by providing the ID of the leaf switch.

Example:

```
apic1(config)# leaf 101
```

Note To apply the same configuration on multiple leaf switches, "-" or "," separated IDs can be used (such as leaf 101-103).

Step 2 Enter the mode shown as follows using the previous example of "interface ethernet 1/2".

Example:

```
apic1(config-leaf)# interface ethernet 1/2
```

Step 3 Enter the command "switchport trunk allowed vlan" followed by the VLAN, then the tenant, application, and the EPG (shown as follows using the previous examples for each):

Example:

```
apic1(config-leaf-if)#switchport trunk allowed vlan 50 tenant exampleCorp application
exampleCorp_web1 epg exampleCorp_webepg1
```

Deploying an Application Policy

Three-Tier Application Deployment

Application profiles enable you to model application requirements that the APIC then automatically renders in the network and data center infrastructure. The application profiles enable administrators to approach the resource pool in terms of applications rather than infrastructure building blocks. The application profile is a container that holds EPGs that are logically related to one another. EPGs can communicate with the other EPGs in the same application profile and with EPGs in other application profiles.

Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

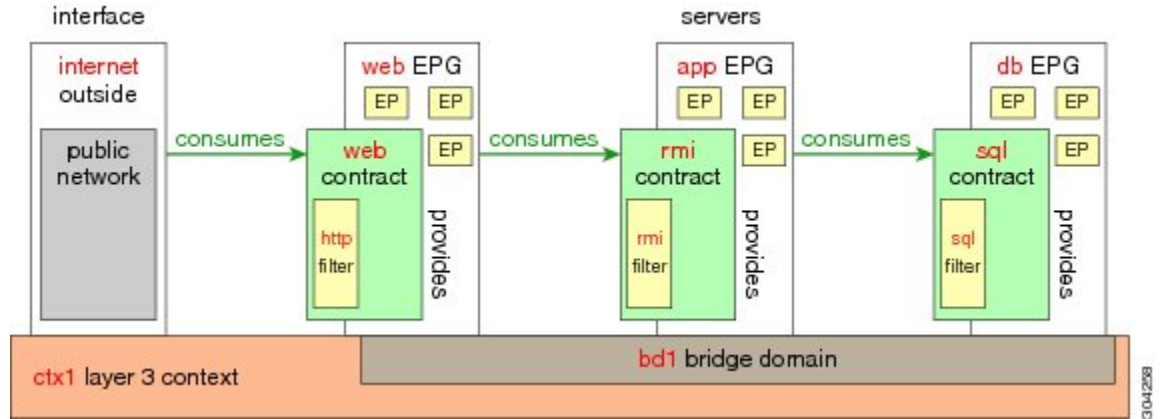
An access list (also referred to as a "filter") specifies the data protocols to be allowed or denied by a contract that contains the access list. A contract can contain multiple subjects. A subject can be used to realize uni- or bidirectional access lists. A unidirectional access list that is used in one direction, either from consumer-to-provider (IN) or from provider-to-consumer (OUT) access list. A bi-directional access list is the same access list that is used in both directions. It is not reflexive.

To deploy an application policy, you must create the required application profiles, access lists (filters), and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

The web server has the HTTP access list, the application server has the Remote Method Invocation (RMI) access list, and the database server has the Structured Query Language (SQL) access list. The application server consumes the SQL contract to communicate with the database server. The web server consumes the RMI contract to communicate with the application server. The traffic enters from the web server and

communicates with the application server. The application server then communicates with the database server, and the traffic can also communicate externally.

Figure 8: Three-Tier Application Diagram



Parameters to Create an Access List for HTTP

The parameters to create an access list (filter) for http in this example is as follows:

Parameter Name	Access List (Filter) for HTTP
Name	http
Number of Entries	2
Entry Name	Dport-80 Dport-443
Ethertype	IP
Protocol	tcp tcp
Destination Port	http https

Parameters to Create an Access List for RMI and SQL

The parameters to create filters for RMI and SQL in this example are as follows:

Parameter Name	Filter for RMI	Filter for SQL
Name	rmi	sql

Parameter Name	Filter for RMI	Filter for SQL
Number of Entries	1	1
Entry Name	Dport-1099	Dport-1521
Ethertype	IP	IP
Protocol	tcp	tcp
Destination Port	1099	1521

Example Application Profile Database

The application profile database in this example is as follows:

EPG	Provided Contracts	Consumed Contracts
web	web	rmi
app	rmi	sql
db	sql	--

Deploying an Application Policy Using the NX-OS Style CLI

The port the EPG uses must belong to one of the VM Managers (VMM) or physical domains associated with the EPG.

Procedure

Step 1 To get into the configuration mode using the NX-OS CLI, enter the following:

Example:

```
apicl#configure
apicl(config)#
```

Step 2 Create an application network profile for the tenant.
The application network profile in this example is OnlineStore.

Example:

```
apicl(config)# tenant exampleCorp
apicl(config-tenant)# application OnlineStore
apicl(config-tenant-app)#
```

Step 3 Create application web, db, and app EPGs for this application network profile of the tenant.

Example:

```
apicl(config-tenant-app)# epg web
apicl(config-tenant-app-epg)# exit
apicl(config-tenant-app)# epg db
apicl(config-tenant-app-epg)# exit
apicl(config-tenant-app)# epg app
apicl(config-tenant-app-epg)# exit
```

- Step 4** Get back into the tenant mode to create an access list (filter) for different traffic types between these EPGs.

Example:

```
apicl(config-tenant-app)# exit
```

- Step 5** Create an access list (filter) for the http and https traffic.

Example:

```
apicl(config-tenant)# access-list http
apicl(config-tenant-acl)# match tcp dest 80
apicl(config-tenant-acl)# match tcp dest 443
apicl(config-tenant-acl)# exit
```

- Step 6** Create an access list (filter) for Remote Method Invocation (RMI) traffic.

Example:

```
apicl(config-tenant)# access-list rmi
apicl(config-tenant-acl)# match tcp dest 1099
apicl(config-tenant-acl)# exit
```

- Step 7** Create an access list (filter) for the SQL/database traffic.

Example:

```
apicl(config-tenant)# access-list sql
apicl(config-tenant-acl)# match tcp dest 1521
apicl(config-tenant-acl)# exit
```

- Step 8** Create the contracts and assign an access group (filters) for RMI traffic between EPGs.

Example:

```
apicl(config)# tenant exampleCorp
apicl(config-tenant)# contract rmi
apicl(config-tenant-contract)# subject rmi
apicl(config-tenant-contract-subj)# access-group rmi both
apicl(config-tenant-contract-subj)# exit
apicl(config-tenant-contract)# exit
```

- Step 9** Create the contracts and assign an access group (filters) for web traffic between EPGs.

Example:

```
apicl(config-tenant)# contract web
apicl(config-tenant-contract)# subject web
apicl(config-tenant-contract-subj)# access-group http both
apicl(config-tenant-contract-subj)# exit
```

- Step 10** Create the contracts and assign an access group (filters) for SQL traffic between EPGs.

Example:

```

apicl (config-tenant) # contract sql
apicl (config-tenant-contract) # subject sql
apicl (config-tenant-contract-subj) # access-group sql both
apicl (config-tenant-contract-subj) # exit
apicl (config-tenant-contract) # exit

```

Step 11 Attach the bridge domain and contracts to the web EPG.

Example:

```

apicl (config-tenant) # application OnlineStore
apicl (config-tenant-app) # epg web
apicl (config-tenant-app-epg) # bridge-domain member exampleCorp_b1
apicl (config-tenant-app-epg) # contract consumer rmi
apicl (config-tenant-app-epg) # contract provider web
apicl (config-tenant-app-epg) # exit

```

Step 12 Attach the bridge domain and contracts to the db EPG.

Example:

```

apicl (config-tenant-app) # epg db
apicl (config-tenant-app-epg) # bridge-domain member exampleCorp_b1
apicl (config-tenant-app-epg) # contract provider sql
apicl (config-tenant-app-epg) # exit

```

Step 13 Attach the bridge domain and contracts to the application EPG.

Example:

```

apicl (config-tenant-app) # epg app
apicl (config-tenant-app-epg) # bridge-domain member exampleCorp_b1

```

Step 14 Associate the provider contracts to the application EPGs.

Example:

```

apicl (config-tenant-app-epg) # contract provider rml
apicl (config-tenant-app-epg) # contract consumer sql
apicl (config-tenant-app-epg) # exit
apicl (config-tenant-app) # exit
apicl (config-tenant) # exit

```

Step 15 Associate the ports and VLANs to the EPGs app, db, and web.

Example:

```

apicl (config) # leaf 103
apicl (config-leaf) # interface ethernet 1/2-4
apicl (config-leaf-if) # vlan-domain member exampleCorp
apicl (config-leaf) # exit
apicl (config) # leaf 103
apicl (config-leaf) # interface ethernet 1/2
apicl (config-leaf-if) # switchport
access trunk vlan
apicl (config-leaf-if) # switchport trunk allowed vlan 100 tenant exampleCorp application
OnlineStore epg app
apicl (config-leaf-if) # exit
apicl (config-leaf) # interface ethernet 1/3
apicl (config-leaf-if) # switchport trunk allowed vlan 101 tenant exampleCorp application
OnlineStore epg db

```

```

apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/4
apic1(config-leaf-if)# switchport trunk allowed vlan 102 tenant exampleCorp application
OnlineStore epg web
apic1(config-leaf-if)# exit

```

Configuring External L3 Connectivity for Tenants

Configuring an MP-BGP Route Reflector for the ACI Fabric

To distribute routes within the ACI fabric, an MP-BGP process must first be operating, and the spine switches must be configured as BGP route reflectors.

The following is an example of an MP-BGP route reflector configuration:



Note

In this example, the BGP fabric ASN is 100. Spine switches 104 and 105 are chosen as MP-BGP route-reflectors.

```

apic1(config)# bgp-fabric
apic1(config-bgp-fabric)# asn 100
apic1(config-bgp-fabric)# route-reflector spine 104,105

```

Creating an OSPF External Routed Network for a Tenant Using the NX-OS CLI

Configuring external routed network connectivity involves the following steps:

- 1 Create a VRF under Tenant.
- 2 Configure L3 networking configuration for the VRF on the border leaf switches, which are connected to the external routed network. This configuration includes interfaces, routing protocols (BGP, OSPF, EIGRP), protocol parameters, route-maps.
- 3 Configure policies by creating external-L3 EPGs under tenant and deploy these EPGs on the border leaf switches. External routed subnets on a VRF which share the same policy within the ACI fabric form one "External L3 EPG" or one "prefix EPG".

Configuration is realized in two modes:

- Tenant mode: VRF creation and external-L3 EPG configuration
- Leaf mode: L3 networking configuration and external-L3 EPG deployment

The following steps are for creating an OSPF external routed network for a tenant. To create an OSPF external routed network for a tenant, you must choose a tenant and then create a VRF for the tenant.



Note The examples in this section show how to provide external routed connectivity to the "web" epg in the "OnlineStore" application for tenant "exampleCorp".

Procedure

Step 1 Configure the VLAN domain.

Example:

```
apicl(config)# vlan-domain dom_exampleCorp
apicl(config-vlan)# vlan 5-1000
apicl(config-vlan)# exit
```

Step 2 Configure the tenant VRF and enable policy enforcement on the VRF.

Example:

```
apicl(config)# tenant exampleCorp
apicl(config-tenant)# vrf context
    exampleCorp_v1
apicl(config-tenant-vrf)# contract enforce
apicl(config-tenant-vrf)# exit
```

Step 3 Configure the tenant BD and mark the gateway IP as "public". The entry "scope public" makes this gateway address available for advertisement through the routing protocol for external-L3 network.

Example:

```
apicl(config-tenant)# bridge-domain exampleCorp_b1
apicl(config-tenant-bd)# vrf member exampleCorp_v1
apicl(config-tenant-bd)# exit
apicl(config-tenant)# interface bridge-domain exampleCorp_b1
apicl(config-tenant-interface)# ip address 172.1.1.1/24 scope public
apicl(config-tenant-interface)# exit
```

Step 4 Configure the VRF on a leaf.

Example:

```
apicl(config)# leaf 101
apicl(config-leaf)# vrf context tenant exampleCorp vrf exampleCorp_v1
```

Step 5 Configure the OSPF area and add the route map.

Example:

```
apicl(config-leaf)# router ospf default
apicl(config-leaf-ospf)# vrf member tenant exampleCorp vrf exampleCorp_v1
apicl(config-leaf-ospf-vrf)# area 0.0.0.1 route-map map100 out
apicl(config-leaf-ospf-vrf)# exit
apicl(config-leaf-ospf)# exit
```

Step 6 Assign the VRF to the interface (sub-interface in this example) and enable the OSPF area.

Example:

Note For the sub-interface configuration, the main interface (ethernet 1/11 in this example) must be converted to an L3 port through “no switchport” and assigned a vlan-domain (dom_exampleCorp in this example) that contains the encapsulation VLAN used by the sub-interface. In the sub-interface ethernet1/11.500, 500 is the encapsulation VLAN.

```
apic1(config-leaf)# interface ethernet 1/11
apic1(config-leaf-if)# no switchport
apic1(config-leaf-if)# vlan-domain member dom_exampleCorp
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/11.500
apic1(config-leaf-if)# vrf member tenant exampleCorp vrf exampleCorp_v1
apic1(config-leaf-if)# ip address 157.10.1.1/24
apic1(config-leaf-if)# ip router ospf default area 0.0.0.1
```

Step 7 Configure the external-L3 EPG policy. This includes the subnet to match for identifying the external subnet and consuming the contract to connect with the epg "web".

Example:

```
apic1(config)# tenant t100
apic1(config-tenant)# external-l3 epg l3epg100
apic1(config-tenant-l3ext-epg)# vrf member v100
apic1(config-tenant-l3ext-epg)# match ip 145.10.1.0/24
apic1(config-tenant-l3ext-epg)# contract consumer web
apic1(config-tenant-l3ext-epg)# exit
apic1(config-tenant)#exit
```

Step 8 Deploy the external-L3 EPG on the leaf switch.

Example:

```
apic1(config)# leaf 101
apic1(config-leaf)# vrf context tenant t100 vrf v100
apic1(config-leaf-vrf)# external-l3 epg l3epg100
```

Configuring Server or Service Policies

Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets. If the client is on an ESX hypervisor with a deployed vShield Domain profile, then the use of a DHCP relay policy configuration is mandatory.

When a vShield controller deploys a Virtual Extensible Local Area Network (VXLAN), the hypervisor hosts create a kernel (vmkN, virtual tunnel end-point [VTEP]) interface. These interfaces need an IP address in the infrastructure tenant that uses DHCP. Therefore, you must configure a DHCP relay policy so that the APIC can act as the DHCP server and provide these IP addresses.

When an ACI fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the ACI fabric must support Option 82.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the NX-OS Style CLI

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Before You Begin

Ensure that Layer 2 or Layer 3 connectivity is configured to reach the DHCP server address.

Procedure

Configure DHCP server policy settings for the APIC infrastructure traffic.

Example: DHCP Relay Policy for an Endpoint Group

```
apic1(config)# tenant infra
apic1(config-tenant)# template dhcp relay policy DhcpRelayP
apic1(config-tenant-template-dhcp-relay)# ip address 10.0.0.1 tenant infra application access epg
default
apic1(config-tenant-template-dhcp-relay)# exit
apic1(config-tenant)# interface bridge-domain default
apic1(config-tenant-interface)# dhcp relay policy tenant DhcpRelayP
apic1(config-tenant-interface)# exit
```

Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.
- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.
- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking** > **VRF** policy configuration in the tenants configuration.

Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be attached to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/fabric/dnsp-default.xml -->
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The `gai.conf` settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a `gai.conf` to prioritize IPv4 over IPv6:

```
# Generated by APIC
label ::1/128      0
label ::/0         1
label 2002::/16   2
label ::/96        3
label ::ffff:0:0/96 4
precedence ::1/128      50
precedence ::/0         40
precedence 2002::/16   30
precedence ::/96        20
# For APICs preferring IPv4 connections, change the value to 100.
precedence ::ffff:0:0/96 10
```

Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPV4) or AAAA records (IPV6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to `/etc/resolv.conf`.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in `/etc/resolv.conf`. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in `/etc/resolv.conf`). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add “options single-request-reopen” to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of `resolv.conf` in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (glibc) will use the IPv6 network by default because `getaddrinfo()` will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (`::ffff/96`). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the glibc IPv6 selection preference for `getaddrinfo()` in `/etc/gai.conf`.

In order to allow glibc to return multiple addresses when using `/etc/hosts`, “multi on” should be added to the `/etc/hosts` file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

Configuring a DNS Service Policy to Connect with DNS Providers Using the NX-OS Style CLI

Procedure

Step 1 In the NX-OS CLI, get into configuration mode, shown as follows:

Example:

```
apic1# configure
apic1(config)#
```

Step 2 Configure a DNS server policy.

Example:

```
apic1(config)# dns
apic1(config-dns)# address 172.21.157.5 preferred
apic1(config-dns)# address 172.21.157.6
apic1(config-dns)# domain company.local default
apic1(config-dns)# use-vrf oob-default
```

Step 3 Configure a DNS profile label on any VRF where you want to use the DNS profile.

Example:

```
apic1(config)# tenant mgmt
apic1(config-tenant)# vrf context oob
apic1(config-tenant-vrf)# dns label default
```

Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the NX-OS Style CLI

Procedure

Step 1 Verify the configuration for the default DNS profile.

Example:

```
apic1# show running-config dns

# Command: show running-config dns
# Time: Sat Oct 3 00:23:52 2015
dns
  address 172.21.157.5 preferred
  address 172.21.157.6
  domain company.local default
  use-vrf oob-default
  exit
```

Step 2 Verify the configurations for the DNS labels.

Example:

```
apic1# show running-config tenant mgmt vrf context oob

# Command: show running-config tenant mgmt vrf context oob
# Time: Sat Oct 3 00:24:36 2015
tenant mgmt
  vrf context oob
    dns label default
  exit
exit
```

Step 3 Verify that the applied configuration is operating on the fabric controllers.

Example:

```
apic1# cat /etc/resolv.conf
# Generated by IFC

nameserver 172.21.157.5
nameserver 172.21.157.6
```



Using the Advanced GUI

This chapter contains the following sections:

- [Toggling Between Basic and Advanced GUI Modes, page 107](#)
- [About Getting Started with APIC Examples, page 108](#)
- [About Switch Discovery with the APIC, page 109](#)
- [Configuring Network Time Protocol, page 112](#)
- [Creating User Accounts, page 115](#)
- [Adding Management Access, page 120](#)
- [Configuring a VMM Domain, page 130](#)
- [Creating Tenants, VRF, and Bridge Domains, page 137](#)
- [Configuring Server or Service Policies, page 140](#)
- [Configuring External Connectivity for Tenants, page 147](#)
- [Deploying an Application Policy, page 150](#)

Toggling Between Basic and Advanced GUI Modes

When logged in to the APIC GUI, you can verify the GUI mode you are in. The mode you have entered is displayed in the top right corner of the GUI. You can choose to operate in one of two modes:

Caution: Cisco recommends that you do not mix configuration modes (Advanced or Basic). When you make a configuration in either mode and change the configuration using the other mode, unintended changes can occur. For example, if you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.

- **Basic Mode**—For information about tasks that you perform in Basic Mode, see the chapter, *Getting Started with APIC Using the Basic GUI*.
- **Advanced Mode**—For information about tasks that you perform in Advanced Mode, see the chapter, *Getting Started with APIC Using the Advanced GUI*.

You can also change from one GUI mode to another or toggle between modes as follows:

- 1 In the GUI, click the **welcome**, **<login_name>** drop-down list and choose Toggle GUI Mode.
- 2 In the **Warning** dialog box, click Yes .
- 3 Wait for the application to complete loading and display the GUI in the changed mode.

**Caution**

Changes made through the APIC Basic GUI can be seen, but cannot be modified in the Advanced GUI, and changes made in the Advanced GUI cannot be rendered in the Basic GUI. The Basic GUI is kept synchronized with the NX-OS style CLI, so that if you make a change from the NX-OS style CLI, these changes are rendered in the Basic GUI, and changes made in the Basic GUI are rendered in the NX-OS style CLI, but the same synchronization does not occur between the Advanced GUI and the NX-OS style CLI. See the following examples:

- Do not mix Basic and Advanced GUI modes. If you apply an interface policy to two ports using Advanced mode and then change the settings of one port using Basic mode, your changes might be applied to both ports.
- Do not mix the Advanced GUI and the CLI, when doing per-interface configuration on APIC. Configurations performed in the GUI, may only partially work in the NX-OS CLI.

