Cisco UCS Director APIC Management Guide, Release 6.7

First Published: 2019-01-09

Last Modified: 2020-06-18

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

CONTENTS

Preface

Audience

This guide is intended primarily for data center administrators who use Cisco UCS Director and who have responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security
- Virtualization and virtual machines

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in **this font**. Main titles such as window, dialog box, and wizard titles appear in **this font**.
Document titles	Document titles appear in *this font*.
TUI elements	In a Text-based User Interface, text the system displays appears in `this font`.
System output	Terminal sessions and information that the system displays appear in `this font`.

Text Type	Indication
CLI commands	CLI command keywords appear in **this font**. Variables in a CLI command appear in *this font*.
[ ]	Elements in square brackets are optional.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
< >	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.

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

Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

Tip Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.

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

Cisco UCS Director Documentation Roadmap

For a complete list of Cisco UCS Director documentation, see the *Cisco UCS Director Documentation Roadmap* available at the following URL: http://www.cisco.com/en/US/docs/unified_computing/ucs/ucs-director/doc-roadmap/b_UCSDirectorDocRoadmap.html.

Cisco UCS Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/b-series-doc.

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/c-series-doc.

Note The *Cisco UCS B-Series Servers Documentation Roadmap* includes links to documentation for Cisco UCS Manager and Cisco UCS Central. The *Cisco UCS C-Series Servers Documentation Roadmap* includes links to documentation for Cisco Integrated Management Controller.

Documentation Feedback

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

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

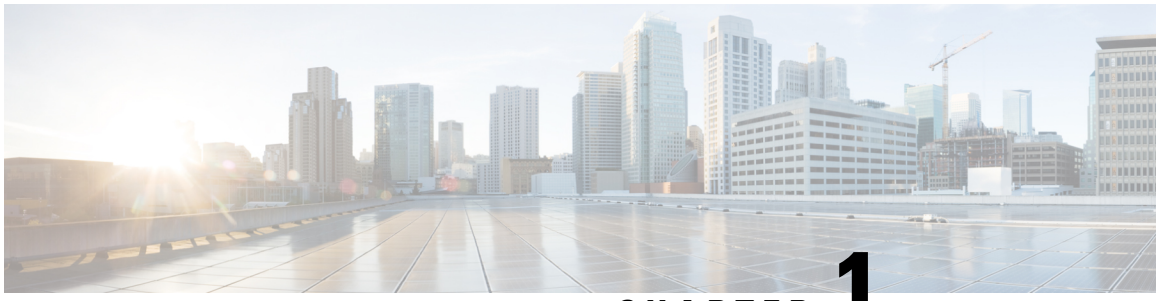

CHAPTER 1

New and Changed Information for this Release

New and Changed Information for this Release

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to this guide or of all new features in this release.

Table 1: New and Changed Information in Cisco UCS Director Connector Pack Release 6.7(4.1)

Feature	Description	Where Documented
Introduction of Interface Policies	Provision to add the following interface policies at fabric level: • Fibre Channel Interface Policy • L2 Interface Policy • Spanning Tree Interface Policy	Viewing APIC Resources, on page 16
Introduction of Port Policies	Provision to add the following port policies at fabric level: • Port Security Policy • Port Channel Member Policy	Viewing APIC Resources, on page 16
Enhancements to MACsec	Provision to define and associate MACsec KeyChain policy and MACsec access parameters policy to MACsec interface policy. MACsec KeyChain policy consists of configuration specific to keychain definition, and MACsec access parameters policy consists of configuration related to MACsec functionality.	Viewing APIC Resources, on page 16

Feature	Description	Where Documented
Support for Virtual Switched Port Analyzer (VSPAN)	Provision to start or stop VSPAN sessions on demand to copy relevant traffic from a virtual switch to a destination group. Also, you can perform the following tssks in a VSPAN session: • Associate a desired destination group to the session • Add an EPG or a client end point as source • Associate a source path to VSPAN source	Viewing APIC Resources, on page 16
Support for Flow Record	Provides support for define NetFlow record at fabric and tenant levels.	Viewing APIC Resources, on page 16 Flow Record, on page 102
Introduction of NetFlow Monitor Policy	Provision to create a NetFlow monitor policy at fabric level and associate it with a flow record. You can also perform the following tasks: • Create an external collector reachability (also known as NetFlow exporter) • Associate a NetFlow exporter with a fabric NetFlow monitor policy • Deploy the NetFlow monitor policy on an existing bridge domain by associating the NetFlow monitor policy with bridge domain	Viewing APIC Resources, on page 16 NetFlow Monitor Policy, on page 78
Support for Flow Control Policy	Provision to create Flow Control Policy at fabric level.	Viewing APIC Resources, on page 16
Support for Slow Drain Policy	Provision to create a slow drain policy at fabric level for handling FCoE packets that are causing traffic congestion on ACI Fabric.	Viewing APIC Resources, on page 16