For example, if you configure a switch port in the GUI at **Tenants > tenant-name > Application Profiles > application-profile-name > Application EPGs > EPG-name > Static Ports > Deploy Static EPG on PC, VPC, or Interface**

Then you use the show running-config command in the NX-OS style CLI, you receive output such as:

```
leaf 102
interface ethernet 1/15
switchport trunk allowed vlan 201 tenant t1 application ap1 epg ep1
exit
exit
```

If you use these commands to configure a static port in the NX-OS style CLI, the following error occurs:

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/15
apic1(config-leaf-if)# switchport trunk allowed vlan 201 tenant t1 application ap1
epg ep1
No vlan-domain associated to node 102 interface ethernet1/15 encap vlan-201
```

This occurs because the CLI has validations that are not performed by the APIC GUI. For the commands from the show running-config command to function in the NX-OS CLI, a vlan-domain must have been previously configured. The order of configuration is not enforced in the GUI.

- Do not make changes with the Basic GUI or the NX-OS CLI before using the Advanced GUI. This may also inadvertently cause objects to be created (with names prepended with `_ui_`) which cannot be changed or deleted in the Advanced GUI.

For the steps to remove such objects, see *Troubleshooting Unwanted `_ui_` Objects* in the *APIC Troubleshooting Guide*.

About Getting Started with APIC Examples

The steps in several examples in this guide include a parameter name. These parameter names are provided as examples for convenience and ease of your understanding, and it is not required for you to use them.

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

**Note**

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>.

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

**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.

Registering the Unregistered Switches 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.

Procedure

- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
- Step 2** In the **Navigation** pane, click **Fabric Membership**.

In the **Work** pane, in the **Fabric Membership** table, a single leaf switch is displayed with an ID of 0. It is the leaf switch that is connected to apic1.

Step 3 Configure the ID by double-clicking the leaf switch row, and performing the following actions:

- a) In the **ID** field, add the appropriate ID (leaf1 is ID 101, and leaf 2 is ID 102).
The ID must be a number that is greater than 100 because the first 100 IDs are for APIC appliance nodes.
- b) In the **Switch Name** field, add the name of the switch, and click **Update**.
Note After an ID is assigned, it cannot be updated. The switch name can be updated by double-clicking the name and updating the **Switch Name** field.

An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod.

Step 4 Monitor the **Work** pane until one or more spine switches appear.

Step 5 Configure the ID by double-clicking the spine switch row, and perform the following actions:

- a) In the **ID** field, add the appropriate ID (spine1 is ID 203 and spine 2 is ID 204).
Note It is recommended that leaf nodes and spine nodes be numbered differently. For example, number spines in the 200 range and number leaves in the 100 range.
- b) In the **Switch Name** field, add the name of the switch, and click **Update**.

An IP address gets assigned to the switch, and in the **Navigation** pane, the switch is displayed under the pod. Wait until all remaining switches appear in the **Node Configurations** table before you go to the next step.

Step 6 For each switch listed in the **Fabric Membership** table, perform the following steps:

- a) Double-click the switch, enter an **ID** and a **Name**, and click **Update**.
- b) Repeat for the next switch in the list.

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

Procedure

Step 1 On the menu bar, choose **FABRIC > INVENTORY**.

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

The switches in the fabric are displayed with their node IDs. In the **Work** pane, 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

Procedure

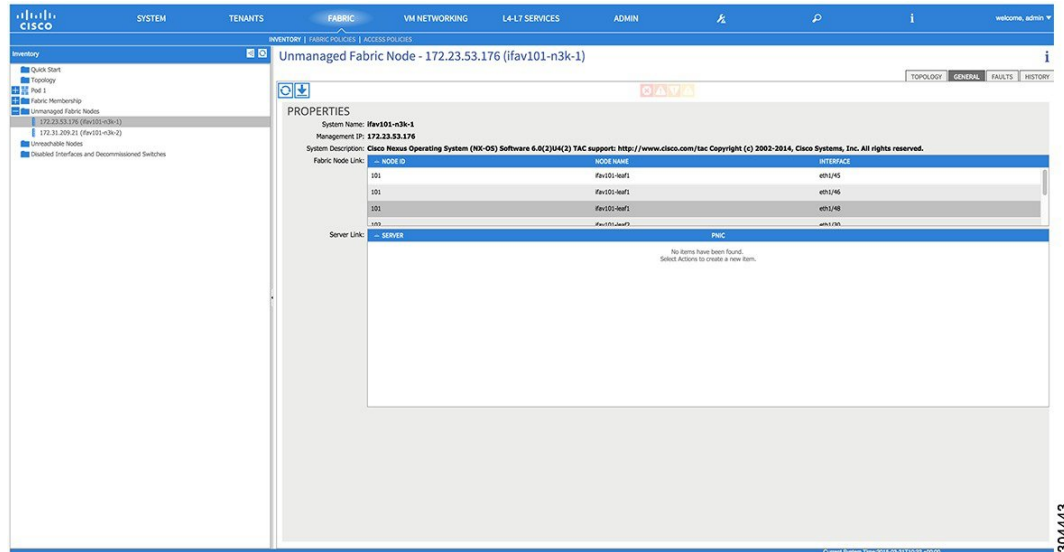
- Step 1** On the menu bar, choose **FABRIC > INVENTORY**.
 - Step 2** In the **Navigation** pane, choose the pod that you want to view.
 - Step 3** In the **Work** pane, click the **TOPOLOGY** tab.
The displayed diagram shows all attached switches, APIC instances, and links.
 - Step 4** (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
To return to the topology diagram, in the upper left corner of the **Work** pane, click the **Previous View** icon.
 - Step 5** (Optional) To refresh the topology diagram, in the upper left corner of the **Work** pane, click the **Refresh** icon.
-

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) or Cisco Discovery Protocol (CDP) 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. The following figure shows the APIC GUI displaying unmanaged switches in the **Fabric > Inventory** view.

Figure 9: Unmanaged Layer 2 Switches in the APIC Fabric Inventory



Configuring Network Time Protocol

Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

In-Band and Out-of-Band Management NTP



Note

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
 - See the Adding Management Access section in this guide for information about in-band management access and out-of-band management access.
-
- Out-of-band management NTP—When an ACI fabric is deployed with out-of-band management, each node of the fabric, inclusive of spines, leaves, and all members of the APIC cluster, is managed from outside the ACI fabric. This IP reachability will be leveraged so that each node can individually query the same NTP server as a consistent clock source. To configure NTP, a Date and Time policy must be created that references an out-of-band management endpoint group. Date and Time policies are confined to a single pod and must be deployed across all pods provisioned in the ACI fabric. Currently only one pod per ACI fabric is allowed.
 - In-Band Management NTP—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication. The steps used to configure in-band management policies are identical to those used to establish an out-of-band management policy. The distinction is around how to allow the fabric to connect to the NTP server.

NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The `gai.conf` can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

Configuring NTP Using the Advanced GUI

Procedure

- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**.
- Step 2** In the **Navigation** pane, choose **Pod Policies > Policies**.
- Step 3** In the **Work** pane, choose **Actions > Create Date and Time Policy**.
- Step 4** In the **Create Date and Time Policy** dialog box, perform the following actions:
 - a) Enter a name for the policy to distinguish between the different NTP configurations in your environment. Click **Next**.
 - b) Click the + sign to specify the NTP server information (provider) to be used.
 - c) In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.

- If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.
- In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.

Repeat the steps for each provider that you want to create.

- Step 5** In the **Navigation** pane, choose **Pod Policies > Policy Groups**.
- Step 6** In the **Work** pane, choose **Actions > Create Pod Policy Group**.
- Step 7** In the **Create Pod Policy Group** dialog box, perform the following actions:
- Enter a name for the policy group.
 - In the **Date Time Policy** field, from the drop down list, choose the NTP policy that you created earlier. Click **Submit**.
The pod policy group is created. Alternatively, you can use the default pod policy group.
- Step 8** In the **Navigation** pane, choose **Pod Policies > Profiles**.
- Step 9** In the **Work** pane, double-click the desired pod selector name.
- Step 10** In the Properties area, from the **Fabric Policy Group** drop down list, choose the pod policy group you created. Click **Submit**.
-

Verifying NTP Operation Using the GUI

Procedure

- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**.
- Step 2** In the **Navigation** pane, choose **Pod Policies > Policies > Date and Time > ntp_policy > server_name**.
The *ntp_policy* is the previously created policy. An IPv6 address is supported in the Host Name/IP address field. If you enter a hostname and it has an IPv6 address set, you must implement the priority of IPv6 address over IPv4 address.
- Step 3** In the **Work** pane, verify the details of the server.
-

Verifying NTP Policy Deployed to Each Node Using the CLI

Procedure

- Step 1** SSH to an APIC in the fabric.
- Step 2** Press the Tab key two times after entering the attach command to list all the available node names:

Example:

```
admin@apic1:~> attach <Tab> <Tab>
```

Step 3 Log in one of the nodes using the same password that you used to access the APIC.

Example:

```
admin@apic1:~> attach node_name
```

Step 4 View the NTP peer status.

Example:

```
leaf-1# show ntp peer-status
```

A reachable NTP server has its IP address prefixed by an asterisk (*), and the delay is a non-zero value.

Step 5 Repeat steps 3 and 4 to verify each node in the fabric.

Creating User Accounts

Configuring a Local User

In the initial configuration script, the admin account is configured and the admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

Configuring a Remote User

Instead of configuring local users, you can point the APIC at the centralized enterprise credential datacenter. The APIC supports Lightweight Directory Access Protocol (LDAP), active directory, RADIUS, and TACACS+.

**Note**

When an APIC is in minority (disconnected from the cluster), remote logins can fail because the ACI is a distributed system and the user information is distributed across APICS. Local logins, however, continue to work because they are local to the APIC.

To configure a remote user authenticated through an external authentication provider, you must meet the following prerequisites:

- The DNS configuration should have already been resolved with the hostname of the RADIUS server.
- You must configure the management subnet.

Configuring a Local User Using the GUI

Before You Begin

- The ACI fabric is installed, APIC controllers are online, and the APIC cluster is formed and healthy.
- As appropriate, the security domain(s) that the user will access are defined. For example, if the new user account will be restricted to accessing a tenant, the tenant domain is tagged accordingly.
- An APIC user account is available that will enable the following:
 - Creating the TACACS+ and TACACS+ provider group.
 - Creating the local user account in the target security domain(s). If the target domain is `all`, the login account used to create the new local user must be a fabric-wide administrator that has access to `all`. If the target domain is a tenant, the login account used to create the new local user must be a tenant administrator that has full read write access rights to the target tenant domain.

Procedure

- Step 1** On the menu bar, choose **ADMIN > AAA**.
- Step 2** In the **Navigation** pane, click **AAA Authentication**.
- Step 3** In the **Work** pane, verify that in the default **Authentication** field, the **Realm** field displays as Local.
- Step 4** In the **Navigation** pane, expand **Security Management > Local Users**.
The admin user is present by default.
- Step 5** In the **Navigation** pane, right-click **Create Local User**.
- Step 6** In the **User Identity** dialog box, enter a **Login ID** and **Password** for the user, and click **Next**.
- Step 7** In the **Security** dialog box, choose the desired security domain for the user, and click **Next**.
- Step 8** In the **Roles** dialog box, click the radio buttons to choose the roles for your user, and click **Next**.
You can provide read-only or read/write privileges.
- Step 9** In the **User Identity** dialog box, perform the following actions:
- a) In the **Login ID** field, add an ID.
 - b) In the **Password** field, enter the password.
- At the time a user sets their password, the APIC validates it against the following criteria:
- Minimum password length is 8 characters.
 - Maximum password length is 64 characters.
 - Has fewer than three consecutive repeated characters.
 - Must have characters from at least three of the following characters types: lowercase, uppercase, digit, symbol.
 - Does not use easily guessed passwords.
 - Cannot be the username or the reverse of the username.

- Cannot be any variation of cisco, isco or any permutation of these characters or variants obtained by changing the capitalization of letters therein.

- In the **Confirm Password** field, confirm the password.
- Click **Finish**.

Step 10 In the **Navigation** pane, click the name of the user that you created. In the **Work** pane, expand the + sign next to your user in the **Security Domains** area. The access privileges for your user are displayed.

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

To configure a Cisco AV Pair on an external authentication server, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair format is as follows:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(16003)
```

The first av-pair format has no UNIX user ID, while the second one does. Both are correct if all remote users have the same role and mutual file access is acceptable. If the UNIX user ID is not specified, ID 23999 is applied by the APIC system, and more than one role/read privilege is specified to any AV Pair user. This can cause users to have higher or lower permissions than configured through the group settings.



Note

The APIC Cisco AV-pair format is compatible and can co-exist with other Cisco AV-pair formats. APIC will pick up the first matching AV-pair from all the AV-pairs.

The APIC supports the following regexes:

```
shell:domains\\s* [=:] \\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31}) (\\(\\d+\\)) $
shell:domains\\s* [=:] \\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31}) $
```

Examples:

- Example 1: A Cisco AV Pair that contains a single Login domain with only writeRoles:

```
shell:domains=domainA/writeRole1|writeRole2/
```

- Example 2: A Cisco AV Pair that contains a single Login domain with only readRoles:

```
shell:domains=domainA//readRole1|readRole2
```

**Note**

The "/" character is a separator between writeRoles and readRoles per Login domain and is required even if only one type of role is to be used.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

```
aaa-network-admin Cleartext-Password := "<password>"
Cisco-avpair = "shell:domains = all/aaa/read-all(16001)"
```

Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs

Procedure

-
- Step 1** On the menu bar, click **ADMIN > AAA**.
 - Step 2** In the **Navigation** pane, click **AAA Authentication**.
 - Step 3** In the **Work** pane, in the **Properties** area, from the **Remote user login policy** drop-down list, choose **Assign Default Role**.

The default value is **No Login**. The **Assign Default Role** option assigns the minimal read-only privileges to users that have missing or bad Cisco AV Pairs. Bad AV Pairs are those AV Pairs that fail the parsing rules.

Best Practice for Assigning AV Pairs

As best practice, Cisco recommends that you assign unique UNIX user ids in the range 16000-23999 for the AV Pairs that are assigned to users when in bash shell (using SSH, Telnet or Serial/KVM consoles). If a situation arises when 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 remote users with a UNIX ID of 23999.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is logged in using Secure Shell (SSH) or Telnet.

Procedure

Configure an AV pair on the external authentication server.

The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

```
* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2 (8101)
```

These are the boost regexes supported by APIC:

```
uid_regex("shell:domains\\s*[:]\\s*(\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31}) (\\d+\\)");
regex("shell:domains\\s*[:]\\s*(\\S+?/\\S*?/\\S*?) (,\\S+?/\\S*?/\\S*?) {0,31})");
```

The following is an example:

```
shell:domains = coke/tenant-admin/read-all,pepsi//read-all (16001)
```

Configuring a Remote User Using the GUI

Before You Begin

- The DNS configuration must have resolved the RADIUS server hostname in order for the fabric controller to reach the server.
- The APIC should have the external management subnet policy configured so that it is able to reach the RADIUS server.

Procedure

-
- Step 1** On the menu bar, choose **ADMIN > AAA**. In the **Navigation** pane, expand **RADIUS Management**.
 - Step 2** Right-click **RADIUS Providers**, and click **Create RADIUS Provider**.
 - Step 3** In the **Create RADIUS Provider** dialog box, and perform the following actions:
 - a) In the **Host Name (or IP Address)** field, add the hostname.
 - b) In the **Authorization Port** field, add the port number required for authorization. This number depends on the RADIUS server configured.
 - c) Click the required **Authorization Protocol** radio button.
 - d) In the **Key** and **Confirm Key** fields, enter the preshared key. This key is the same information that is shared with the server key configured on the RADIUS server.
 - Step 4** In the **Navigation** pane, under **RADIUS Providers**, click the RADIUS provider that you created. Details about the configurations for the RADIUS provider are displayed in the **Work** pane.
 - Step 5** In the **Navigation** pane, right-click **RADIUS Provider Groups**, and click **Create RADIUS Provider Group**.
 - Step 6** In the **Create RADIUS Provider Group** dialog box, perform the following actions:
 - a) In the **Name** field, enter a name.
 - b) Expand the **Providers** field, and from the **Name** field drop-down list, choose the provider created earlier.
 - c) In the **Priority** field, assign a priority. Click **Update**, and click **Submit**. The radius provider group is created.
 - Step 7** In the **Navigation** pane, expand **AAA Authentication**, and right-click **Login Domain** to click **Create Login Domain**.
 - Step 8** In the **Create Login Domain** dialog box, perform the following actions:
 - a) In the **Name** field, enter a domain name.
 - b) In the **Realm** field drop-down list, choose the RADIUS realm.

- c) In the **RADIUS Provider Group** field drop-down list, choose the provider group that was created earlier. Click **Submit**.

The login domain is created and is now available for remote user login and configuration.

Adding Management Access

An APIC controller has two routes to reach the management network, one is by using the in-band management interface and the other is by using the out-of-band management interface.

- In-band management access—You can configure in-band management connectivity to the APIC and the ACI fabric. You first configure the VLANs that will be used by APIC when the APIC is communicating with the leaf switches, and then you configure the VLANs that the VMM servers will use to communicate with the leaf switches.
- Out-of-band management access—You can configure out-of-band management connectivity to the APIC and the ACI fabric. You configure an out-of-band contract that is associated with an out-of-band endpoint group (EPG), and attach the contract to the external network profile.



Note The APIC out-of-band management connection link must be 1 Gbps.

The APIC controller always selects the in-band management interface over the out-of-band management interface, if the in-band management interface is configured. The out-of-band management interface is used only when the in-band management interface is not configured, or if the destination address is on the same subnet as the out-of-band management subnet of the APIC.

The APIC management interface does not support an IPv6 address and cannot connect to an external IPv6 server through this interface.

Configuring the external management instance profile under the management tenant for in-band or out-of-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet.

IPv4/IPv6 Addresses and In-Band Policies

In-band management addresses can be provisioned on the APIC controller only through a policy (Postman REST API, NX-OS Style CLI, or GUI). Additionally, the in-band management addresses must be configured statically on each node.

IPv4/IPv6 Addresses in Out-of-Band Policies

Out-of-band management addresses can be provisioned on the APIC controller either at the time of bootstrap or by using a policy (Postman REST API, NX-OS Style CLI, GUI). Additionally, the out-of-band management addresses must be configured statically on each node or by specifying a range of addresses (IPv4/IPv6) to the entire cluster. IP addresses are randomly assigned from a range to the nodes in the cluster.

Configuring Management Access

Configuring In-Band Management Access Using the Advanced GUI



Note

IPv4 and IPv6 addresses are supported for in-band management access. IPv6 configurations are supported using static configurations (for both in-band and out-of-band). IPv4 and IPv6 dual in-band and out-of-band configurations are supported only through static configuration. For more information, see the KB article, *Configuring Static Management Access in Cisco APIC*.

Procedure

- Step 1** On the menu bar, choose **FABRIC > Access Policies**. In the **Navigation** pane, expand **Interface Policies**.
- Step 2** In the **Navigation** pane, right-click **Switch Policies** and choose **Configure Interface, PC and VPC**.
- Step 3** In the **Configure Interface, PC, and VPC** dialog box, to configure switch ports connected to APICs, perform the following actions:
- Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the APIC.
 - From the **Switches** field drop-down list, check the check boxes for the switches to which the APICs are connected. (leaf1 and leaf2).
 - In the **Switch Profile Name** field, enter a name for the profile (apicConnectedLeaves).
 - Click the + icon to configure the ports.
- A dialog box similar to the following image is displayed for the user to enter the content:

- Verify that in the **Interface Type** area, the **Individual** radio button is selected.
- In the **Interfaces** field, enter the ports to which APICs are connected.
- In the **Interface Selector Name** field, enter the name of the port profile (apicConnectedPorts).
- In the **Interface Policy Group** field, click the **Create One** radio button.
- In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).

- j) In the **Domain** field, click the **Create One** radio button.
- k) In the **Domain Name** field, enter the domain name. (**inband**)
- l) In the **VLAN** field, choose the **Create One** radio button.
- m) In the **VLAN Range** field, enter the VLAN range. Click **Save**, and click **Save** again. Click **Submit**.

Step 4 In the **Navigation** pane, right-click **Switch Policies** and choose **Configure Interface, PC and VPC**.

Step 5 In the **Configure Interface, PC, and VPC** dialog box, perform the following actions:

- a) Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the server.
- b) In the **Switches** field, from drop-down list, check the check boxes for the switches to which the servers are connected. (leaf1).
- c) In the **Switch Profile Name** field, enter a name for the profile (vmmConnectedLeaves).
- d) Click the + icon to configure the ports.

A dialog box similar to the following image is displayed for the user to enter the content:

- e) Verify that in the **Interface Type** area, the **Individual** radio button is selected.
- f) In the **Interfaces** field, enter the ports to which the servers are connected (1/40).
- g) In the **Interface Selector Name** field, enter the name of the port profile.
- h) In the **Interface Policy Group** field, click the **Create One** radio button.
- i) In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
- j) In the **Domain** field, from the drop-down list click the **Choose One** radio button
- k) From the **Physical Domain** drop-down list, choose the domain created earlier.
- l) In the **Domain Name** field, enter the domain name.
- m) Click **Save**, and click **Save** again.

Step 6 In the **Configure Interface, PC, and VPC** dialog box, click **Submit**.

Step 7 On the menu bar, click **TENANTS > mgmt**. In the **Navigation** pane, expand **Tenant mgmt > Networking > Bridge Domains** to configure the bridge domain on the in-band connection.

Step 8 Expand the in-band bridge domain (inb). Right-click **Subnets**. Click **Create Subnet** and perform the following actions to configure the in-band gateway:

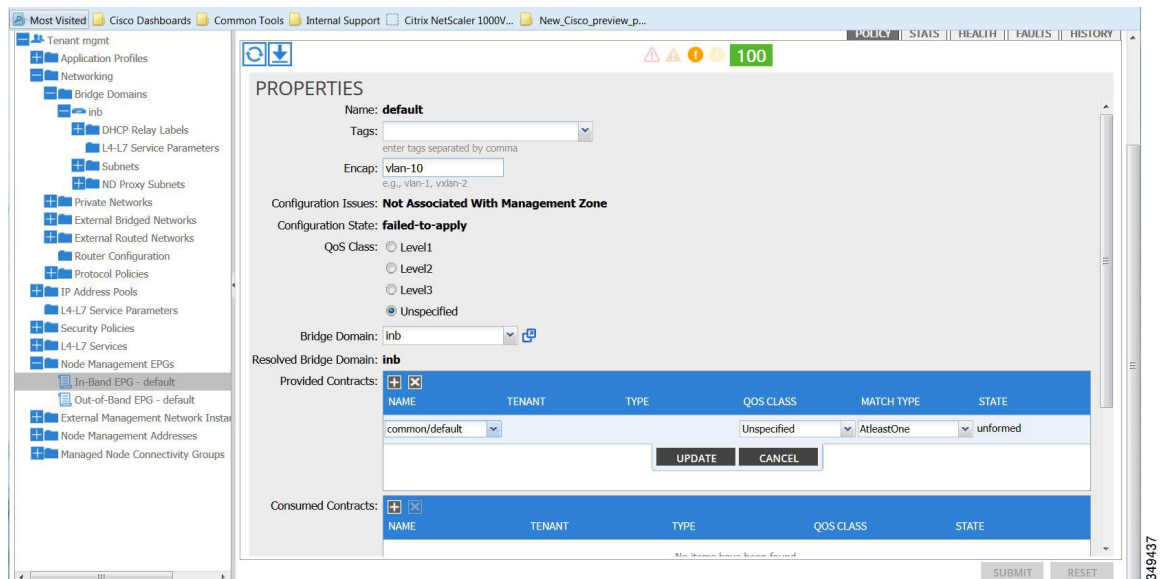
- a) In the **Create Subnet** dialog box, in the **Gateway IP** field, enter the in-band management gateway IP address and mask.

b) Click **Submit**.

Step 9 In the **Navigation** pane, expand **Tenant mgmt > Node Management EPGs**. Right-click **Node Management EPGs** and choose **Create In-Band Management EPG**. Perform the following actions to set the VLAN on the in-band EPG used to communicate with the APIC:

- In the **Name** field, enter the in-band management EPG name.
- In the **Encap** field, enter the VLAN (vlan-10).
- From the **Bridge Domain** drop-down field, choose the bridge domain. Click **Submit**.
- In the **Navigation** pane, choose the newly created in-band EPG.
- Expand **Provided Contracts**. In the **Name** field, from the drop-down list, choose the default contract to enable EPG to provide the default contract that will be consumed by the EPGs on which the VMM servers are located.
- Click **Update**, and click **Submit**.

A dialog box similar to the following image is displayed:



Step 10 In the **Navigation** pane, right-click **Node Management Addresses** and click **Create Node Management Addresses**, and perform the following actions to configure the IP addresses to be assigned to APIC controllers in the fabric:

- In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (apicInb).
- In the **Nodes** field, **Select** column, check the check boxes for the nodes that will be part of this fabric (apic1, apic2, apic3).
- In the **Config** field, check the **In-Band Addresses** check box.
- In the **Node Range** fields, enter the range.
- In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. This associates the default in-band Management EPG.
- In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.

g) Click **Submit**. The IP addresses for the APICs are now configured.

- Step 11** In the **Navigation** pane, right-click **Node Management Addresses**. Click **Create Node Management Addresses**, and perform the following actions to configure the IP addresses for the leaf and spine switches in the fabric:
- In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (switchInb).
 - In the **Nodes** field, **Select** column, check the check boxes next to the nodes that will be part of this fabric (leaf1, leaf2, spine1, spine2).
 - In the **Config** field, click the **In-Band Addresses** checkbox.
 - In the **Node Range** fields, enter the range.
 - In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. The default in-band management EPG is now associated.
 - In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired.
 - Click **Submit**. In the **Confirm** dialog box, click **Yes**. The IP addresses for the leaf and spine switches are now configured.
- Step 12** In the **Navigation** pane, under **Node Management Addresses**, click the APIC policy name (apicInb) to verify the configurations. In the **Work** pane, the IP addresses assigned to various nodes are displayed.
- Step 13** In the **Navigation** pane, under **Node Management Addresses**, click the switches policy name (switchInb). In the **Work** pane, the IP addresses that are assigned to switches and the gateway addresses they are using are displayed.
- Note** You can make out-of-band management access the default management connectivity mode for the APIC server by clicking **System > System Settings > APIC Connectivity Preferences**. Then on the **Connectivity Preferences** page, click **inband**.

Configuring Out-of-Band Management Access Using the Advanced GUI



Note IPv4 and IPv6 addresses are supported for out-of-band management access.

Before You Begin

The APIC out-of-band management connection link must be 1 Gbps.

Procedure

- Step 1** On the menu bar, choose **TENANTS > mgmt**. In the **Navigation** pane, expand **Tenant mgmt**.
- Step 2** Right-click **Node Management Addresses**, and click **Create Node Management Addresses**.
- Step 3** In the **Create Node Management Addresses** dialog box, perform the following actions:
- In the **Policy Name** field, enter a policy name (switchOob).
 - In the **Nodes** field, check the check boxes next to the appropriate leaf and spine switches (leaf1, leaf2, spine1).
 - In the **Config** field, check the check box for **Out of-Band Addresses**.
- Note** The **Out-of-Band IP addresses** area is displayed.

- d) In the **Out-of-Band Management EPG** field, choose the EPG from the drop-down list (default).
- e) In the **Out-of-Band IP Addresses** and **Out-of-Band Gateway** fields, enter the desired IPv4 or IPv6 addresses that will be assigned to the switches. Click **OK**.

The node management IP addresses are configured. You must configure out-of-band management access addresses for the leaf and spine switches as well as for APIC.

- Step 4** In the **Navigation** pane, expand **Node Management Addresses**, and click the policy that you created. In the **Work** pane, the out-of-band management addresses are displayed against the switches.
- Step 5** In the **Navigation** pane, expand **Security Policies > Out-of-Band Contracts**.
- Step 6** Right-click **Out-of-Band Contracts**, and click **Create Out-of-Band Contract**.
- Step 7** In the **Create Out-of-Band Contract** dialog box, perform the following tasks:
- a) In the **Name** field, enter a name for the contract (oob-default).
 - b) Expand **Subjects**. In the **Create Contract Subject** dialog box, in the **Name** field, enter a subject name (oob-default).
 - c) Expand **Filters**, and in the **Name** field, from the drop-down list, choose the name of the filter (default). Click **Update**, and click **OK**.
 - d) In the **Create Out-of-Band Contract** dialog box, click **Submit**.
- An out-of-band contract that can be applied to the out-of-band EPG is created.
- Step 8** In the **Navigation** pane, expand **Node Management EPGs > Out-of-Band EPG - default**.
- Step 9** In the **Work** pane, expand **Provided Out-of-Band Contracts**.
- Step 10** In the **OOB Contract** column, from the drop-down list, choose the out-of-band contract that you created (oob-default). Click **Update**, and click **Submit**.
The contract is associated with the node management EPG.
- Step 11** In the **Navigation** pane, right-click **External Network Instance Profile**, and click **Create External Management Entity Instance**.
- Step 12** In the **Create External Management Entity Instance** dialog box, perform the following actions:
- a) In the **Name** field, enter a name (oob-mgmt-ext).
 - b) Expand the **Consumed Out-of-Band Contracts** field. From the **Out-of-Band Contract** drop-down list, choose the contract that you created (oob-default). Click **Update**.
Choose the same contract that was provided by the out-of-band management.
 - c) In the **Subnets** field, enter the subnet address. Click **Submit**.
Only the subnet addresses you choose here will be used to manage the switches. The subnet addresses that are not included cannot be used to manage the switches.

The node management EPG is attached to the external network instance profile. The out-of-band management connectivity is configured.

Note You can make out-of-band management access the default management connectivity mode for the APIC server by clicking **System > System Settings > APIC Connectivity Preferences**. Then on the **Connectivity Preferences** page, click **ooband**.

Modifying the IP Address of an APIC Controller Using the GUI

Before You Begin

Procedure

-
- Step 1** On the menu bar, choose **TENANTS > mgmt**. In the **Navigation** pane, expand **Tenant mgmt**.
- Step 2** Right-click **Node Management Addresses**, and click **Create Node Management Addresses**.
- Step 3** In the **Create Node Management Addresses** dialog box, perform the following actions:
- In the **Policy Name** field, enter a policy name.
 - In the **Nodes** field, click the check box for the appropriate controller for which you want to change the IP address.
 - In the **Config** field, click the check box for **Out-of-Band Addresses**.
The **Out-of-Band Addresses** area expands.
 - In the **Out-of-Band Management EPG** field, from the drop-down list, choose the appropriate EPG.
 - In the **Out-of-Band Gateway** field, enter the gateway for the new IP address you want to assign.
The **Mask** field auto populates.
 - In the **Out-of-Band IP Addresses** range field, enter the appropriate IP address/addresses. Click **Submit**.
 - In the **Confirm** dialog box, click **Yes**, when asked to proceed with assigning new management IP addresses.
The new management IP address is assigned to the APIC controller.
-

What to Do Next

- You must use the new IP address to reconnect to the APIC controller.
- You must delete the old IP address of the controller once a new IP address is assigned to it.

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

All IPv6 tables mirror the existing IP tables functionality, except for Network Address Translation (NAT).

Existing IP Tables

- Earlier, every rule in the IPv6 tables were executed one at a time and a system call was made for every rule addition or deletion.
- Whenever a new policy was added, rules were appended to the existing IP tables file and no extra modifications were done to the file.
- When a new source port was configured in the out-of-band policy, it added source and destination rules with the same port number.

Modifications to IP Tables

- When IP tables are created, they are first written into hash maps that are then written into intermediate file IP tables-new which are restored. When saved, a new IP tables file is created in the `/etc/sysconfig/`

folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.

- 2 When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (/etc/sysconfig/iptables-new) and saved.
- 3 It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.
- 4 When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.


```
-A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
```
- 5 When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.

**Note**

- If only IPv4 is enabled, do not configure an IPv6 policy.
- If only IPv6 is enabled, do not configure an IPv4 policy.
- If both IPv4 and IPv6 are enabled and a policy is added, it will be configured to both the versions . So when you add an IPv4 subnet, it will be added to IP tables and similarly an IPv6 subnet is added to IPv6 tables.

Management Connectivity Modes

Establish connection to external entities using the out-of-band or in-band network depending upon whether you have configured out-of-band and/or in-band management connectivity. The following two modes are available to establish connectivity to external entities such as the vCenter server:


- Layer 2 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 2.
- Layer 3 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 3 through a router. The leaf is connected to a router through which external entities can be reached.

**Note**

- The inband IP address range must be separate and distinct from the IP address range used on the Layer 3 connection from the leaf node to outside the fabric.
- The Layer 3 inband management design does not provide inband management access to the spine fabric nodes in the topology.

The following diagram displays the two modes available to establish connectivity.

Figure 10: Layer 2 and Layer 3 Management Connectivity Examples



Configuring Layer 2 Management Connectivity Using the Advanced GUI



Note

Before You Begin

Before you create a vCenter domain profile, you must establish connectivity to establish an external network using in-band management network.

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

- Step 1** On the menu bar, choose **Tenants > mgmt**.
- Step 2** In the **Navigation** pane, expand **Tenant mgmt > Networking**, right-click **Bridge Domains**, and click **Create Bridge Domain**.
- Step 3** In the **Create Bridge Domain** dialog box, perform the following actions:
 - a) In the **Name** field, enter a bridge domain name.
 - b) In the **VRF** field, from the drop-down list, choose the network (mgmt/inb). Click **Next**.
 - c) Click the **L3 Configuration** tab, and in the **Subnets** field, click the + icon to add a subnet. Add the Gateway IP address as required.
 - d) In the **Create Bridge Domain** dialog box, click **Next** and then click **Submit**.

The bridge domain created.

- Step 4** In the **Navigation** pane, expand **Tenant mgmt > Application Profiles**.
- Step 5** Right-click **Application Profiles** and click **Create Application Profile**.
- Step 6** In the **Create Application Profile** dialog box, perform the following actions:
- In the **EPGs** field, click the + icon to add an EPG, and in the **Name** field, enter a name.
 - From the **BD** drop-down list, choose the appropriate BD.
 - From the **Domain** field drop-down list, choose the appropriate domain.
 - In the **Static Path** field, (enter the appropriate values similar to the following example, 101/1/40).
 - In the **Static Path VLAN** field, enter the appropriate VLAN (enter the appropriate value similar to the following example vlan-11).
 - In the **Consumed Contract** field, from the drop-down list, choose the appropriate value. Click **Update** and **Submit**.

In the **Navigation** pane, under **Networking** a bridge domain is created, and under **Application Profiles**, an application profile and an application EPG are created. The layer 2 management connectivity is now configured.

Configuring Layer 3 Management Connectivity Using the Advanced GUI



Note

- The name vmm is used as an example string in this task.

Before You Begin

Before you create a VMM domain profile, you must establish connectivity to an external network using the inband-management network.

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

- Step 1** On the menu bar, choose **TENANTS > mgmt**.
- Step 2** In the **Navigation** pane, perform the following actions:
- Expand **Tenant mgmt > Networking > External Routed Networks**.
 - Right-click **Create Routed Outside**.
- Step 3** In the **Create Routed Outside** dialog box, perform the following actions:
- In the **Name** field, enter the name of the Layer 3 routed outside policy (vmm).
This name can be up to 64 alphanumeric characters. You cannot change the name after the object is saved.
 - From the **VRF** drop-down list, choose the in-band default network (mgmt/inb).
Note You must choose the default in-band network.
- Step 4** Expand the **Nodes and Interfaces Protocol Profiles** area. In the **Create Node Profile** dialog box, perform the following actions:
- In the **Name** field, enter a name. (borderLeaf)

- b) Expand **Nodes** to display the **Select Node** dialog box. In the **Node ID** field, choose a leaf switch from the drop-down list (leaf1).
- c) In the **Router ID** field, enter the router ID.
- d) Expand **Static Routes**.
- e) In the **Create Static Route** dialog box, in the **Prefix** field, enter the subnet prefix for the static route of the external management system (for example, the VMware vCenter, the syslog server, or the AAA server) with which you are trying to communicate.
- f) Expand **Next Hop Addresses** and in the **Next Hop IP** field, enter the IP address of the router that is connected to the leaf switch. In the **Preference** field, choose a preference. Click **Update**.
- g) Click **OK**. In the **Select Node** dialog box, click **OK**.

Step 5 Expand **Interface Profiles**. In the **Create Interface Profile** dialog box, perform the following actions:

- a) In the **Name** field, enter a name. (portProfile1)
- b) Expand **Routed Interfaces**. In the **Select Routed Interface** area, in the **Path** field, from the drop-down list, choose the path that associates with leaf1.
- c) In the **IPv4 Primary/IPv6 Preferred Address** field, enter the IP address and subnet mask for the routed interface on the leaf. Click **OK**.
- d) In the **Create Interface Profile** dialog box, click **OK**. In the **Create Node Profile** dialog box, click **OK**.

Step 6 In the **Create Routed Outside** dialog box, click **Next**, and expand **External EPG Networks**.

Step 7 In the **Create External Network** dialog box, perform the following actions:

- a) In the **Name** field, enter a name (vmmMgmt).
- b) Expand the + icon for **Subnet**.
- c) In the **Create Subnet** dialog box, in the **IP address** field, enter the subnet address.
- d) Click **OK** two times, and click **Finish**.

The L3 management connectivity is configured.

Validating Management Connectivity

This validation process applies to both Layer 2 and Layer 3 modes and can be used to verify connectivity that is established by using the APIC GUI, REST API, or CLI.

After completing the steps to establish management connectivity, log in to the APIC console. Ping to the IP address of the vCenter server that is reachable (for example, 192.168.81.2) and verify that the ping works. This action indicates that the policies have been successfully applied.

Configuring a VMM Domain

Configuring Virtual Machine Networking Policies

The APIC integrates with third-party VM manager (VMM) (for example, VMware vCenter and SCVMM) to extend the benefits of ACI to the virtualized infrastructure. The APIC enables the ACI policies inside the VMM system to be used by its administrator.

This section provides examples of VMM integration using VMware vCenter and vShield. For details about the different modes of Cisco ACI and VMM integration, see the *ACI Virtualization Guide*.

About the VM Manager



Note

Information about the necessary configuration of the APIC for integration with the vCenter is described here. For instructions about configuring the VMware components, see the VMware documentation.

The following are details of some VM manager terms:

- A VM controller is an external virtual machine management entity such as VMware vCenter, and the VMware vShield. The APIC communicates with the controller to publish network policies that are applied to virtual workloads. A VM controller administrator provides an APIC administrator with a VM controller authentication credential; multiple controllers of the same type can use the same credential.
- Credentials represent the authentication credentials to communicate with VM controllers. Multiple controllers can use the same credentials.
- A virtual machine mobility domain (vCenter mobility domain) is a grouping of VM controllers with similar networking policy requirements. This mandatory container holds one or more VM controllers with policies such as for a VLAN pool, server to network MTU policy, or server to network access LACP policy. When an endpoint group gets associated with a vCenter domain, network policies get pushed to all the VM controllers in the vCenter domain.
- A pool represents a range of traffic encapsulation identifiers (for example, VLAN IDs, VNIDs, and multicast addresses). A pool is a shared resource and can be consumed by multiple domains such as VMM and Layer 4 to Layer 7 services. A leaf switch does not support overlapping VLAN pools. You must not associate different overlapping VLAN pools with the VMM domain. The two types of VLAN-based pools are as follows:
 - Dynamic pools—Managed internally by the APIC to allocate VLANs for endpoint groups (EPGs). A vCenter Domain can associate only to a dynamic pool.
 - Static pools—The EPG has a relation to the domain, and the domain has a relation to the pool. The pool contains a range of encapsulated VLANs and VXLANs. For static EPG deployment, the user defines the interface and the encapsulation. The encapsulation must be within the range of a pool that is associated with a domain with which the EPG is associated.
- For a VMware vCenter to be deployed, it must operate in VLAN mode or VXLAN mode. A VMM domain must be associated with a VLAN pool and a vShield must be associated with the vCenter.

About Attachable Entity Profile

Attach Entity Profiles

The ACI fabric provides multiple **attachment points** that connect through leaf ports to various **external entities** such as baremetal servers, hypervisors, Layer 2 switches (for example, the Cisco UCS fabric interconnect), and Layer 3 routers (for example Cisco Nexus 7000 Series switches). These attachment points can be physical ports, port channels, or a virtual port channel (vPC) on the leaf switches.

An **attachable entity profile** (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies, for example, Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), maximum transmission unit (MTU), and Link Aggregation Control Protocol (LACP).

A VM manager (VMM) domain automatically derives the physical interfaces policies from the interface policy groups that are associated with an AEP.

- An override policy at AEP can be used to specify a different physical interface policy for a VMM domain. This policy is useful in scenarios where a hypervisor is connected to the leaf switch through an intermediate Layer 2 node, and a different policy is desired at the leaf switch and hypervisor physical ports. For example, you can configure LACP between a leaf switch and a Layer 2 node. At the same time, you can disable LACP between the hypervisor and the Layer 2 switch by disabling LACP under the AEP override policy.

An AEP is required to deploy any VLAN pools on the leaf switches. It is possible to reuse the encapsulation pools (for example, VLAN) across different leaf switches. An AEP implicitly provides the scope of the VLAN pool (associated to the domain) to the physical infrastructure.



Note

- An AEP provisions the VLAN pool (and associated VLANs) on the leaf. The VLANs are not actually enabled on the port. No traffic flows unless an EPG is deployed on the port.
- Without VLAN pool deployment using an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
 - A particular VLAN is provisioned or enabled on the leaf port based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter.
 - If you wish to set the VMM encapsulation statically in the EPG, you must use a static pool. If you have a mix of static and dynamic allocations, create a dynamic pool and add a block within that pool with static mode.
- A leaf switch does not support overlapping VLAN pools. Different overlapping VLAN pools must not be associated with the same AEP that is associated through a domain.

For information about configuring LLDP and CDP, see the chapter related to Working with Blade Servers in the guide.

Prerequisites for Creating a VMM Domain Profile

To configure a VMM domain profile, you must meet the following prerequisites:

- All fabric nodes are discovered and configured.
- Inband (inb) or out-of-band (oob) management has been configured on the APIC.
- A Virtual Machine Manager (VMM) is installed, configured, and reachable through the inb/oob management network (for example, a vCenter).
- You have the administrator/root credentials to the VMM (for example vCenter).



Note If you prefer not to use the vCenter admin/root credentials, you can create a custom user account with minimum required permissions. See [Custom User Account with Minimum VMware vCenter Privileges](#), on page 47 for a list of the required user privileges.

- A DNS policy for the APIC must be configured if you plan to reference the VMM by hostname rather than an IP address.

Custom User Account with Minimum VMware vCenter Privileges

This allows the APIC to send VMware API commands to vCenter to allow the creation of the DVS/AVS, creation of the VMK interface (AVS), publish port groups and relay all necessary alerts.

To configure the vCenter from Cisco APIC, your credentials must allow the following minimum set of privileges within the vCenter:

- **Alarms**

APIC creates two alarms on the folder. One for DVS and another for port-group. The alarm is raised when the EPG or Domain policy is deleted on APIC, but for port-group or DVS it cannot be deleted due to the VMs are attached.

- **Distributed Switch**

- **dvPort Group**

- **Folder**

- **Network**

APIC manages the network settings such as add or delete port-groups, setting host/DVS MTU, LLDP/CDP, LACP etc.