Feature	Description	Where Documented
Support for Data Plane Policing	Provision to create a data plane policing (DPP) policy at fabric level to manage bandwidth consumption on ACI fabric access interfaces. DPP policies can apply to egress traffic, ingress traffic, or both.	Viewing APIC Resources, on page 16
Enhancements to the Interface Policy	Provision to create alias for the following interface policies: • CDP Interface Policy • Link Level Policy • LLDP Interface Policy • MACsec Access Interface Policy • Port Channel Member Policy • Port Channel Policy • Spanning Tree Interface Policy • Storm Control Policy	Viewing APIC Resources, on page 16
Support for APIC Monitoring Policy	Provides support to define a monitoring policy as a default policy to be applied to all the tenants in an APIC account to monitor EPGs, application profiles, services, and so on.	Viewing APIC Resources, on page 16
Enhancements to VRF	Provides support to create BGP context per address family, OSPF context per address family, SNMP context, community profile, and BGP route target profile to VRFs. Also, provides support to add BGP route target to the BGP route target profile.	• Adding a BGP Context per Address Family, on page 45 • Adding a OSPF Context per Address Family, on page 46 • Adding a SNMP Context, on page 46 • Creating a Community Profile, on page 46 • Adding a BGP Route Target Profile, on page 44 • Adding a BGP Route Target to BGP Route Target Profile, on page 45

New and Changed Information for this Release

Feature	Description	Where Documented
Introduction of BGP Timers Policy	Provision to define a BGP timers policy at tenant level.	BGP Timers, on page 47
Enhancements to First-Hop Security (FHS)	Provides support for associating an FHS policy to a tenant while adding a bridge domain to VRF.	Creating a FHS Trust Policy, on page 76
Support for DHCP Policy	Provides support to add DHCP option policy to a tenant.	Creating a DHCP Option Policy, on page 94
Enhancements to Enhanced Interior Gateway Routing Protocol (EIGRP)	Provision to define an eigrpCtxAfPol policy under tenant protocol policies and apply the policy to one or more VRFs under the tenant.	EIGRP Address Family Context Policy, on page 96
Support for Snoop Policy	Provides support to define Internet Group Management Protocol (IGMP) Snoop policy and Multicast Listener Discovery (MLD) Snoop policy at tenant level.	IGMP Snoop Policy, on page 98 MLD Snoop Policy, on page 99

Table 2: New and Changed Information in Cisco UCS Director Release 6.7(4.0)

Feature	Description	Where Documented
Support for PC/vPC Leaf Policy	Provision to create a port channel (PC) and virtual port channel (vPC) leaf policy and associate it with: • Netflow monitor policy • Virtual destination groups • Virtual source groups • Override policy group	Viewing APIC Resources, on page 16
Support for Access Port Selector	Provision to add an access port selector to a fabric interface profile.	Viewing APIC Resources, on page 16
Support for VMM Domain	Provides support for creating a virtual machine manager (VMM) domain to integrate APIC with a third-party VMM (for example, VMware vCenter) to extend the benefits of ACI to the virtualized infrastructure.	Viewing APIC Resources, on page 16

Feature	Description	Where Documented
Support for Creating VRF in APIC	Provision to define IPv4 unicast address family or IPv6 unicast address family as the EIGRP address family type, to configure an EIGRP routing instance.	• Creating a VRF, on page 43 • Adding an EIGRP to the VRF, on page 44
Enhancements to add Domain to an EPG	Provision to configure a default port binding type for all new vEthernet port profiles.	Adding a Domain to an EPG, on page 56
Support for Route Tag Policy	Provides support for creating a route tag policy with a tag value which is used to prevent routing loops.	Creating a Route Tag Policy, on page 96

Table 3: New and Changed Information in Cisco UCS Director Release 6.7

Feature	Description	Where Documented
Enhancements to tenant management	Provision to define globally unique identifier (GUID) for SCVMM provider, and define an alias name for the tenant. While the tenant name cannot be changed after creation, the alias name of the tenant can be changed as required.	• Creating a Tenant, on page 40 • Adding a GUID to a Tenant, on page 41
Support for Neighbor Discovery Router Advertisement (ND RA) prefixes	You can create ND RA prefixes for Layer 3 interfaces.	• Creating an ND RA Prefix Policy, on page 52
Extension of support for APIC account	You can perform the following tasks in APIC account: • Adding an EPG • Adding a Domain to an EPG • Adding a Static Node to EPG • Adding a Static Path to EPG	• Adding an EPG, on page 55 • Adding a Domain to an EPG, on page 56 • Adding a Static Path to EPG, on page 58 • Adding a Static Node to EPG, on page 60
Enhancements to EPG	Extended the support for EPG to define data plane policy, forwarding control, preferred group member, flood on encapsulation, and FHS trust control policy during creation of EPG	• Adding a Domain to an EPG, on page 56 • Adding a Static Node to EPG, on page 60 • Adding a Static Path to EPG, on page 58