- **Host**

If you use AVS in addition to above, you need the Host privilege on the data center where APIC will create DVS.

- **Host.Configuration.Advanced settings**
- **Host.Local operations.Reconfigure virtual machine**
- **Host.Configuration.Network configuration**

This is needed for AVS and the auto-placement feature for virtual Layer 4 to Layer 7 Service VMs. For AVS, APIC creates VMK interface and places it in 'vtep' port-group which is used for OpFlex.

- **Virtual machine**

If you use Service Graph in addition to above, you need the Virtual machine privilege for the virtual appliances which will be used for Service Graph.

- **Virtual machine.Configuration.Modify device settings**
- **Virtual machine.Configuration.Settings**

Creating a VMM Domain Profile

In this section, examples of a VMM domain are vCenter domain.

Creating a vCenter Domain Profile Using the GUI

An overview of the tasks performed in the creation of a vCenter Domain are as follows (details are in the steps that follow):

- Create/select a switch profile
- Create/select an interface profile
- Create/select an interface policy group
- Create/select VLAN pool
- Create vCenter domain
- Create vCenter credentials

Procedure

-
- Step 1** On the menu bar, click **FABRIC > Access Policies**.
- Step 2** In the **Navigation** pane, right-click **Switch Policies**, and then click **Configured Interfaces, PC, and VPC**.
- Step 3** In the **Configured Interfaces, PC, and VPC** dialog box, perform the following actions:
- a) Expand **Configured Switch Interfaces**.
 - b) Click the + icon.
 - c) Make sure that the **Quick** radio button is chosen.
 - d) From the **Switches** drop-down list, choose the appropriate leaf ID.
In the **Switch Profile Name** field, the switch profile name automatically populates.
 - e) Click the + icon to configure the switch interfaces.
 - f) In the **Interface Type** area, check the appropriate radio button.
 - g) In the **Interfaces** field, enter the desired interface range.
 - h) In the **Interface Selector Name** field, the selector name automatically populates.
 - i) In the **Interface Policy Group** area, choose the **Create One** radio button.
 - j) From the **Link Level Policy** drop-down list, choose the desired link level policy.
 - k) From the **CDP Policy** drop-down list, choose the desired CDP policy.
Note Similarly choose the desired interface policies from the available policy areas.
 - l) In the **Attached Device Type** area, choose **ESX Hosts**.
 - m) In the **Domain** area, make sure that the **Create One** radio button is chosen.
 - n) In the **Domain Name** field, enter the domain name.
 - o) In the **VLAN** area, make sure that the **Create One** radio button is chosen.
 - p) In the **VLAN Range** field, enter the VLAN range as appropriate.
Note We recommend a range of at least 200 VLAN numbers. Do not define a range that includes the reserved VLAN ID for infrastructure network, because that VLAN is for internal use.
 - q) In the **vCenter Login Name** field, enter the login name.

- r) (Optional) From the **Security Domains** drop-down list, choose the appropriate security domain.
- s) In the **Password** field, enter a password.
- t) In the **Confirm Password** field, reenter the password.
- u) Expand **vCenter**.

Step 4 In the **Create vCenter Controller** dialog box, enter the appropriate information, and click **OK**.

Step 5 In the **Configure Interface, PC, And VPC** dialog box, complete the following actions:
If you do not specify policies in the **Port Channel Mode** and the **vSwitch Policy** areas, the same policies that you configured earlier in this procedure will take effect for the vSwitch.

- a) From the **Port Channel Mode** drop-down list, choose a mode.
- b) In the **vSwitch Policy** area, click the desired radio button to enable CDP or LLDP.
- c) From the **NetFlow Exporter Policy** drop-down list, choose a policy or create one.
A NetFlow exporter policy configures the external collector reachability.
- d) Choose values from the **Active Flow Timeout**, **Idle Flow Timeout**, and **Sampling Rate** drop-down lists.
- e) Click **SAVE** twice and then click **SUBMIT**.

Step 6 Verify the new domain and profiles, by performing the following actions:

- a) On the menu bar, choose **VM Networking > Inventory**.
- b) In the **Navigation** pane, expand **VMware > Domain_name > vCenter_name**.

In the **Work** pane, under **Properties**, view the VMM domain name to verify that the controller is online. In the **Work** pane, the vCenter properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the vCenter server is established, and the inventory is available.

Creating a vCenter and a vShield Domain Profile Using the Advanced GUI

An overview of the tasks performed in the creation of a vCenter and vShield domains are as follows (details are in the steps that follow):

- Create/select a switch profile.
- Create/select an interface profile.
- Create/select an interface policy group.
- Create/select VLAN pool.
- Create vCenter and vShield domains.
- Create vCenter and vShield credentials.

Procedure

Step 1 On the menu bar, click **FABRIC > Access Policies**.

Step 2 In the **Navigation** pane, click **Switch Policies**, and then click **Configure Interfaces, PC, and VPC**.

Step 3 In the **Configure Interface, PC, and VPC** dialog box, perform the following actions:

- a) Expand **Configured Switch Interfaces**.

- b) Click the + icon.
- c) In the **Select Switches to Configure Interfaces** area, make sure that the **Quick** radio button is chosen.
- d) From the **Switches** drop-down list, choose the appropriate leaf IDs.
In the **Switch Profile Name** field, the switch profile name automatically populates.
- e) Click the + icon to configure the switch interfaces.
- f) In the **Interface Type** area, click the appropriate radio button.
- g) In the **Interfaces** field, enter the desired interface range.
- h) In the **Interface Selector Name** field, the selector name automatically populates.
- i) In the **Interface Policy Group** area, make sure that the **Create One** radio button is chosen.
- j) From the **Link Level Policy** drop-down list, choose the desired link level policy.
- k) From the **CDP Policy** drop-down list, choose the desired CDP policy.
Note Similarly choose the desired interface policies in the available policy areas.
- l) From the **Attached Device Type** drop-down list, choose the appropriate device type.
- m) In the **Domain** area, make sure that the **Create One** radio button is chosen.
- n) In the **Domain Name** field, enter the domain name.
- o) In the **VLAN** area, make sure that the **Create One** radio button is chosen.
- p) In the **VLAN Range** field, enter the VLAN range as appropriate.
Note We recommend a range of at least 200 VLAN numbers. Do not define a range that includes the reserved VLAN ID for infrastructure network, because that VLAN is for internal use.
- q) In the **vCenter Login Name** field, enter the login name.
- r) In the **Password** field, enter a password.
- s) In the **Confirm Password** field, reenter the password.
- t) Expand **vCenter/vShield**.

Step 4 In the **Create vCenter/vShield Controller** dialog box, enter the appropriate information and click **OK**.

Step 5 In the **Configure Interface, PC, And VPC** dialog box, complete the following actions:
If you do not specify policies in the **Port Channel Mode** and the **vSwitch Policy** areas, the same policies that you configured earlier in this procedure will take effect for the vSwitch.

- a) From the **Port Channel Mode** drop-down list, choose a mode.
- b) In the **vSwitch Policy** area, click the desired radio button to enable CDP or LLDP.
- c) From the **NetFlow Exporter Policy** drop-down list, choose a policy or create one.
A NetFlow exporter policy configures the external collector reachability.
- d) Choose values from the **Active Flow Timeout**, **Idle Flow Timeout**, and **Sampling Rate** drop-down lists.
- e) Click **SAVE** twice and then click **SUBMIT**.

Step 6 Verify the new domain and profiles, by performing the following actions:

- a) On the menu bar, choose **VM Networking > Inventory**.
- b) In the **Navigation** pane, expand , and click **VMware > Domain_name > vCenter_name**.

In the **Work** pane, under **Properties**, view the VMM domain name to verify that the controller is online. In the **Work** pane, the vCenter properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the vCenter server is established, and the inventory is available.

Creating Tenants, VRF, and Bridge Domains

Tenants Overview

- 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.

Tenant Creation

A tenant contains primary elements such as filters, contracts, bridge domains, and application profiles that you can create after you first create a tenant.

VRF and Bridge Domains

You can create and specify a VRF and a bridge domain for the tenant. The defined bridge domain element subnets reference a corresponding Layer 3 context.

For details about enabling IPv6 Neighbor Discovery see *IPv6 and Neighbor Discovery* in *Cisco APIC Layer 3 Networking Guide*.

Creating a Tenant, VRF, and Bridge Domain Using the GUI

If you have a public subnet when you configure the routed outside, you must associate the bridge domain with the outside configuration.

Procedure

-
- Step 1** On the menu bar, click **TENANT > Add Tenant**.
 - Step 2** In the **Create Tenant** dialog box, perform the following tasks:
 - a) In the **Name** field, enter a name.
 - b) Click the **Security Domains +** icon to open the **Create Security Domain** dialog box.
 - c) In the **Name** field, enter a name for the security domain. Click **Submit**.
 - d) In the **Create Tenant** dialog box, check the check box for the security domain that you created, and click **Submit**.
 - Step 3** In the **Navigation** pane, expand **Tenant-name > Networking**, and in the **Work** pane, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:

- a) In the **Name** field, enter a name.
- b) Click **Submit** to complete the VRF configuration.

Step 4 In the **Networking** pane, drag the **BD** icon to the canvas while connecting it to the **VRF** icon. In the **Create Bridge Domain** dialog box that displays, perform the following tasks:

- a) In the **Name** field, enter a name.
- b) Click the **L3 Configurations** tab.
- c) Expand **Subnets** to open the **Create Subnet** dialog box, enter the subnet mask in the **Gateway IP** field and click **OK**.
- d) Click **Submit** to complete bridge domain configuration.

Step 5 In the **Networks** pane, drag the **L3** icon down to the canvas while connecting it to the **VRF** icon. In the **Create Routed Outside** dialog box that displays, perform the following tasks:

- a) In the **Name** field, enter a name.
- b) Expand **Nodes And Interfaces Protocol Profiles** to open the **Create Node Profile** dialog box.
- c) In the **Name** field, enter a name.
- d) Expand **Nodes** to open the **Select Node** dialog box.
- e) In the **Node ID** field, choose a node from the drop-down list.
- f) In the **Router ID** field, enter the router ID.
- g) Expand **Static Routes** to open the **Create Static Route** dialog box.
- h) In the **Prefix** field, enter the IPv4 or IPv6 address.
- i) Expand **Next Hop Addresses** and in the **Next Hop IP** field, enter the IPv4 or IPv6 address.
- j) In the **Preference** field, enter a number, then click **UPDATE** and then **OK**.
- k) In the **Select Node** dialog box, click **OK**.
- l) In the **Create Node Profile** dialog box, click **OK**.
- m) Check the **BGP**, **OSPF**, or **EIGRP** check boxes if desired, and click **NEXT**. Click **OK** to complete the Layer 3 configuration.

To confirm L3 configuration, in the **Navigation** pane, expand **Networking > VRFs**.

Configuring an Enforced Bridge Domain Using the Basic GUI

Procedure

Step 1 Log in to the APIC GUI, and on the menu bar, click **TENANT > Add Tenant**.

Step 2 In the **Create Tenant** dialog box, perform the following tasks:

- a) In the **Name** field, enter a tenant name.
- b) Click the **Security Domains +** icon to open the **Create Security Domain** dialog box.
- c) In the **Name** field, enter a security domain name and click **Submit**.
- d) In the **Create Tenant** dialog box, check the check box for the security domain that you created, and click **Submit**.

Step 3 In the **Navigation** pane, expand **Tenant-name > Networking**, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:

- a) In the **Name** field, enter the VRF name.

- b) Select the **BD Enforcement Status** check box.
- c) Click **Submit** to complete the VRF configuration.

To confirm enforced bridge domain configuration, expand **Fabric > Fabric Policies > Global policies > Exception List** and confirm the presence of the BD Enforced Exception list.

Configuring an Enforced Bridge Domain Using the NX-OS Style CLI

This section provides information on how to configure your enforced bridge domain using the NX-OS style command line interface (CLI).

Procedure

Step 1 Create and enable the tenant:

Example:

In the following example ("cokeVrf") is created and enabled.

```
apic1(config-tenant)# vrf context cokeVrf
apic1(config-tenant-vrf)# bd-enforce enable
apic1(config-tenant-vrf)# exit
apic1(config-tenant)#exit
```

Step 2 Add the subnet to the exception list.

Example:

```
apic1(config)#bd-enf-exp-ip add1.2.3.4/24
apic1(config)#exit
```

You can confirm if the enforced bridge domain is operational using the following type of command:

```
apic1# show running-config all | grep bd-enf
bd-enforce enable
bd-enf-exp-ip add 1.2.3.4/24
```

The following command removes the subnet from the exception list:

```
apic1(config)# no bd-enf-exp-ip 1.2.3.4/24
apic1(config)#tenant coke
apic1(config-tenant)#vrf context cokeVrf
```

What to Do Next

To disable the enforced bridge domain run the following command:

```
apic1(config-tenant-vrf)# no bd-enforce enable
```

Configuring an Enforced Bridge Domain Using the REST API

Procedure

	Command or Action	Purpose
Step 1	Create a tenant. Example: POST https://apic-ip-address/api/mo/uni.xml <code><fvTenant name="ExampleCorp"/></code>	When the POST succeeds, you see the object that you created in the output.
Step 2	Create a VRF and bridge domain. Example: URL for POST: https://apic-ip-address/api/mo/uni/tn-ExampleCorp.xml <code><fvTenant name="ExampleCorp"></code> <code> <fvCtx name="pvn1"/></code> <code> <fvBD name="bd1"></code> <code> <fvRsCtx tnFvCtxName="pvn1"></code> <code> bdEnforceEnable="yes"/></code> <code> <fvSubnet ip="10.10.100.1/24"/></code> <code> </fvBD></code> <code></fvTenant></code> For adding an exception IP, use the following post: https://apic-ip-address/api/node/mo/uni/infra.xml <code><bdEnforceExceptionCont></code> <code><bdEnforceExceptIp ip="11.0.1.0/24"/></code> <code></bdEnforceExceptionCont></code> Note If you have a public subnet when you configure the routed outside, you must associate the bridge domain with the outside configuration.	Note The Gateway Address can be an IPv4 or an IPv6 address. For more about details IPv6 gateway address, see the related KB article, <i>KB: Creating a Tenant, VRF, and Bridge Domain with IPv6 Neighbor Discovery</i> .

Configuring Server or Service Policies

Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets. If the client is on an ESX hypervisor with a deployed vShield Domain profile, then the use of a DHCP relay policy configuration is mandatory.

When a vShield controller deploys a Virtual Extensible Local Area Network (VXLAN), the hypervisor hosts create a kernel (vmkN, virtual tunnel end-point [VTEP]) interface. These interfaces need an IP address in the infrastructure tenant that uses DHCP. Therefore, you must configure a DHCP relay policy so that the APIC can act as the DHCP server and provide these IP addresses.

When an ACI fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from

a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the ACI fabric must support Option 82.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the GUI

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Deploying DHCP Relay Policy for an Endpoint Group

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

-
- Step 1** On the menu bar, choose **TENANTS > infra**. In the **Navigation** pane, under **Tenant infra**, expand **Networking > Protocol Policies > DHCP > Relay Policies**.
- Step 2** Right-click **Relay Policies** and click **Create DHCP Relay Policy**.
- Step 3** In the **Create DHCP Relay Policy** dialog box, perform the following actions:
- In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).
 - Expand **Providers**. In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.
 - In the **Application EPG** area, in the **Tenant** field, from the drop-down list, choose the tenant. (infra)
 - In the **Application Profile** field, from the drop-down list, choose the application. (access)
 - In the **EPG** field, from the drop-down list, choose the EPG. (default)
 - In the **DHCP Server Address** field, enter the IP address for the infra DHCP server. Click **Update**.

Note The infra DHCP IP address is the infra IP address of APIC1. You must enter the default IP address of 10.0.0.1 if deploying for vShield controller configuration.
 - Click **Submit**.
- The DHCP relay policy is created.
- Step 4** In the **Navigation** pane, expand **Networking > Bridge Domains > default > DHCP Relay Labels**.
- Step 5** Right-click **DHCP Relay Labels**, and click **Create DHCP Relay Label**.
- Step 6** In the **Create DHCP Relay Label** dialog box, perform the following actions:
- In the **Scope** field, click the tenant radio button.

This action displays, in the **Name** field drop-down list, the DHCP relay policy created earlier.
 - In the **Name** field, from the drop-down list, choose the name of the DHCP policy created (DhcpRelayP) or create a new relay policy by choosing **Create DHCP Relay Policy**.
 - In the **DHCP Option Policy**, select an existing option policy, or create a new one by choosing **Create DHCP Option Policy**.
 - Click **Submit**.

The DHCP server is associated with the bridge domain.

- Step 7** In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels** to view the DHCP server created.

Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.
- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.
- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking** > **VRF** policy configuration in the tenants configuration.

Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere

Source	In-Band Management	Out-of-Band Management	External Server Location
Spine switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be attached to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/fabric/dnsp-default.xml -->
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The `gai.conf` settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a `gai.conf` to prioritize IPv4 over IPv6:

```
# Generated by APIC
label ::1/128      0
label ::/0        1
label 2002::/16   2
label ::/96       3
label ::ffff:0:0/96 4
precedence ::1/128      50
precedence ::/0        40
precedence 2002::/16   30
precedence ::/96       20
# For APICs preferring IPv4 connections, change the value to 100.
precedence ::ffff:0:0/96 10
```

Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPv4) or AAAA records (IPv6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to `/etc/resolv.conf`.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in `/etc/resolv.conf`. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in `/etc/resolv.conf`). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add “options single-request-reopen” to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of `resolv.conf` in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (glibc) will use the IPv6 network by default because `getaddrinfo()` will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (`::ffff/96`). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the glibc IPv6 selection preference for `getaddrinfo()` in `/etc/gai.conf`.

In order to allow glibc to return multiple addresses when using `/etc/hosts`, “multi on” should be added to the `/etc/hosts` file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

Configuring a DNS Service Policy to Connect with DNS Providers Using the Advanced GUI

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**. In the **Navigation** pane, expand **Global Policies > DNS Profiles**, and click the default DNS profile.
- Step 2** In the **Work** pane, in the **Management EPG** field, from the drop-down list, choose the appropriate management EPG (default (Out-of-Band)).
- Step 3** Expand **DNS Providers**, and perform the following actions:
- In the **Address** field, enter the provider address.
 - In the **Preferred** column, check the check box if you want to have this address as the preferred provider. You can have only one preferred provider.
 - Click **Update**.
 - (Optional) To add a secondary DNS provider, expand **DNS Providers**, and in the **Address** field, type the provider address. Click **Update**.
- Step 4** Expand **DNS Domains**, and perform the following actions:
- In the **Name** field, enter the domain name (cisco.com).
 - In the **Default** column, check the check box to make this domain the default domain. You can have only one domain name as the default.
 - Click **Update**.
 - (Optional) To add a secondary DNS domain, expand **DNS Domains**. In the **Address** field, enter the secondary domain name. Click **Update**.
- Step 5** Click **Submit**.
The DNS server is configured.
- Step 6** On the menu bar, click **TENANTS > mgmt**.
- Step 7** In the **Navigation** pane, expand **Networking > VRF > oob**, and click **oob**.
- Step 8** In the **Work** pane, under **Properties**, in the **DNS labels** field, enter the appropriate DNS label (default). Click **Submit**.
The DNS profile label is now configured on the tenant and VRF.
-

Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI

Procedure

- Step 1** Verify the configuration for the default DNS profile.

Example:

```

admin@apic1:~> cd /aci/fabric/fabric-policies/global-policies/dns-profiles/default
admin@apic1:default> cat summary
# dns-profile
name : default
description : added via CLI by tdeleon@cisco.com
ownerkey :
ownertag :

dns-providers:
address preferred
-----
10.44.124.122 no
10.70.168.183 no
10.37.87.157 no
10.102.6.247 yes
dns-domains:
name default description
-----
cisco.com yes
management-epg : tenants/mgmt/node-management-epgs/default/out-of-band/default

```

Step 2 Verify the configurations for the DNS labels.

Example:

```

admin@apic1:default> cd
/aci/tenants/mgmt/networking/private-networks/oob/dns-profile-labels/default
admin@apic1:default> cat summary
# dns-lbl
name : default
description :
ownerkey :
ownertag :
tag : yellow-green

```

Step 3 Verify that the applied configuration is operating on the fabric controllers.

Example:

```

admin@apic1:~> cat /etc/resolv.conf
# Generated by IFC
search cisco.com
nameserver 10.102.6.247
nameserver 10.44.124.122
nameserver 10.37.87.157
nameserver 10.70.168.183
admin@apic1:~> ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=238 time=35.4 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=238 time=29.0 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=238 time=29.2 ms

```

Step 4 Verify that the applied configuration is operating on the leaf and spine switches.

Example:

```

leaf1# cat /etc/resolv.conf
search cisco.com
nameserver 10.102.6.247
nameserver 10.70.168.183
nameserver 10.44.124.122
nameserver 10.37.87.157
leaf1# cat /etc/dcos_resolv.conf
# DNS enabled
leaf1# ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161): 56 data bytes
64 bytes from 72.163.4.161: icmp_seq=0 ttl=238 time=29.255 ms

```

```
64 bytes from 72.163.4.161: icmp_seq=1 ttl=238 time=29.212 ms
64 bytes from 72.163.4.161: icmp_seq=2 ttl=238 time=29.343 ms
```

Configuring External Connectivity for Tenants

Before you can distribute the static route to the other leaf switches on the Application Centric Infrastructure (ACI) fabric, a multiprotocol BGP (MP-BGP) process must first be operating, and the spine switches must be configured as BGP route reflectors.

To integrate the ACI fabric into an external routed network, you can configure Open Shortest Path First (OSPF) for management tenant Layer 3 connectivity.

Configuring an MP-BGP Route Reflector Using the Advanced GUI

Procedure

- Step 1** On the menu bar, choose **System > System Settings**.
 - Step 2** In the **Navigation** pane, right-click **BGP Route Reflector**, and click **Create Route Reflector Node Policy EP**.
 - Step 3** In the **Create Route Reflector Node Policy EP** dialog box, from the **Spine Node** drop-down list, choose the appropriate spine node. Click **Submit**.
Note Repeat the above steps to add additional spine nodes as required.
The spine switch is marked as the route reflector node.
 - Step 4** In the **BGP Route Reflector** properties area, in the **Autonomous System Number** field, choose the appropriate number. Click **Submit**.
Note The autonomous system number must match the leaf connected router configuration if Border Gateway Protocol (BGP) is configured on the router. If you are using routes learned using static or Open Shortest Path First (OSPF), the autonomous system number value can be any valid value.
 - Step 5** On the menu bar, choose **Fabric > Fabric Policies > POD Policies**.
 - Step 6** In the **Navigation** pane, expand and right-click **Policy Groups**, and click **Create POD Policy Group**.
 - Step 7** In the **Create POD Policy Group** dialog box, in the **Name** field, enter the name of a pod policy group.
 - Step 8** In the **BGP Route Reflector Policy** drop-down list, choose the appropriate policy (default). Click **Submit**.
The BGP route reflector policy is associated with the route reflector pod policy group, and the BGP process is enabled on the leaf switches.
 - Step 9** In the **Navigation** pane, choose **Pod Policies > Profiles > default**. In the **Work** pane, from the **Fabric Policy Group** drop-down list, choose the pod policy that was created earlier. Click **Submit**.
The pod policy group is now applied to the fabric policy group.
-

Verifying the MP-BGP Route Reflector Configuration

Procedure

- Step 1** Verify the configuration by performing the following actions:
- Use secure shell (SSH) to log in as an administrator to each leaf switch as required.
 - Enter the **show processes | grep bgp** command to verify the state is S.
If the state is NR (not running), the configuration was not successful.
- Step 2** Verify that the autonomous system number is configured in the spine switches by performing the following actions:
- Use the SSH to log in as an administrator to each spine switch as required.
 - Execute the following commands from the shell window

Example:
`cd /mit/sys/bgp/inst`

Example:
`grep asn summary`

The configured autonomous system number must be displayed. If the autonomous system number value displays as 0, the configuration was not successful.

Creating an OSPF External Routed Network for Management Tenant Using the Advanced GUI

- You must verify that the router ID and the logical interface profile IP address are different and do not overlap.
- The following steps are for creating an OSPF external routed network for a management tenant. To create an OSPF external routed network for a tenant, you must choose a tenant and create a VRF for the tenant.
- For more details, see *Cisco APIC and Transit Routing*.

Procedure

- Step 1** On the menu bar, choose **TENANTS > mgmt**.
- Step 2** In the **Navigation** pane, expand **Networking > External Routed Networks**.
- Step 3** Right-click **External Routed Networks**, and click **Create Routed Outside**.
- Step 4** In the **Create Routed Outside** dialog box, perform the following actions:
- In the **Name** field, enter a name (RtdOut).

- b) Check the **OSPF** check box.
- c) In the **OSPF Area ID** field, enter an area ID.
- d) In the **OSPF Area Control** field, check the appropriate check box.
- e) In the **OSPF Area Type** field, choose the appropriate area type.
- f) In the **OSPF Area Cost** field, choose the appropriate value.
- g) In the **VRF** field, from the drop-down list, choose the VRF (inb).
Note This step associates the routed outside with the in-band VRF.
- h) From the **External Routed Domain** drop-down list, choose the appropriate domain.
- i) Click the + icon for **Nodes and Interfaces Protocol Profiles** area.

Step 5 In the **Create Node Profile** dialog box, perform the following actions:

- a) In the **Name** field, enter a name for the node profile. (borderLeaf).
- b) In the **Nodes** field, click the + icon to display the **Select Node** dialog box.
- c) In the **Node ID** field, from the drop-down list, choose the first node. (leaf1).
- d) In the **Router ID** field, enter a unique router ID.
- e) Uncheck the **Use Router ID as Loopback Address** field.
Note By default, the router ID is used as a loopback address. If you want them to be different, uncheck the **Use Router ID as Loopback Address** check box.
- f) Expand **Loopback Addresses**, and enter the IP address in the **IP** field. Click **Update**, and click **OK**. Enter the desired IPv4 or IPv6 IP address.
- g) In the **Nodes** field, expand the + icon to display the **Select Node** dialog box.
Note You are adding a second node ID.
- h) In the **Node ID** field, from the drop-down list, choose the next node. (leaf2).
- i) In the **Router ID** field, enter a unique router ID.
- j) Uncheck the **Use Router ID as Loopback Address** field.
Note By default, the router ID is used as a loopback address. If you want them to be different, uncheck the **Use Router ID as Loopback Address** check box.
- k) Expand **Loopback Addresses**, and enter the IP address in the **IP** field. Click **Update**, and click **OK**. Click **OK**. Enter the desired IPv4 or IPv6 IP address.

Step 6 In the **Create Node Profile** dialog box, in the **OSPF Interface Profiles** area, click the + icon.

Step 7 In the **Create Interface Profile** dialog box, perform the following tasks:

- a) In the **Name** field, enter the name of the profile (portProf).
- b) In the **Interfaces** area, click the **Routed Interfaces** tab, and click the + icon.
- c) In the **Select Routed Interfaces** dialog box, in the **Path** field, from the drop-down list, choose the first port (leaf1, port 1/40).
- d) In the **IP Address** field, enter an IP address and mask. Click **OK**.
- e) In the **Interfaces** area, click the **Routed Interfaces** tab, and click the + icon.
- f) In the **Select Routed Interfaces** dialog box, in the **Path** field, from the drop-down list, choose the second port (leaf2, port 1/40).
- g) In the **IP Address** field, enter an IP address and mask. Click **OK**.
Note This IP address should be different from the IP address you entered for leaf1 earlier.
- h) In the **Create Interface Profile** dialog box, click **OK**.

The interfaces are configured along with the OSPF interface.

Step 8 In the **Create Node Profile** dialog box, click **OK**.

Step 9 In the **Create Routed Outside** dialog box, click **Next**.
The **Step 2 External EPG Networks** area is displayed.

Step 10 In the **External EPG Networks** area, click the + icon.

Step 11 In the **Create External Network** dialog box, perform the following actions:

- a) In the **Name** field, enter a name for the external network (extMgmt).
- b) Expand **Subnet** and in the **Create Subnet** dialog box, in the **IP address** field, enter an IP address and mask for the subnet.
- c) In the **Scope** field, check the desired check boxes. Click **OK**.
- d) In the **Create External Network** dialog box, click **OK**.
- e) In the **Create Routed Outside** dialog box, click **Finish**.

Note In the **Work** pane, in the **External Routed Networks** area, the external routed network icon (RtdOut) is now displayed.

Deploying an Application Policy

Three-Tier Application Deployment

A filter specifies the data protocols to be allowed or denied by a contract that contains the filter. A contract can contain multiple subjects. A subject can be used to realize uni- or bidirectional filters. A unidirectional filter is a filter that is used in one direction, either from consumer-to-provider (IN) or from provider-to-consumer (OUT) filter. A bidirectional filter is the same filter that is used in both directions. It is not reflexive.

Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

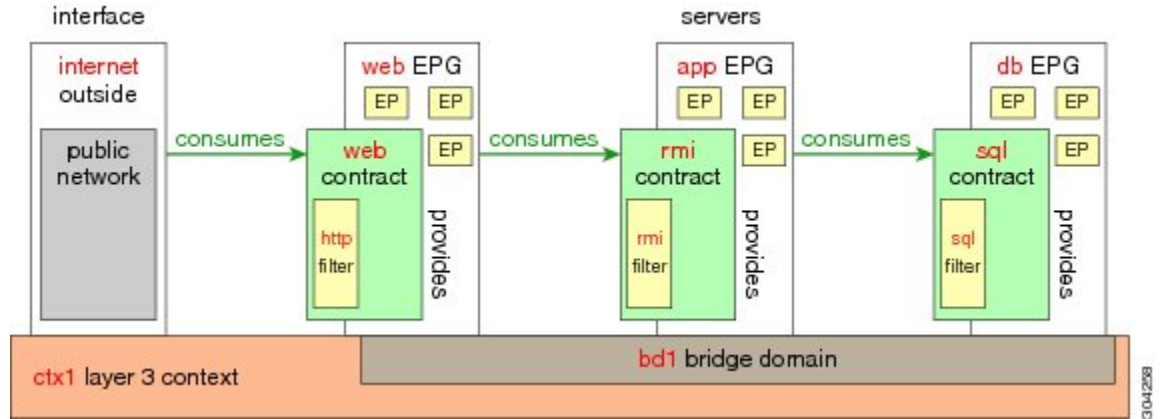
Application profiles enable you to model application requirements that the APIC then automatically renders in the network and data center infrastructure. The application profiles enable administrators to approach the resource pool in terms of applications rather than infrastructure building blocks. The application profile is a container that holds EPGs that are logically related to one another. EPGs can communicate with other EPGs in the same application profile and with EPGs in other application profiles.

To deploy an application policy, you must create the required application profiles, filters, and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

The web server has the HTTP filter, the application server has the Remote Method Invocation (RMI) filter, and the database server has the Structured Query Language (SQL) filter. The application server consumes the SQL contract to communicate with the database server. The web server consumes the RMI contract to communicate with the application server. The traffic enters from the web server and communicates with the

application server. The application server then communicates with the database server, and the traffic can also communicate externally.

Figure 11: Three-Tier Application Diagram



Parameters to Create a Filter for http

The parameters to create a filter for http in this example is as follows:

Parameter Name	Filter for http
Name	http
Number of Entries	2
Entry Name	Dport-80 Dport-443
Ethertype	IP
Protocol	tcp tcp
Destination Port	http https

Parameters to Create Filters for rmi and sql

The parameters to create filters for rmi and sql in this example are as follows:

Parameter Name	Filter for rmi	Filter for sql
Name	rmi	sql

Parameter Name	Filter for rmi	Filter for sql
Number of Entries	1	1
Entry Name	Dport-1099	Dport-1521
Ethertype	IP	IP
Protocol	tcp	tcp
Destination Port	1099	1521

Example Application Profile Database

The application profile database in this example is as follows:

EPG	Provided Contracts	Consumed Contracts
web	web	rmi
app	rmi	sql
db	sql	--

Deploying an Application Policy Using the GUI

Creating a Filter Using the GUI

Create three separate filters. In this example they are HTTP, RMI, SQL. This task shows how to create the HTTP filter. The task is identical for creating the other filters.

Before You Begin

Verify that the tenant, network, and bridge domain have been created.

Procedure

Step 1 On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the **tenant > Security Policies**, right-click **Filters**, and click **Create Filter**.

Note In the **Navigation** pane, you expand the tenant where you want to add filters.

Step 2 In the **Create Filter** dialog box, perform the following actions:

- In the **Name** field, enter the filter name (http).
- Expand **Entries**, and in the **Name** field, enter the name (Dport-80).
- From the **EtherType** drop-down list, choose the EtherType (IP).

- d) From the **IP Protocol** drop-down list, choose the protocol (tcp).
 - e) From the **Destination Port/Range** drop-down lists, choose **http** in the **From** and **To** fields. (http)
 - f) Click **Update**, and click **Submit**.
The newly added filter appears in the **Navigation** pane and in the **Work** pane.
- Step 3** Expand **Entries** in the **Name** field. Follow the same process to add another entry with HTTPS as the **Destination** port, and click **Update**.
This new filter rule is added.
- Step 4** Follow the same process in the earlier steps to create two more filters (rmi and sql) and use the parameters provided in [Parameters to Create Filters for rmi and sql](#), on page 63.
-

Creating a Contract Using the GUI

Procedure

- Step 1** On the menu bar, choose **TENANTS** and the tenant name on which you want to operate. In the **Navigation** pane, expand the **tenant > Security Policies**.
- Step 2** Right-click **Contracts > Create Contract**.
- Step 3** In the **Create Contract** dialog box, perform the following tasks:
- a) In the **Name** field, enter the contract name (web).
 - b) Click the + sign next to **Subjects** to add a new subject.
 - c) In the **Create Contract Subject** dialog box, enter a subject name in the **Name** field. (web)
 - d) **Note** This step associates the filters created that were earlier with the contract subject.
In the **Filter Chain** area, click the + sign next to **Filters**.
 - e) In the dialog box, from the drop-down menu, choose the filter name (http), and click **Update**.
- Step 4** In the **Create Contract Subject** dialog box, click **OK**.
- Step 5** Create two more contracts for rmi and for sql following the same steps in this procedure. For the rmi contract, choose the rmi subject and for sql, choose the sql subject.
-

Creating an Application Profile Using the GUI

Procedure

- Step 1** On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the tenant, right-click **Application Profiles**, and click **Create Application Profile**.
- Step 2** In the **Create Application Profile** dialog box, in the **Name** field, add the application profile name (OnlineStore).
-

Creating EPGs Using the GUI

The port the EPG uses must belong to one of the VM Managers (VMM) or physical domains associated with the EPG.

Procedure

-
- Step 1** On the menu bar, choose **Tenants** and the tenant where you want to create an EPG.
- Step 2** In the navigation pane, expand the folder for the tenant, the **Application Profiles** folder, and the folder for the application profile.
- Step 3** Right-click the **Application EPG** folder, and in the **Create Application EPG** dialog box, perform the following actions:
- a) In the **Name** field, add the EPG name (db).
 - b) In the **Bridge Domain** field, choose the bridge domain from the drop-down list (bd1).
 - c) Check the **Associate to VM Domain Profiles** check box. Click **Next**.
 - d) In the **Step 2 for Specify the VM Domains** area, expand **Associate VM Domain Profiles** and from the drop-down list, choose the desired VMM domain.
 - e) (Optional) In the **Delimiter** field, enter one of the following symbols: |, ~, !, @, ^, +, or =.
If you do not enter a symbol, the system will use the default | delimiter in the VMware portgroup name.
 - f) If you have Cisco AVS, from the **Encap Mode** drop-down list, choose an encapsulation mode.
You can choose one of the following encap modes:
 - **VXLAN**—This overrides the domain's VLAN configuration, and the EPG will use VXLAN encapsulation. However, a fault will be triggered for the EPG if a multicast pool is not configured on the domain.
 - **VLAN**—This overrides the domain's VXLAN configuration, and the EPG will use VLAN encapsulation. However, a fault will be triggered for the EPG if a VLAN pool is not configured on the domain.
 - **Auto**—This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
 - g) Click **Update** and then click **FINISH**.
- Step 4** In the **Create Application Profile** dialog box, create two more EPGs. The three EPGs should be db, app, and web in the same bridge domain and data center.
-

Consuming and Providing Contracts Using the GUI

You can associate contracts that were created earlier to create policy relationships between the EPGs.

When you name the provided and consumed contracts, verify that you give the same name for both provided and consumed contracts.

Procedure

- Step 1** **Note** The db, app, and web EPGs are displayed as icons.
Click and drag across the APIC GUI window from the db EPG to the app EPG.
The **Add Consumed Contract** dialog box is displayed.
- Step 2** In the **Name** field, from the drop-down list, choose **sql** contract. Click **OK**.
This step enables the db EPG to provide the sql contract and the app EPG to consume the sql contract.
- Step 3** Click and drag across the APIC GUI screen from the app ePG to the web EPG.
The **Add Consumed Contract** dialog box is displayed.
- Step 4** In the **Name** field, from the drop-down list, choose **rmi** contract. Click **OK**.
This step enables the app EPG to provide the rmi contract and the web EPG to consume the rmi contract.
- Step 5** Click the web EPG icon, and click the + sign in the **Provided Contracts** area.
The **Add Provided Contract** dialog box is displayed.
- Step 6** In the **Name** field, from the drop-down list, choose **web** contract. Click **OK**. Click **Submit**.
You have created a three-tier application profile called OnlineStore.
- Step 7** To verify, in the **Navigation** pane, navigate to and click **OnlineStore** under **Application Profiles**.
In the **Work** pane, you can see the three EPGs app, db, and web are displayed.
- Step 8** In the **Work** pane, choose **Operational > Contracts**.
You can see the EPGs and contracts displayed in the order that they are consumed and provided.
-



CHAPTER 6

Using the REST API

This chapter contains the following sections:

- [About Getting Started with APIC Examples, page 157](#)
- [About Switch Discovery with the APIC, page 157](#)
- [Configuring Network Time Protocol, page 161](#)
- [Creating User Accounts, page 164](#)
- [Adding Management Access, page 167](#)
- [Configuring a VMM Domain, page 178](#)
- [Creating Tenants, VRF, and Bridge Domains, page 185](#)
- [Configuring Server or Service Policies, page 186](#)
- [Configuring External Connectivity for Tenants, page 192](#)
- [Deploying an Application Policy, page 194](#)

About Getting Started with APIC Examples

The steps in several examples in this guide include a parameter name. These parameter names are provided as examples for convenience and ease of your understanding, and it is not required for you to use them.

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



Note 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>.

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



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.

Registering the Unregistered Switches Using the REST API



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.

Procedure

Register the switches.

Example:

```
POST: https://<apic-ip>/api/node/mo/uni/controller.xml
<fabricNodeIdentPol>
<fabricNodeIdentP serial="FGE173900ZD" name="leaf1" nodeId="101"/>
<fabricNodeIdentP serial="FGE1740010A" name="leaf2" nodeId="102"/>
<fabricNodeIdentP serial="FGE1740010H" name="spine1" nodeId="203"/>
<fabricNodeIdentP serial="FGE1740011B" name="spine2" nodeId="204"/>
</fabricNodeIdentPol>
```

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 REST API

Procedure

Validate switch registration using the REST API.

Example:

GET: `https://<apic-ip>/api/node/class/topSystem.xml?`

```
<?xml version="1.0" encoding="UTF-8"?>
<imdata>
  <topSystem address="10.0.0.1" dn="topology/pod-1/node-1/sys" fabricId="1" id="1"
name="apic1"
  oobMgmtAddr="10.30.13.44" podId="1" role="apic" serial="" state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.0.2" dn="topology/pod-1/node-2/sys" fabricId="1" id="2"
name="apic2"
  oobMgmtAddr="10.30.13.45" podId="1" role="apic" serial="" state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.0.3" dn="topology/pod-1/node-3/sys" fabricId="1" id="3"
name="apic3"
  oobMgmtAddr="10.30.13.46" podId="1" role="apic" serial="" state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.98.127" dn="topology/pod-1/node-101/sys" fabricId="1" id="101"
name="leaf1" oobMgmtAddr="0.0.0.0" podId="1" role="leaf" serial="FOX-270308"
state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.98.124" dn="topology/pod-1/node-102/sys" fabricId="1" id="102"
name="leaf2" oobMgmtAddr="0.0.0.0" podId="1" role="leaf" serial="FOX-270308"
state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.98.125" dn="topology/pod-1/node-203/sys" fabricId="1" id="203"
name="spine2" oobMgmtAddr="0.0.0.0" podId="1" role="spine" serial="FOX-616689"
state="in-service"
systemUpTime="00:00:00:02.199" .../>
  <topSystem address="10.0.98.126" dn="topology/pod-1/node-204/sys" fabricId="1" id="204"
name="spine1" oobMgmtAddr="0.0.0.0" podId="1" role="spine" serial="FOX-616689"
state="in-service"
systemUpTime="00:00:00:02.199" .../>
</imdata>
```

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 REST API

Procedure

Validate the fabric topology using the REST API.

Example:

Identifiers for user reference are as follows:

- n1 = Identifier of the first node
- s1 = Slot on the first node
- p1 = Port on slot s1
- n2 = Identifier of the second node
- s2 = Slot on the second node
- p2 = Port on slot s2

GET: <https://<apic-ip>/api/node/class/fabricLink.xml?>


```
<?xml version="1.0" encoding="UTF-8"?>
<imdata>
  <fabricLink dn="topology/lkcnt-19/lk-18-1-50-to-19-5-2" n1="18" n2="19" p1="50" p2="2"
s1="1" s2="5" status="" .../>
  <fabricLink dn="topology/lkcnt-20/lk-18-1-49-to-20-5-1" n1="18" n2="20" p1="49" p2="1"
s1="1" s2="5" status="" .../>
  <fabricLink dn="topology/lkcnt-3/lk-18-1-1-to-3-1-1" n1="18" n2="3" p1="1" p2="1"
s1="1" s2="1" status="" .../>
  <fabricLink dn="topology/lkcnt-19/lk-17-1-49-to-19-5-1" n1="17" n2="19" p1="49" p2="1"
s1="1" s2="5" status="" .../>
  <fabricLink dn="topology/lkcnt-20/lk-17-1-50-to-20-5-2" n1="17" n2="20" p1="50" p2="2"
s1="1" s2="5" status="" .../>
  <fabricLink dn="topology/lkcnt-1/lk-17-1-1-to-1-1-1" n1="17" n2="1" p1="1" p2="1"
s1="1" s2="1" status="" .../>
  <fabricLink dn="topology/lkcnt-2/lk-17-1-2-to-2-1-1" n1="17" n2="2" p1="2" p2="1"
s1="1" s2="1" status="" .../>
</imdata>
```

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) or Cisco Discovery Protocol (CDP) 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. The following figure shows the APIC GUI displaying unmanaged switches in the **Fabric > Inventory** view.

Figure 12: Unmanaged Layer 2 Switches in the APIC Fabric Inventory



Configuring Network Time Protocol

Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

In-Band and Out-of-Band Management NTP



Note

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
- See the Adding Management Access section in this guide for information about in-band management access and out-of-band management access.

- Out-of-band management NTP—When an ACI fabric is deployed with out-of-band management, each node of the fabric, inclusive of spines, leaves, and all members of the APIC cluster, is managed from outside the ACI fabric. This IP reachability will be leveraged so that each node can individually query the same NTP server as a consistent clock source. To configure NTP, a Date and Time policy must be created that references an out-of-band management endpoint group. Date and Time policies are confined to a single pod and must be deployed across all pods provisioned in the ACI fabric. Currently only one pod per ACI fabric is allowed.
- In-Band Management NTP—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication. The steps used to configure in-band management policies are identical to those used to establish an out-of-band management policy. The distinction is around how to allow the fabric to connect to the NTP server.

NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The `gai.conf` can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

Configuring NTP Using the REST API

Procedure

Step 1 Configure NTP.

Example:

POST url: `https://APIC-IP/api/node/mo/uni/fabric/time-test.xml`

```
<imdata totalCount="1">
  <datetimePol adminSt="enabled" authSt="disabled" descr="" dn="uni/fabric/time-CiscoNTPPol"
    name="CiscoNTPPol" ownerKey="" ownerTag="">
    <datetimeNtpProv descr="" keyId="0" maxPoll="6" minPoll="4" name="10.10.10.11"
      preferred="yes">
      <datetimeRsNtpProvToEpg tDn="uni/tn-mgmt/mgmttp-default/inb-default"/>
    </datetimeNtpProv>
  </datetimePol>
</imdata>
```

```

    </datetimePol>
  </imdata>

```

Step 2 Add the default Date Time Policy to the pod policy group.

Example:

```

POST url: https://APIC-IP/api/node/mo/uni/fabric/funcprof/podpgrp-cal01/rsTimePol.xml

POST payload: <imdata totalCount="1">
<fabricRsTimePol tnDatetimePolName="CiscoNTPPol">
</fabricRsTimePol>
</imdata>

```

Step 3 Add the pod policy group to the default pod profile.

Example:

```

POST url:
https://APIC-IP/api/node/mo/uni/fabric/podprof-default/pods-default-ty-ALL/rspodPGrp.xml

payload: <imdata totalCount="1">
<fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-cal01" status="created">
</fabricRsPodPGrp>
</imdata>

```

Verifying NTP Operation Using the GUI

Procedure

-
- Step 1** On the menu bar, choose **FABRIC > Fabric Policies**.
 - Step 2** In the **Navigation** pane, choose **Pod Policies > Policies > Date and Time > ntp_policy > server_name**. The *ntp_policy* is the previously created policy. An IPv6 address is supported in the Host Name/IP address field. If you enter a hostname and it has an IPv6 address set, you must implement the priority of IPv6 address over IPv4 address.
 - Step 3** In the **Work** pane, verify the details of the server.
-

Verifying NTP Policy Deployed to Each Node Using the CLI

Procedure

-
- Step 1** SSH to an APIC in the fabric.
 - Step 2** Press the Tab key two times after entering the attach command to list all the available node names:

Example:

```
admin@apic1:~> attach <Tab> <Tab>
```

Step 3 Log in one of the nodes using the same password that you used to access the APIC.

Example:

```
admin@apic1:~> attach node_name
```

Step 4 View the NTP peer status.

Example:

```
leaf-1# show ntp peer-status
```

A reachable NTP server has its IP address prefixed by an asterisk (*), and the delay is a non-zero value.

Step 5 Repeat steps 3 and 4 to verify each node in the fabric.

Creating User Accounts

Configuring a Local User

In the initial configuration script, the admin account is configured and the admin is the only user when the system starts. The APIC supports a granular, role-based access control system where user accounts can be created with various roles including non-admin users with fewer privileges.

Configuring a Remote User

Instead of configuring local users, you can point the APIC at the centralized enterprise credential datacenter. The APIC supports Lightweight Directory Access Protocol (LDAP), active directory, RADIUS, and TACACS+.

**Note**

When an APIC is in minority (disconnected from the cluster), remote logins can fail because the ACI is a distributed system and the user information is distributed across APICS. Local logins, however, continue to work because they are local to the APIC.

To configure a remote user authenticated through an external authentication provider, you must meet the following prerequisites:

- The DNS configuration should have already been resolved with the hostname of the RADIUS server.
- You must configure the management subnet.

Configuring a Local User Using the REST API

Procedure

Create a local user.

Example:

URL: `https://apic-ip-address/api/policymgr/mo/uni/userext.xml`

POST CONTENT:

```
<aaaUser name="operations" phone="" pwd="<strong_password>" >
  <aaaUserDomain childAction="" descr="" name="all" rn="userdomain-all" status="">
    <aaaUserRole childAction="" descr="" name="Ops" privType="writePriv"/>
  </aaaUserDomain>
</aaaUser>
```

AV Pair on the External Authentication Server

The Cisco APIC requires that an administrator configure a Cisco AV Pair on an external authentication server. The Cisco AV pair specifies the APIC required RBAC roles and privileges for the user. The Cisco AV Pair format is the same for RADIUS, LDAP, or TACACS+.

To configure a Cisco AV Pair on an external authentication server, an administrator adds a Cisco AV pair to the existing user record. The Cisco AV pair format is as follows:

```
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2
shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,
domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(16003)
```

The first av-pair format has no UNIX user ID, while the second one does. Both are correct if all remote users have the same role and mutual file access is acceptable. If the UNIX user ID is not specified, ID 23999 is applied by the APIC system, and more than one role/read privilege is specified to any AV Pair user. This can cause users to have higher or lower permissions than configured through the group settings.



Note

The APIC Cisco AV-pair format is compatible and can co-exist with other Cisco AV-pair formats. APIC will pick up the first matching AV-pair from all the AV-pairs.

The APIC supports the following regexes:

```
shell:domains\\s* [=:]\\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31}) (\\ (\\d+\\))$
shell:domains\\s* [=:]\\s* ((\\S+?/\\S*?/\\S*?) (, \\S+?/\\S*?/\\S*?) {0, 31})$
```

Examples:

- Example 1: A Cisco AV Pair that contains a single Login domain with only writeRoles:

```
shell:domains=domainA/writeRole1|writeRole2/
```

- Example 2: A Cisco AV Pair that contains a single Login domain with only readRoles:

```
shell:domains=domainA//readRole1|readRole2
```

**Note**

The "/" character is a separator between writeRoles and readRoles per Login domain and is required even if only one type of role is to be used.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

```
aaa-network-admin Cleartext-Password := "<password>"
Cisco-avpair = "shell:domains = all/aaa/read-all(16001)"
```

Changing the Default Behavior for Remote Users with Missing or Bad Cisco AV Pairs

Procedure

-
- Step 1** On the menu bar, click **ADMIN > AAA**.
 - Step 2** In the **Navigation** pane, click **AAA Authentication**.
 - Step 3** In the **Work** pane, in the **Properties** area, from the **Remote user login policy** drop-down list, choose **Assign Default Role**.

The default value is **No Login**. The **Assign Default Role** option assigns the minimal read-only privileges to users that have missing or bad Cisco AV Pairs. Bad AV Pairs are those AV Pairs that fail the parsing rules.

Best Practice for Assigning AV Pairs

As best practice, Cisco recommends that you assign unique UNIX user ids in the range 16000-23999 for the AV Pairs that are assigned to users when in bash shell (using SSH, Telnet or Serial/KVM consoles). If a situation arises when 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 remote users with a UNIX ID of 23999.

The Cisco AVpair string is case sensitive. Although a fault may not be seen, using mismatching cases for the domain name or roles could lead to unexpected privileges being given.

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is logged in using Secure Shell (SSH) or Telnet.

Procedure

Configure an AV pair on the external authentication server.

The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

```
* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2

* shell:domains =
domainA/writeRole1|writeRole2|writeRole3/readRole1|readRole2,domainB/writeRole1|writeRole2|writeRole3/readRole1|readRole2(8101)

These are the boost regexes supported by APIC:
uid_regex("shell:domains\\s*[:]=:]\\s*(\\S+?/\\S*?/\\S*?)(,\\S+?/\\S*?/\\S*?){0,31})(\\(\\d+\\))$");
regex("shell:domains\\s*[:]=:]\\s*(\\S+?/\\S*?/\\S*?)(,\\S+?/\\S*?/\\S*?){0,31})$");
```

The following is an example:

```
shell:domains = coke/tenant-admin/read-all,pepsi//read-all(16001)
```

Configuring a Remote User Using the REST API

Procedure

Step 1 Create a RADIUS provider.

Example:

```
URL: https://apic-ip-address/api/policymgr/mo/uni/userext/radiusext.xml
POST Content:
<aaaRadiusProvider name="radius-auth-server.org.com" key="test123" />
```

Step 2 Create a login domain.

Example:

```
URL: https://apic-ip-address/api/policymgr/mo/uni/userext.xml
POST Content:
<aaaLoginDomain name="rad"> <aaaDomainAuth realm="radius"/> </aaaLoginDomain>
```

Adding Management Access

An APIC controller has two routes to reach the management network, one is by using the in-band management interface and the other is by using the out-of-band management interface.

- In-band management access—You can configure in-band management connectivity to the APIC and the ACI fabric. You first configure the VLANs that will be used by APIC when the APIC is communicating with the leaf switches, and then you configure the VLANs that the VMM servers will use to communicate with the leaf switches.
- Out-of-band management access—You can configure out-of-band management connectivity to the APIC and the ACI fabric. You configure an out-of-band contract that is associated with an out-of-band endpoint group (EPG), and attach the contract to the external network profile.



Note The APIC out-of-band management connection link must be 1 Gbps.

The APIC controller always selects the in-band management interface over the out-of-band management interface, if the in-band management interface is configured. The out-of-band management interface is used only when the in-band management interface is not configured, or if the destination address is on the same subnet as the out-of-band management subnet of the APIC.

The APIC management interface does not support an IPv6 address and cannot connect to an external IPv6 server through this interface.

Configuring the external management instance profile under the management tenant for in-band or out-of-band has no effect on the protocols that are configured under the fabric-wide communication policies. The subnets and contracts specified under the external management instance profile do not affect HTTP/HTTPS or SSH/Telnet.

IPv4/IPv6 Addresses and In-Band Policies

In-band management addresses can be provisioned on the APIC controller only through a policy (Postman REST API, NX-OS Style CLI, or GUI). Additionally, the in-band management addresses must be configured statically on each node.

IPv4/IPv6 Addresses in Out-of-Band Policies

Out-of-band management addresses can be provisioned on the APIC controller either at the time of bootstrap or by using a policy (Postman REST API, NX-OS Style CLI, GUI). Additionally, the out-of-band management addresses must be configured statically on each node or by specifying a range of addresses (IPv4/IPv6) to the entire cluster. IP addresses are randomly assigned from a range to the nodes in the cluster.

Configuring Management Access

Configuring In-Band Management Access Using the REST API

IPv4 and IPv6 addresses are supported for in-band management access. IPv6 configurations are supported using static configurations (for both in-band and out-of-band). IPv4 and IPv6 dual in-band and out-of-band configurations are supported only through static configuration. For more information, see the KB article, *Configuring Static Management Access in Cisco APIC*.

Procedure

Step 1 Create a VLAN namespace.

Example:

```
POST
https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/uni.xml -->
<polUni>
  <infraInfra>
    <!-- Static VLAN range -->
    <fvnsVlanInstP name="inband" allocMode="static">
      <fvnsEncapBlk name="encap" from="vlan-10" to="vlan-11"/>
    </fvnsVlanInstP>
  </infraInfra>
</polUni>
```

```

    </fvnsVlanInstP>
  </infraInfra>
</polUni>

```

Step 2 Create a physical domain.

Example:

```

POST
https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/uni.xml -->
<polUni>
  <physDomP name="inband">
    <infraRsVlanNs tDn="uni/infra/vlanns-inband-static"/>
  </physDomP>
</polUni>

```

Step 3 Create selectors for the in-band management.

Example:

```

POST
https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/.xml -->
<polUni>
  <infraInfra>
    <infraNodeP name="vmmNodes">
      <infraLeafS name="leafS" type="range">
        <infraNodeBlk name="single0" from_"101" to_"101"/>
      </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-vmmPorts"/>
    </infraNodeP>

    <!-- Assumption is that VMM host is reachable via eth1/40. -->
    <infraAccPortP name="vmmPorts">
      <infraHPortS name="portS" type="range">
        <infraPortBlk name="block1"
          fromCard="1" toCard="1"
          fromPort="40" toPort="40"/>
        <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-inband" />
      </infraHPortS>
    </infraAccPortP>

    <infraNodeP name="apicConnectedNodes">
      <infraLeafS name="leafS" type="range">
        <infraNodeBlk name="single0" from_"101" to_"102"/>
      </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-apicConnectedPorts"/>
    </infraNodeP>

    <!-- Assumption is that APIC is connected to eth1/1. -->
    <infraAccPortP name="apicConnectedPorts">
      <infraHPortS name="portS" type="range">
        <infraPortBlk name="block1"
          fromCard="1" toCard="1"
          fromPort="1" toPort="3"/>
        <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-inband" />
      </infraHPortS>
    </infraAccPortP>

    <infraFuncP>
      <infraAccPortGrp name="inband">
        <infraRsAttEntP tDn="uni/infra/attentp-inband"/>
      </infraAccPortGrp>
    </infraFuncP>

    <infraAttEntityP name="inband">

```

```

        <infraRsDomP tDn="uni/phys-inband"/>
    </infraAttEntityP>
</infraInfra>
</polUni>

```

Step 4 Configure an in-band bridge domain and endpoint group (EPG).

Example:

```

POST https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/.xml -->
<polUni>
  <fvTenant name="mgmt">
    <!-- Configure the in-band management gateway address on the
         in-band BD. -->
    <fvBD name="inb">
      <fvSubnet ip="10.13.1.254/24"/>
    </fvBD>

    <mgmtMgmtP name="default">
      <!-- Configure the encap on which APICs will communicate on the
           in-band network. -->
      <mgmtInB name="default" encap="vlan-10">
        <fvRsProv tnVzBrCPName="default"/>
      </mgmtInB>
    </mgmtMgmtP>
  </fvTenant>
</polUni>

```

Step 5 Create an address pool.

Example:

```

POST
https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/.xml -->
<polUni>
  <fvTenant name="mgmt">
    <!-- Addresses for APIC in-band management network -->
    <fvnsAddrInst name="apicInb" addr="10.13.1.254/24">
      <fvnsUcastAddrBlk from="10.13.1.1" to="10.13.1.10"/>
    </fvnsAddrInst>

    <!-- Addresses for switch in-band management network -->
    <fvnsAddrInst name="switchInb" addr="10.13.1.254/24">
      <fvnsUcastAddrBlk from="10.13.1.101" to="10.13.1.120"/>
    </fvnsAddrInst>
  </fvTenant>
</polUni>

```

Note Dynamic address pools for IPv6 is not supported.

Step 6 Create management groups.

Example:

```

POST
https://apic-ip-address/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/.xml -->
<polUni>
  <infraInfra>
    <!-- Management node group for APICs -->
    <mgmtNodeGrp name="apic">
      <infraNodeBlk name="all" from_"1" to_"3"/>
      <mgmtRsGrp tDn="uni/infra/funcprof/grp-apic"/>
    </mgmtNodeGrp>
  </infraInfra>
</polUni>

```

```

<!-- Management node group for switches-->
<mgmtNodeGrp name="switch">
  <infraNodeBlk name="all" from_="101" to_="104"/>
  <mgmtRsGrp tDn="uni/infra/funcprof/grp-switch"/>
</mgmtNodeGrp>

<!-- Functional profile -->
<infraFuncP>
  <!-- Management group for APICs -->
  <mgmtGrp name="apic">
    <!-- In-band management zone -->
    <mgmtInBZone name="default">
      <mgmtRsInbEpg tDn="uni/tn-mgmt/mgmt-default/inb-default"/>
      <mgmtRsAddrInst tDn="uni/tn-mgmt/addrinst-apicInb"/>
    </mgmtInBZone>
  </mgmtGrp>

  <!-- Management group for switches -->
  <mgmtGrp name="switch">
    <!-- In-band management zone -->
    <mgmtInBZone name="default">
      <mgmtRsInbEpg tDn="uni/tn-mgmt/mgmt-default/inb-default"/>
      <mgmtRsAddrInst tDn="uni/tn-mgmt/addrinst-switchInb"/>
    </mgmtInBZone>
  </mgmtGrp>
</infraFuncP>
</infraInfra>
</polUni>

```

Note Dynamic address pools for IPv6 is not supported.

Configuring Out-of-Band Management Access Using the REST API

IPv4 and IPv6 addresses are supported for out-of-band management access.

Before You Begin

The APIC out-of-band management connection link must be 1 Gbps.

Procedure

Step 1 Create an out-of-band contract.

Example:

POST <https://apic-ip-address/api/mo/uni.xml>

```

<polUni>
  <fvTenant name="mgmt">
    <!-- Contract -->
    <vzOOBBrCP name="oob-default">
      <vzSubj name="oob-default">
        <vzRsSubjFiltAtt tnVzFilterName="default" />
      </vzSubj>
    </vzOOBBrCP>
  </fvTenant>
</polUni>

```

Step 2 Associate the out-of-band contract with an out-of-band EPG.

Example:

```
POST https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
  <fvTenant name="mgmt">
    <mgmtMgmtP name="default">
      <mgmtOoB name="default">
        <mgmtRsOoBProv tnVzOOBBrCPName="oob-default" />
      </mgmtOoB>
    </mgmtMgmtP>
  </fvTenant>
</polUni>
```

Step 3 Associate the out-of-band contract with an external management EPG.**Example:**

```
POST https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
  <fvTenant name="mgmt">
    <mgmtExtMgmtEntity name="default">
      <mgmtInstP name="oob-mgmt-ext">
        <mgmtRsOoBCons tnVzOOBBrCPName="oob-default" />
        <!-- SUBNET from where switches are managed -->
        <mgmtSubnet ip="10.0.0.0/8" />
      </mgmtInstP>
    </mgmtExtMgmtEntity>
  </fvTenant>
</polUni>
```

Step 4 Create a management address pool.**Example:**

```
POST https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
  <fvTenant name="mgmt">
    <fvnsAddrInst name="switchOoboobaddr" addr="172.23.48.1/21">
      <fvnsUcastAddrBlk from="172.23.49.240" to="172.23.49.244"/>
    </fvnsAddrInst>
  </fvTenant>
</polUni>
```

Step 5 Create node management groups.**Example:**

```
POST https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
  <infraInfra>
    <infraFuncP>
      <mgmtGrp name="switchOob">
        <mgmtOoBZone name="default">
          <mgmtRsAddrInst tDn="uni/tn-mgmt/addrinst-switchOoboobaddr" />
          <mgmtRsOobEpg tDn="uni/tn-mgmt/mgmtp-default/oob-default" />
        </mgmtOoBZone>
      </mgmtGrp>
    </infraFuncP>
    <mgmtNodeGrp name="switchOob">
      <mgmtRsGrp tDn="uni/infra/funcprof/grp-switchOob" />
      <infraNodeBlk name="default" from_="101" to_="103" />
    </mgmtNodeGrp>
  </infraInfra>
```

```
</polUni>
```

Note You can configure the APIC server to use out-of-band management connectivity as the default connectivity mode.

```
POST https://apic-ip-address/api/node/mo/.xml
<polUni>
<fabricInst>
  <mgmtConnectivityPrefs interfacePref="ooband"/>
</fabricInst>
</polUni>
```

Modifying the IP Address of an APIC Controller Using the REST API

Procedure

Modify the IP address of the APIC controller.

Example:

```
POST
https://APIC-IP/api/mo/uni.xml

<?xml version="1.0" encoding="UTF-8"?>
<!-- /api/policymgr/mo/.xml -->
<polUni>
  <infraInfra>
    <infraFuncP>
      <mgmtGrp name="mgmtGroupApic">
        <mgmtOobBZone name="mgmtOobZoneApic">
          <mgmtRsOobEpg tDn="uni/tn-mgmt/mgmtp-default/oob-default"/>
          <mgmtRsAddrInst tDn="uni/tn-mgmt/addrinst-oobAddrApic"/>
        </mgmtOobBZone>
      </mgmtGrp>
    </infraFuncP>
    <mgmtNodeGrp name="mgmtNodeGroupApic">
      <mgmtRsGrp tDn="uni/infra/funcprof/grp-mgmtGroupApic"/>
      <infraNodeBlk name="default" from_"1" to_"1"/>
    </mgmtNodeGrp>
  </infraInfra>
</polUni>

<?xml version="1.0" encoding="UTF-8"?>
<!-- /api/policymgr/mo/.xml -->
<polUni>
  <fvTenant name="mgmt">
    <fvnsAddrInst name="oobAddrApic" addr="172.23.48.1/21">
      <fvnsUcastAddrBlk from="172.23.48.16" to="172.23.48.16"/>
    </fvnsAddrInst>
  </fvTenant>
</polUni>
```

What to Do Next

- You must use the new IP address to reconnect to the APIC controller.
- You must delete the old IP address of the controller once a new IP address is assigned to it.

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

All IPv6 tables mirror the existing IP tables functionality, except for Network Address Translation (NAT).

Existing IP Tables

- 1 Earlier, every rule in the IPv6 tables were executed one at a time and a system call was made for every rule addition or deletion.
- 2 Whenever a new policy was added, rules were appended to the existing IP tables file and no extra modifications were done to the file.
- 3 When a new source port was configured in the out-of-band policy, it added source and destination rules with the same port number.

Modifications to IP Tables

- 1 When IP tables are created, they are first written into hash maps that are then written into intermediate file IP tables-new which are restored. When saved, a new IP tables file is created in the `/etc/sysconfig/` folder. You can find both these files at the same location. Instead of making a system call for every rule, you must make a system call only while restoring and saving the file.
- 2 When a new policy is added instead of appending it to the file, an IP table is created from scratch, that is by loading default policies into the hashmaps, checking for new policies, and adding them to hashmaps. Later, they are written to the intermediate file (`/etc/sysconfig/iptables-new`) and saved.
- 3 It is not possible to configure source ports alone for a rule in out-of-band policy. Either destination port or source port along with a destination port can be added to the rules.
- 4 When a new policy is added, a new rule will be added to the IP tables file. This rule changes the access flow of IP tables default rules.


```
-A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
```
- 5 When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.



Note

- If only IPv4 is enabled, do not configure an IPv6 policy.
- If only IPv6 is enabled, do not configure an IPv4 policy.
- If both IPv4 and IPv6 are enabled and a policy is added, it will be configured to both the versions . So when you add an IPv4 subnet, it will be added to IP tables and similarly an IPv6 subnet is added to IPv6 tables.

Management Connectivity Modes

Establish connection to external entities using the out-of-band or in-band network depending upon whether you have configured out-of-band and/or in-band management connectivity. The following two modes are available to establish connectivity to external entities such as the vCenter server:

- Layer 2 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 2.
- Layer 3 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 3 through a router. The leaf is connected to a router through which external entities can be reached.

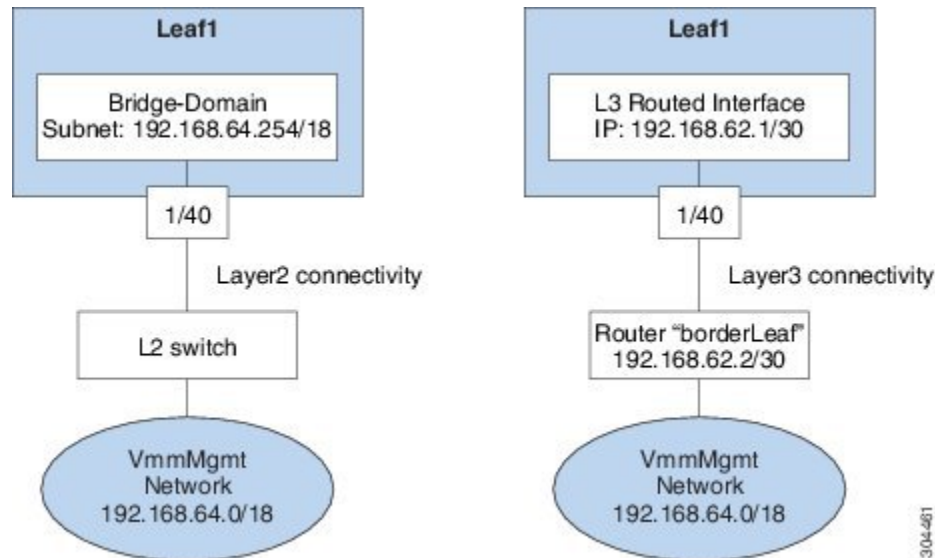


Note

- The inband IP address range must be separate and distinct from the IP address range used on the Layer 3 connection from the leaf node to outside the fabric.
- The Layer 3 inband management design does not provide inband management access to the spine fabric nodes in the topology.

The following diagram displays the two modes available to establish connectivity.

Figure 13: Layer 2 and Layer 3 Management Connectivity Examples



Configuring Layer 2 Management Connectivity Using the REST API



Note

The name vmm is used as an example string in this task.

The policy creates the following objects under Tenant-mgmt:

Creates bridge domain (vmm) and the following related objects as follows:

- Creates the subnet object with this IP prefix (192.168.64.254/18) in this bridge domain. This IP address (192.168.64.254) is assigned to the bridge domain that typically is used as the switch virtual interface (SVI) in a traditional switch configuration.
- Creates an association to the in-band network (ctx).

Creates an application profile (vmm) and management EPG (vmmMgmt) with related objects as follows:

- Creates an association to the bridge domain (vmm).
- Creates a policy to deploy this EPG on leaf1. The encapsulation used for this EPG is vlan-11.

Before You Begin

Before you create a vCenter domain profile, you must establish connectivity to establish an external network using an in-band management network.

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

You can establish connectivity from the APIC to external routes by using a router that is connected to a leaf port.

Example:

```
POST
https://APIC-IP/api/mo/uni.xml

<!-- api/policymgr/mo/.xml -->
<polUni>
  <fvTenant name="mgmt">
    <fvBD name="vmm">
      <fvRsCtx tnFvCtxName="inb"/>
      <fvSubnet ip='192.168.64.254/18'/>
    </fvBD>

    <fvAp name="vmm">
      <fvAEPg name="vmmMgmt">
        <fvRsBd tnFvBDName="vmm" />
        <fvRsPathAtt tDn="topology/pod-1/paths-101/pathep-[eth1/40]" encap="vlan-11"/>
        <fvRsCons tnVzBrCPName="default"/>
        <fvRsDomAtt tDn="uni/phys-inband"/>
      </fvAEPg>
    </fvAp>
  </fvTenant>
</polUni>
```

Configuring Layer 3 Management Connectivity Using the REST API

The name vmm is used as an example string in this task.

The policy creates the following objects under Tenant-mgmt:

- Creates a routed outside policy (vmm) with the following details:
 - 1 Creates a Layer 3 external network instance profile object (vmmMgmt).

- 2 Creates a route for the remote network (192.168.64.0/18) with the IP address of the next-hop router 192.168.62.2.
- 3 Creates a logical node profile object (borderLeaf) that is attached to leaf1.
- 4 Creates a port profile (portProfile1) with the routed interface 1/40 with the IP address 192.168.62.1/30.
- 5 Creates an association to inband network (ctx).

Before You Begin

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

You can establish connectivity from the APIC to external routes by using a router that is connected to a leaf port.

Example:

```

<!-- api/policymgr/mo/.xml -->
POST
https://APIC-IP/api/mo/uni.xml

<polUni>
  <fvTenant name="mgmt">
    <l3extOut name="vmm">
      <l3extInstP name="vmmMgmt">
        <l3extSubnet ip="192.168.0.0/16" />
        <fvRsCons tnVzBrCPName="default" />
      </l3extInstP>
      <l3extLNodeP name="borderLeaf">
        <l3extRsNodeL3OutAtt tDn="topology/pod-1/node-101" rtrId="1.2.3.4">
          <ipRouteP ip="192.168.64.0/18">
            <ipNextHopP nhAddr="192.168.62.2" />
          </ipRouteP>
        </l3extRsNodeL3OutAtt>
        <l3extLIfP name="portProfile">
          <l3extRsPathL3OutAtt tDn="topology/pod-1/paths-101/pathep-[eth1/40]"
ifInstT="l3-port" addr="192.168.62.1/30" />
        </l3extLIfP>
      </l3extLNodeP>
      <l3extRsEctx tnFvCtxName="inb" />
    </l3extOut>
  </fvTenant>
</polUni>

```

Validating Management Connectivity

This validation process applies to both Layer 2 and Layer 3 modes and can be used to verify connectivity that is established by using the APIC GUI, REST API, or CLI.

After completing the steps to establish management connectivity, log in to the APIC console. Ping to the IP address of the vCenter server that is reachable (for example, 192.168.81.2) and verify that the ping works. This action indicates that the policies have been successfully applied.

Configuring a VMM Domain

Configuring Virtual Machine Networking Policies

The APIC integrates with third-party VM manager (VMM) (for example, VMware vCenter and SCVMM) to extend the benefits of ACI to the virtualized infrastructure. The APIC enables the ACI policies inside the VMM system to be used by its administrator.

This section provides examples of VMM integration using VMware vCenter and vShield. For details about the different modes of Cisco ACI and VMM integration, see the *ACI Virtualization Guide*.

About the VM Manager

**Note**

Information about the necessary configuration of the APIC for integration with the vCenter is described here. For instructions about configuring the VMware components, see the VMware documentation.

The following are details of some VM manager terms:

- A VM controller is an external virtual machine management entity such as VMware vCenter, and the VMware vShield. The APIC communicates with the controller to publish network policies that are applied to virtual workloads. A VM controller administrator provides an APIC administrator with a VM controller authentication credential; multiple controllers of the same type can use the same credential.
- Credentials represent the authentication credentials to communicate with VM controllers. Multiple controllers can use the same credentials.
- A virtual machine mobility domain (vCenter mobility domain) is a grouping of VM controllers with similar networking policy requirements. This mandatory container holds one or more VM controllers with policies such as for a VLAN pool, server to network MTU policy, or server to network access LACP policy. When an endpoint group gets associated with a vCenter domain, network policies get pushed to all the VM controllers in the vCenter domain.
- A pool represents a range of traffic encapsulation identifiers (for example, VLAN IDs, VNIDs, and multicast addresses). A pool is a shared resource and can be consumed by multiple domains such as VMM and Layer 4 to Layer 7 services. A leaf switch does not support overlapping VLAN pools. You must not associate different overlapping VLAN pools with the VMM domain. The two types of VLAN-based pools are as follows:
 - Dynamic pools—Managed internally by the APIC to allocate VLANs for endpoint groups (EPGs). A vCenter Domain can associate only to a dynamic pool.
 - Static pools—The EPG has a relation to the domain, and the domain has a relation to the pool. The pool contains a range of encapsulated VLANs and VXLANs. For static EPG deployment, the user defines the interface and the encapsulation. The encapsulation must be within the range of a pool that is associated with a domain with which the EPG is associated.
- For a VMware vCenter to be deployed, it must operate in VLAN mode or VXLAN mode. A VMM domain must be associated with a VLAN pool and a vShield must be associated with the vCenter.

About Attachable Entity Profile

Attach Entity Profiles

The ACI fabric provides multiple **attachment points** that connect through leaf ports to various **external entities** such as baremetal servers, hypervisors, Layer 2 switches (for example, the Cisco UCS fabric interconnect), and Layer 3 routers (for example Cisco Nexus 7000 Series switches). These attachment points can be physical ports, port channels, or a virtual port channel (vPC) on the leaf switches.

An **attachable entity profile** (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies, for example, Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), maximum transmission unit (MTU), and Link Aggregation Control Protocol (LACP).

A VM manager (VMM) domain automatically derives the physical interfaces policies from the interface policy groups that are associated with an AEP.

- An override policy at AEP can be used to specify a different physical interface policy for a VMM domain. This policy is useful in scenarios where a hypervisor is connected to the leaf switch through an intermediate Layer 2 node, and a different policy is desired at the leaf switch and hypervisor physical ports. For example, you can configure LACP between a leaf switch and a Layer 2 node. At the same time, you can disable LACP between the hypervisor and the Layer 2 switch by disabling LACP under the AEP override policy.

An AEP is required to deploy any VLAN pools on the leaf switches. It is possible to reuse the encapsulation pools (for example, VLAN) across different leaf switches. An AEP implicitly provides the scope of the VLAN pool (associated to the domain) to the physical infrastructure.



Note

- An AEP provisions the VLAN pool (and associated VLANs) on the leaf. The VLANs are not actually enabled on the port. No traffic flows unless an EPG is deployed on the port.
- Without VLAN pool deployment using an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
 - A particular VLAN is provisioned or enabled on the leaf port based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter.
 - If you wish to set the VMM encapsulation statically in the EPG, you must use a static pool. If you have a mix of static and dynamic allocations, create a dynamic pool and add a block within that pool with static mode.
- A leaf switch does not support overlapping VLAN pools. Different overlapping VLAN pools must not be associated with the same AEP that is associated through a domain.

For information about configuring LLDP and CDP, see the chapter related to Working with Blade Servers in the guide.

Prerequisites for Creating a VMM Domain Profile

To configure a VMM domain profile, you must meet the following prerequisites:

- All fabric nodes are discovered and configured.
- Inband (inb) or out-of-band (oob) management has been configured on the APIC.
- A Virtual Machine Manager (VMM) is installed, configured, and reachable through the inb/oob management network (for example, a vCenter).
- You have the administrator/root credentials to the VMM (for example vCenter).



Note If you prefer not to use the vCenter admin/root credentials, you can create a custom user account with minimum required permissions. See [Custom User Account with Minimum VMware vCenter Privileges](#), on page 47 for a list of the required user privileges.

- A DNS policy for the APIC must be configured if you plan to reference the VMM by hostname rather than an IP address.

Custom User Account with Minimum VMware vCenter Privileges

This allows the APIC to send VMware API commands to vCenter to allow the creation of the DVS/AVS, creation of the VMK interface (AVS), publish port groups and relay all necessary alerts.

To configure the vCenter from Cisco APIC, your credentials must allow the following minimum set of privileges within the vCenter:

- **Alarms**

APIC creates two alarms on the folder. One for DVS and another for port-group. The alarm is raised when the EPG or Domain policy is deleted on APIC, but for port-group or DVS it cannot be deleted due to the VMs are attached.

- **Distributed Switch**

- **dvPort Group**

- **Folder**

- **Network**

APIC manages the network settings such as add or delete port-groups, setting host/DVS MTU, LLDP/CDP, LACP etc.

- **Host**

If you use AVS in addition to above, you need the Host privilege on the data center where APIC will create DVS.

- **Host.Configuration.Advanced settings**
- **Host.Local operations.Reconfigure virtual machine**
- **Host.Configuration.Network configuration**

This is needed for AVS and the auto-placement feature for virtual Layer 4 to Layer 7 Service VMs. For AVS, APIC creates VMK interface and places it in 'vtep' port-group which is used for OpFlex.

- **Virtual machine**

If you use Service Graph in addition to above, you need the Virtual machine privilege for the virtual appliances which will be used for Service Graph.

- **Virtual machine.Configuration.Modify device settings**
- **Virtual machine.Configuration.Settings**

Creating a VMM Domain Profile

In this section, examples of a VMM domain are vCenter domain.

Creating a vCenter Domain Profile Using the REST API

Procedure

Step 1 Configure a VMM domain name, a controller, and user credentials.

Example:

POST URL: `https://<api-ip>/api/node/mo/.xml`

```
<polUni>
<vmmProvP vendor="VMware">
<!-- VMM Domain -->
<vmmDomP name="productionDC">
<!-- Association to VLAN Namespace -->
<infraRsVlanNs tDn="uni/infra/vlanns-VlanRange-dynamic"/>
<!-- Credentials for vCenter -->
<vmmUsrAccP name="admin" usr="administrator" pwd="admin" />
<!-- vCenter IP address -->
<vmmCtrlrP name="vcenter1" hostOrIp="<vcenter ip address>" rootContName="<Datacenter Name
in vCenter>">
<vmmRsAcc tDn="uni/vmmp-VMware/dom-productionDC/usracc-admin"/>
</vmmCtrlrP>
</vmmDomP>
</vmmProvP>
```

Example:

```
<polUni>
<vmmProvP vendor="VMware">
  <vmmDomP name="mininet" delimiter="@ " >
    </vmmDomP>
</vmmProvP>
</polUni>
```

Step 2 Create an attachable entity profile for VLAN namespace deployment.

Example:

POST URL: `https://<apic-ip>/api/policymgr/mo/uni.xml`
`<infraInfra>`

```
<infraAttEntityP name="profile1">
<infraRsDomP tDn="uni/vmmp-VMware/dom-productionDC"/>
</infraAttEntityP>
</infraInfra>
```

Step 3 Create an interface policy group and selector.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```
<infraInfra>
  <infraAccPortP name="swprofilelifselector">
    <infraHPortS name="selector1" type="range">
      <infraPortBlk name="blk"
        fromCard="1" toCard="1" fromPort="1" toPort="3">
      </infraPortBlk>
    </infraHPortS>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-group1" />
    </infraHPortS>
  </infraAccPortP>

  <infraFuncP>
    <infraAccPortGrp name="group1">
      <infraRsAttEntP tDn="uni/infra/attentp-profile1" />
    </infraAccPortGrp>
  </infraFuncP>
</infraInfra>
```

Step 4 Create a switch profile.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```
<infraInfra>
  <infraNodeP name="swprofile1">
    <infraLeafS name="selectorswprofile11718" type="range">
      <infraNodeBlk name="single0" from_"101" to_"101"/>
      <infraNodeBlk name="single1" from_"102" to_"102"/>
    </infraLeafS>
    <infraRsAccPortP tDn="uni/infra/accportprof-swprofilelifselector"/>
  </infraNodeP>
</infraInfra>
```

Step 5 Configure the VLAN pool.

Example:

POST URL: <https://<apic-ip>/api/node/mo/.xml>

```
<polUni>
<infraInfra>
<fvnsVlanInstP name="VlanRange" allocMode="dynamic">
  <fvnsEncapBlk name="encap" from="vlan-100" to="vlan-400"/>
</fvnsVlanInstP>
</infraInfra>
</polUni>
```

Step 6 Locate all the configured controllers and their operational state.

Example:

```
GET:
https://<apic-ip>/api/node/class/compCtrlr.xml?
<imdata>
<compCtrlr apiVer="5.1" ctrlrPKey="uni/vmmp-VMware/dom-productionDC/ctrlr-vcenter1"
deployIssues="" descr="" dn="comp/prov-VMware/ctrlr-productionDC-vcenter1" domName="
productionDC"
hostOrIp="esx1" mode="default" model="VMware vCenter Server 5.1.0 build-756313"
```



```
name="vcenter1" operSt="online" port="0" pwd="" remoteOperIssues="" scope="vm"
usr="administrator" vendor="VMware, Inc." ... />
</imdata>
```

- Step 7** Locate the hypervisor and VMs for a vCenter with the name 'vcenter1' under a VMM domain called 'ProductionDC'.

Example:

```
GET:
https://<apic-ip>/api/node/mo/comp/prov-VMware/ctrlr-productionDC-vcenter1.xml?query-target=children

<imdata>
<compHv descr="" dn="comp/prov-VMware/ctrlr-productionDC-vcenter1/hv-host-4832" name="esx1"
state="poweredOn" type="hv" ... />
<compVm descr="" dn="comp/prov-VMware/ctrlr-productionDC-vcenter1/vm-vm-5531" name="AppVM1"
state="poweredOff" type="virt" .../>
<hvsLNode dn="comp/prov-VMware/ctrlr-productionDC-vcenter1/sw-dvs-5646" lACPEnable="yes"
lACPMode="passive" lDPConfigOperation="both" lDPConfigProtocol="lldp" maxMtu="1500"
mode="default" name="apicVswitch" .../>
</imdata>
```

Creating a vCenter and a vShield Domain Profile Using the REST API

Procedure

- Step 1** Create a VLAN pool.

Example:

```
POST URL: https://<apic-ip>/api/policymgr/mo/uni.xml

<polUni>
  <infraInfra>
    <fvnsVlanInstP name="vlan1" allocMode="dynamic">
      <fvnsEncapBlk name="encap" from="vlan-100" to="vlan-400"/>
    </fvnsVlanInstP>
  </infraInfra>
</polUni>
```

- Step 2** Create a vCenter domain, and assign a VLAN pool.

Example:

```
POST URL: https://<apic-ip>/api/policymgr/mo/uni.xml
<vmmProvP dn="uni/vmmp-VMware">
<vmmDomP name="productionDC">
<infraRsVlanNs tDn="uni/infra/vlanns-vlan1-dynamic"/>
</vmmDomP>
</vmmProvP>
```

- Step 3** Create an attachable entity profile for infrastructure VLAN deployment.

Example:

```
POST URL: https://<apic-ip>/api/policymgr/mo/uni.xml

<infraInfra>
  <infraAttEntityP name="profile1">
    <infraRsDomP tDn="uni/vmmp-VMware/dom-productionDC"/>
    <infraProvAcc name="provfunc"/>
  </infraAttEntityP>
</infraInfra>
```

```

    </infraAttEntityP>
  </infraInfra>

```

Step 4 Create an interface policy group and selector.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```

<infraInfra>
  <infraAccPortP name="swprofilelifselector">
    <infraHPortS name="selector1" type="range">
      <infraPortBlk name="blk"
        fromCard="1" toCard="1" fromPort="1" toPort="3">
      </infraPortBlk>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-group1" />
    </infraHPortS>
  </infraAccPortP>

  <infraFuncP>
    <infraAccPortGrp name="group1">
      <infraRsAttEntP tDn="uni/infra/attentp-profile1" />
    </infraAccPortGrp>
  </infraFuncP>
</infraInfra>

```

Step 5 Create a switch profile.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```

<infraInfra>
  <infraNodeP name="swprofile1">
    <infraLeafS name="selectorswprofile11718" type="range">
      <infraNodeBlk name="single0" from_="101" to_="101"/>
      <infraNodeBlk name="single1" from_="102" to_="102"/>
    </infraLeafS>
    <infraRsAccPortP tDn="uni/infra/accportprof-swprofilelifselector"/>
  </infraNodeP>
</infraInfra>

```

Step 6 Create credentials for controllers.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```

<vmmProvP dn="uni/vmmp-VMware">
  <vmmDomP name="productionDC">
    <vmmUsrAccP name="vcenter_user" usr="administrator" pwd="default"/>
    <vmmUsrAccP name="vshield_user" usr="admin" pwd="default"/>
  </vmmDomP>
</vmmProvP>

```

Step 7 Create a vCenter controller

Example:

```

<vmmProvP dn="uni/vmmp-VMware">
  <vmmDomP name="productionDC">
    <vmmCtrlrP name="vcenter1" hostOrIp="172.23.50.85" rootContName="Datacenter1">
      <vmmRsAcc tDn="uni/vmmp-VMware/dom-productionDC/usracc-vcenter_user"/>
    </vmmCtrlrP>
  </vmmDomP>
</vmmProvP>

```

Step 8 Create a VXLAN pool and a multicast address range.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```

<infraInfra>
  <fvnsVxlanInstP name="vxlan1">
    <fvnsEncapBlk name="encap" from="vxlan-6000" to="vxlan-6200"/>
  </fvnsVxlanInstP>
  <fvnsMcastAddrInstP name="multicast1">
    <fvnsMcastAddrBlk name="mcast" from="224.0.0.1" to="224.0.0.20"/>
  </fvnsMcastAddrInstP>
</infraInfra>

```

Step 9 Create a vShield controller.

Example:

POST URL: <https://<apic-ip>/api/policymgr/mo/uni.xml>

```

<vmmProvP dn="uni/vmmp-VMware">
  <vmmDomP name="productionDC">
    <vmmCtrlrP name="vshield1" hostOrIp="172.23.54.62" scope="iaas">
      <vmmRsAcc tDn="uni/vmmp-VMware/dom-productionDC/usracc-vshield_user"/>
      <vmmRsVmmCtrlrP tDn="uni/vmmp-VMware/dom-productionDC/ctrlr-vcenter1"/>
      <vmmRsVxlanNs tDn="uni/infra/vxlanns-vxlan1"/>
      <vmmRsMcastAddrNs tDn="uni/infra/maddrns-multicast1"/>
    </vmmCtrlrP>
  </vmmDomP>
</vmmProvP>

```

Creating Tenants, VRF, and Bridge Domains

Tenants Overview

- 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.

Tenant Creation

A tenant contains primary elements such as filters, contracts, bridge domains, and application profiles that you can create after you first create a tenant.

VRF and Bridge Domains

You can create and specify a VRF and a bridge domain for the tenant. The defined bridge domain element subnets reference a corresponding Layer 3 context.

For details about enabling IPv6 Neighbor Discovery see *IPv6 and Neighbor Discovery* in *Cisco APIC Layer 3 Networking Guide*.

Creating a Tenant, VRF, and Bridge Domain Using the REST API

Procedure

Step 1 Create a tenant.

Example:

```
POST https://apic-ip-address/api/mo/uni.xml
<fvTenant name="ExampleCorp"/>
```

When the POST succeeds, you see the object that you created in the output.

Step 2 Create a VRF and bridge domain.

Note The Gateway Address can be an IPv4 or an IPv6 address. For more about details IPv6 gateway address, see the related KB article, *KB: Creating a Tenant, VRF, and Bridge Domain with IPv6 Neighbor Discovery*.

Example:

URL for POST: `https://apic-ip-address/api/mo/uni/tn-ExampleCorp.xml`

```
<fvTenant name="ExampleCorp">
  <fvCtx name="pvn1"/>
  <fvBD name="bd1">
    <fvRsCtx tnFvCtxName="pvn1"/>
    <fvSubnet ip="10.10.100.1/24"/>
  </fvBD>
</fvTenant>
```

Note If you have a public subnet when you configure the routed outside, you must associate the bridge domain with the outside configuration.

Configuring Server or Service Policies

Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets. If the client is on an ESX hypervisor with a deployed vShield Domain profile, then the use of a DHCP relay policy configuration is mandatory.

When a vShield controller deploys a Virtual Extensible Local Area Network (VXLAN), the hypervisor hosts create a kernel (vmkN, virtual tunnel end-point [VTEP]) interface. These interfaces need an IP address in the infrastructure tenant that uses DHCP. Therefore, you must configure a DHCP relay policy so that the APIC can act as the DHCP server and provide these IP addresses.

When an ACI fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the ACI fabric must support Option 82.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the REST API

- This task is a prerequisite for users who want to create a vShield Domain Profile.
- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Configure the APIC as the DHCP server policy for the infrastructure tenant.

Note This relay policy will be pushed to all the leaf ports that are connected hypervisors using the attach entity profile configuration. For details about configuring with attach entity profile, see the examples related to creating VMM domain profiles.

Example:

DHCP Relay Policy for EPG

```
<!-- api/policymgr/mo/.xml -->
<polUni>
```

```
POST https://apic-ip-address/api/mo/uni.xml
```

```
<fvTenant name="infra">
  <dhcpRelayP name="DhcpRelayP" owner="tenant">
    <dhcpRsProv tDn="uni/tn-infra/ap-access/epg-default" addr="10.0.0.1" />
  </dhcpRelayP>
  <fvBD name="default">
    <dhcpLbl name="DhcpRelayP" owner="tenant"/>
  </fvBD>
</fvTenant>
</polUni>
```

Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.
- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.

- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking > VRF** policy configuration in the tenants configuration.

Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN Note The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be attached to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/node/mo/uni/fabric/dnsp-default.xml -->
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The `gai.conf` settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a `gai.conf` to prioritize IPv4 over IPv6:

```
# Generated by APIC
label ::1/128      0
label ::/0        1
label 2002::/16   2
label ::/96       3
label ::ffff:0:0/96 4
precedence  ::1/128      50
precedence  ::/0        40
precedence  2002::/16   30
precedence  ::/96       20
# For APICs preferring IPv4 connections, change the value to 100.
precedence  ::ffff:0:0/96 10
```

Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPv4) or AAAA records (IPv6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to `/etc/resolv.conf`.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in `/etc/resolv.conf`. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in `/etc/resolv.conf`). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add "options single-request-reopen" to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of resolv.conf in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (glibc) will use the IPv6 network by default because getaddrinfo() will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (::ffff/96). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the glibc IPv6 selection preference for getaddrinfo() in /etc/gai.conf.

In order to allow glibc to return multiple addresses when using /etc/hosts, “multi on” should be added to the /etc/hosts file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

Configuring a DNS Service Policy to Connect with DNS Providers Using the REST API

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Step 1 Configure the DNS service policy.

Example:

```
POST URL :
https://apic-IP-address/api/node/mo/uni/fabric.xml

<dnsProfile name="default">

  <dnsProv addr="172.21.157.5" preferred="yes"/>
  <dnsProv addr="172.21.157.6"/>

  <dnsDomain name="cisco.com" isDefault="yes"/>

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

</dnsProfile>
```

Step 2 Configure the DNS label under the out-of-band management tenant.

Example:

```
POST URL: https://apic-IP-address/api/node/mo/uni/tn-mgmt/ctx-oob.xml
<dnsLbl name="default" tag="yellow-green"/>
```

Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the CLI

Procedure

Step 1 Verify the configuration for the default DNS profile.

Example:

```
admin@apic1:~> cd /aci/fabric/fabric-policies/global-policies/dns-profiles/default
admin@apic1:default> cat summary
# dns-profile
name : default
description : added via CLI by tdeleon@cisco.com
ownerkey :
ownertag :

dns-providers:
address preferred
-----
10.44.124.122 no
10.70.168.183 no
10.37.87.157 no
10.102.6.247 yes
dns-domains:
name default description
-----
cisco.com yes
management-epg : tenants/mgmt/node-management-epgs/default/out-of-band/default
```

Step 2 Verify the configurations for the DNS labels.

Example:

```
admin@apic1:default> cd
/aci/tenants/mgmt/networking/private-networks/oob/dns-profile-labels/default
admin@apic1:default> cat summary
# dns-lbl
name : default
description :
ownerkey :
ownertag :
tag : yellow-green
```

Step 3 Verify that the applied configuration is operating on the fabric controllers.

Example:

```
admin@apic1:~> cat /etc/resolv.conf
# Generated by IFC
search cisco.com
nameserver 10.102.6.247
nameserver 10.44.124.122
nameserver 10.37.87.157
nameserver 10.70.168.183
admin@apic1:~> ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=238 time=35.4 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=238 time=29.0 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=238 time=29.2 ms
```

Step 4 Verify that the applied configuration is operating on the leaf and spine switches.

Example:

```
leaf1# cat /etc/resolv.conf
search cisco.com
nameserver 10.102.6.247
nameserver 10.70.168.183
nameserver 10.44.124.122
nameserver 10.37.87.157
leaf1# cat /etc/dcos_resolv.conf
# DNS enabled
leaf1# ping www.cisco.com
PING origin-www.cisco.com (72.163.4.161): 56 data bytes
64 bytes from 72.163.4.161: icmp_seq=0 ttl=238 time=29.255 ms
64 bytes from 72.163.4.161: icmp_seq=1 ttl=238 time=29.212 ms
64 bytes from 72.163.4.161: icmp_seq=2 ttl=238 time=29.343 ms
```

Configuring External Connectivity for Tenants

Before you can distribute the static route to the other leaf switches on the Application Centric Infrastructure (ACI) fabric, a multiprotocol BGP (MP-BGP) process must first be operating, and the spine switches must be configured as BGP route reflectors.

To integrate the ACI fabric into an external routed network, you can configure Open Shortest Path First (OSPF) for management tenant Layer 3 connectivity.

Configuring an MP-BGP Route Reflector Using the REST API

Procedure

Step 1 Mark the spine switches as route reflectors.

Example:

```
POST https://apic-ip-address/api/policymgr/mo/uni/fabric.xml

<bgpInstPol name="default">
  <bgpAsP asn="1" />
  <bgpRRP>
    <bgpRRNodePEp id="<spine_id1>" />
    <bgpRRNodePEp id="<spine_id2>" />
  </bgpRRP>
</bgpInstPol>
```

Step 2 Set up the pod selector using the following post.

Example:

For the FuncP setup—

```
POST https://apic-ip-address/api/policymgr/mo/uni.xml

<fabricFuncP>
  <fabricPodPGrp name="bgpRRPodGrp">
    <fabricRsPodPGrpBGP tnbGpInstPolName="default" />
  </fabricPodPGrp>
</fabricFuncP>
```

Example:

For the PodP setup—

```
POST https://apic-ip-address/api/policymgr/mo/uni.xml

<fabricPodP name="default">
  <fabricPodS name="default" type="ALL">
    <fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-bgpRRPodGrp"/>
  </fabricPodS>
</fabricPodP>
```

Verifying the MP-BGP Route Reflector Configuration

Procedure

- Step 1** Verify the configuration by performing the following actions:
- Use secure shell (SSH) to log in as an administrator to each leaf switch as required.
 - Enter the **show processes | grep bgp** command to verify the state is S.
If the state is NR (not running), the configuration was not successful.
- Step 2** Verify that the autonomous system number is configured in the spine switches by performing the following actions:
- Use the SSH to log in as an administrator to each spine switch as required.
 - Execute the following commands from the shell window

Example:

```
cd /mit/sys/bgp/inst
```

Example:

```
grep asn summary
```

The configured autonomous system number must be displayed. If the autonomous system number value displays as 0, the configuration was not successful.

Creating OSPF External Routed Network for Management Tenant Using REST API

- You must verify that the router ID and the logical interface profile IP address are different and do not overlap.
- The following steps are for creating an OSPF external routed network for a management tenant. To create an OSPF external routed network for a tenant, you must choose a tenant and create a VRF for the tenant.

- For more details, see *Cisco APIC and Transit Routing*.

Procedure

Create an OSPF external routed network for management tenant.

Example:

POST: <https://apic-ip-address/api/mo/uni/tn-mgmt.xml>

```
<fvTenant name="mgmt">
  <fvBD name="bd1">
    <fvRsBDToOut tnL3extOutName="RtdOut" />
    <fvSubnet ip="1.1.1.1/16" />
    <fvSubnet ip="1.2.1.1/16" />
    <fvSubnet ip="40.1.1.1/24" scope="public" />
    <fvRsCtx tnFvCtxName="inb" />
  </fvBD>
  <fvCtx name="inb" />

  <l3extOut name="RtdOut">
    <l3extRsL3DomAtt tDn="uni/l3dom-extdom"/>
    <l3extInstP name="extMgmt">
      </l3extInstP>
    <l3extLNodeP name="borderLeaf">
      <l3extRsNodeL3OutAtt tDn="topology/pod-1/node-101" rtrId="10.10.10.10"/>
      <l3extRsNodeL3OutAtt tDn="topology/pod-1/node-102" rtrId="10.10.10.11"/>
      <l3extLIIfP name='portProfile'>
        <l3extRsPathL3OutAtt tDn="topology/pod-1/paths-101/pathep-[eth1/40]"
ifInstT='l3-port' addr="192.168.62.1/24"/>
        <l3extRsPathL3OutAtt tDn="topology/pod-1/paths-102/pathep-[eth1/40]"
ifInstT='l3-port' addr="192.168.62.5/24"/>
        <ospfIfP/>
      </l3extLIIfP>
    </l3extLNodeP>
    <l3extRsEctx tnFvCtxName="inb"/>
    <ospfExtP areaId="57" />
  </l3extOut>
</fvTenant>
```

Deploying an Application Policy

Three-Tier Application Deployment

A filter specifies the data protocols to be allowed or denied by a contract that contains the filter. A contract can contain multiple subjects. A subject can be used to realize uni- or bidirectional filters. A unidirectional filter is a filter that is used in one direction, either from consumer-to-provider (IN) or from provider-to-consumer (OUT) filter. A bidirectional filter is the same filter that is used in both directions. It is not reflexive.

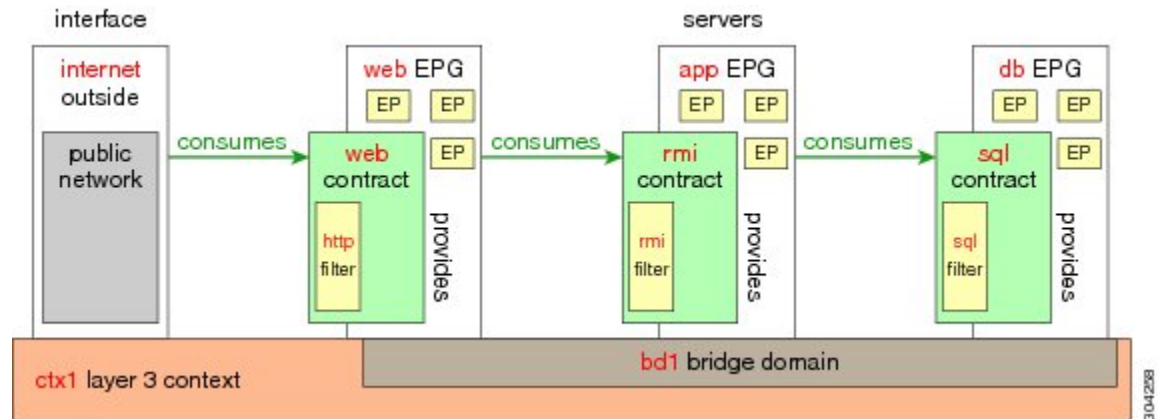
Contracts are policies that enable inter-End Point Group (inter-EPG) communication. These policies are the rules that specify communication between application tiers. If no contract is attached to the EPG, inter-EPG communication is disabled by default. No contract is required for intra-EPG communication because intra-EPG communication is always allowed.

Application profiles enable you to model application requirements that the APIC then automatically renders in the network and data center infrastructure. The application profiles enable administrators to approach the resource pool in terms of applications rather than infrastructure building blocks. The application profile is a container that holds EPGs that are logically related to one another. EPGs can communicate with other EPGs in the same application profile and with EPGs in other application profiles.

To deploy an application policy, you must create the required application profiles, filters, and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

The web server has the HTTP filter, the application server has the Remote Method Invocation (RMI) filter, and the database server has the Structured Query Language (SQL) filter. The application server consumes the SQL contract to communicate with the database server. The web server consumes the RMI contract to communicate with the application server. The traffic enters from the web server and communicates with the application server. The application server then communicates with the database server, and the traffic can also communicate externally.

Figure 14: Three-Tier Application Diagram



Parameters to Create a Filter for http

The parameters to create a filter for http in this example is as follows:

Parameter Name	Filter for http
Name	http
Number of Entries	2
Entry Name	Dport-80 Dport-443
Ethertype	IP
Protocol	tcp tcp
Destination Port	http https

Parameters to Create Filters for rmi and sql

The parameters to create filters for rmi and sql in this example are as follows:

Parameter Name	Filter for rmi	Filter for sql
Name	rmi	sql
Number of Entries	1	1
Entry Name	Dport-1099	Dport-1521
Ethertype	IP	IP
Protocol	tcp	tcp
Destination Port	1099	1521

Example Application Profile Database

The application profile database in this example is as follows:

EPG	Provided Contracts	Consumed Contracts
web	web	rmi
app	rmi	sql
db	sql	--

Deploying an Application Profile Using the REST API

The port the EPG uses must belong to one of the VM Managers (VMM) or physical domains associated with the EPG.

Procedure

Step 1 Send this HTTP POST message to deploy the application using the XML API.

Example:

```
POST https://apic-ip-address/api/mo/uni/tn-ExampleCorp.xml
```

Step 2 Include this XML structure in the body of the POST message.

Example:

```

<fvTenant name="ExampleCorp">
  <fvAp name="OnlineStore">
    <fvAEPg name="web">
      <fvRsBd tnFvBDName="bd1"/>
      <fvRsCons tnVzBrCPName="rmi"/>
      <fvRsProv tnVzBrCPName="web"/>
      <fvRsDomAtt tDn="uni/vmmp-VMware/dom-datacenter"delimiter=@/>
    </fvAEPg>
    <fvAEPg name="db">
      <fvRsBd tnFvBDName="bd1"/>
      <fvRsProv tnVzBrCPName="sql"/>
      <fvRsDomAtt tDn="uni/vmmp-VMware/dom-datacenter"/>
    </fvAEPg>
    <fvAEPg name="app">
      <fvRsBd tnFvBDName="bd1"/>
      <fvRsProv tnVzBrCPName="rmi"/>
      <fvRsCons tnVzBrCPName="sql"/>
      <fvRsDomAtt tDn="uni/vmmp-VMware/dom-datacenter"/>
    </fvAEPg>
  </fvAp>
  <vzFilter name="http" >
  <vzEntry dFromPort="80" name="DPort-80" prot="tcp" etherT="ip"/>
  <vzEntry dFromPort="443" name="DPort-443" prot="tcp" etherT="ip"/>
  </vzFilter>
  <vzFilter name="rmi" >
  <vzEntry dFromPort="1099" name="DPort-1099" prot="tcp" etherT="ip"/>
  </vzFilter>
  <vzFilter name="sql">
  <vzEntry dFromPort="1521" name="DPort-1521" prot="tcp" etherT="ip"/>
  </vzFilter>
  <vzBrCP name="web">
    <vzSubj name="web">
      <vzRsSubjFiltAtt tnVzFilterName="http"/>
    </vzSubj>
  </vzBrCP>
  <vzBrCP name="rmi">
    <vzSubj name="rmi">
      <vzRsSubjFiltAtt tnVzFilterName="rmi"/>
    </vzSubj>
  </vzBrCP>
  <vzBrCP name="sql">
    <vzSubj name="sql">
      <vzRsSubjFiltAtt tnVzFilterName="sql"/>
    </vzSubj>
  </vzBrCP>
</fvTenant>

```

In the string `fvRsDomAtt tDn="uni/vmmp-VMware/dom-datacenter"delimiter=@/`, **delimiter=@** is optional. If you do not enter a delimiter, the system will use the default | delimiter.

In the XML structure, the first line modifies, or creates if necessary, the tenant named ExampleCorp.

```

<fvTenant name="ExampleCorp">

```

This line creates an application network profile named OnlineStore.

```
<fvAp name="OnlineStore">
```

The elements within the application network profile create three endpoint groups, one for each of the three servers. The following lines create an endpoint group named web and associate it with an existing bridge domain named bd1. This endpoint group is a consumer, or destination, of the traffic allowed by the binary contract named rmi and is a provider, or source, of the traffic allowed by the binary contract named web. The endpoint group is associated with the VMM domain named datacenter.

```
<fvAEPg name="web">
  <fvRsBd tnFvBDName="bd1"/>
  <fvRsCons tnVzBrCPName="rmi"/>
  <fvRsProv tnVzBrCPName="web"/>
  <fvRsDomAtt tDn="uni/vmmp-VMware/dom-datacenter"/>
</fvAEPg>
```

The remaining two endpoint groups, for the application server and the database server, are created in a similar way.

The following lines define a traffic filter named http that specifies TCP traffic of types HTTP (port 80) and HTTPS (port 443).

```
<vzFilter name="http" >
<vzEntry dFromPort="80" name="DPort-80" prot="tcp" etherT="ip"/>
<vzEntry dFromPort="443" name="DPort-443" prot="tcp" etherT="ip"/>
</vzFilter>
```

The remaining two filters, for application data and database (sql) data, are created in a similar way.

The following lines create a binary contract named web that incorporates the filter named http:

```
<vzBrCP name="web">
  <vzSubj name="web">
    <vzRsSubjFiltAtt tnVzFilterName="http"/>
  </vzSubj>
</vzBrCP>
```

The remaining two contracts, for rmi and sql data protocols, are created in a similar way.

The final line closes the structure:

```
</fvTenant>
```