Feature	Description	Where Documented
Support for EPG Contract Master	Provision to define an EPG as a contract master for another EPG in the same tenant. To streamline associating contracts to new EPGs, you can enable EPG to inherit all the (provided and consumed) contracts from master EPG.	• Adding an EPG Contract Master, on page 61
Enhancements to APIC Contracts	Cisco UCS Director introduces fields to define alias name, DSCP target, and tag for a contract subject during creation. When a contract is applied to both inbound and outbound traffic while creating a contract subject, the user gets the additional fields to define the service graph, QoS priority, and target DSCP for the in term and out term properties. If the selected contract does not apply to both directions, then the filter chain must be configured for consumer to provider and provider to consumer separately. Cisco UCS Director has the provision to define the filter chain for consumer to provider and provider to consumer.	• Creating a Contract Subject, on page 64 • Adding a Consumed Label to a Contract Subject, on page 68 • Adding a Provided Label to a Contract Subject, on page 69 • Adding a Filter Chain, on page 73 • Adding a Filter Chain for Consumer to Provider, on page 74 • Adding a Filter Chain for Provider to Consumer, on page 74
Support for data plane policing (DPP)	You can use DPP to manage bandwidth consumption on ACI fabric access interfaces.	• Creating a Data Plane Policing, on page 75
Support for First-Hop Security (FHS) feature	You can use FHS feature to achieve a better IPv4 and IPv6 link security and management over the layer 2 links.	• Creating a FHS Trust Policy, on page 76

Feature	Description	Where Documented
Enhancements to Routed Outside	To support protocol and QoS in an external routed network, this release introduces additional fields in the following actions: • Create a routed outside • Add a route map or profile to an external routed network • Add a logical node profile to an external routed network • Add a logical node to a logical node profile of an external routed network • Add a static route to a logical node • Add an external network to an external routed network	• Creating a Routed Outside, on page 83 • Adding a Route Map or Profile to an External Routed Network, on page 84 • Adding a Logical Node Profile to an External Routed Network, on page 85 • Adding a Logical Node to a Logical Node Profile of an External Routed Network, on page 85 • Adding a Static Route to a Logical Node, on page 86 • Adding an External Network to an External Routed Network, on page 86
Extension of support for APIC L3out tasks changes	New fields have been added to the following tasks to extend the support of L3out in APIC account: • Adding an external routed network in APIC account • Adding a logical node profile to external routed network • Adding an external network to APIC external routed network • Adding a static route to a logical node in APIC account • Adding a routed profile to an external routed network • Adding a logical node to a logical node profile of an external routed network	• Adding a Logical Node Profile to an External Routed Network, on page 85 • Adding an External Network to an External Routed Network, on page 86 • Adding a Static Route to a Logical Node, on page 86 • Adding a Route Map or Profile to an External Routed Network, on page 84 • Adding a Logical Node to a Logical Node Profile of an External Routed Network, on page 85
Introduction of Logical NetFlow Monitoring Policy	Provision to deploy and enable NetFlow policies on a per-interface basis, depending on the traffic-type or address family to be monitored (IPv4, IPv6, or Layer 2 (CE type)).	• Adding a Logical NetFlow Monitoring Policy , on page 89

Feature	Description	Where Documented
Support for IGMP interface policy and route map	Provision to add an IGMP interface policy and create route map policy for route redistribution or policy-based routing.	• Adding an IGMP Interface Policy, on page 95 • Adding a Route Map Entry, on page 105
Support for route control context	Provision to define match action rules and set action rules for a route map. Also, you can create an action rule profile which is used to define the route-map set clauses for the L3out.	• Creating a Match Rule for a Route Map, on page 103 • Adding a Match Regex Community Term to a Route Map Match Rule, on page 103 • Adding a Match Prefix to a Match Rule, on page 103 • Adding a Match Community Term to a Route Map Match Rule, on page 104 • Adding a Match Community Factor to a Match Community Term, on page 104 • Creating a Set Rules for Route Map, on page 106 • Adding an Additional Community, on page 107 • Adding a Set AS Path to the Action Rule Profile, on page 107 • Adding an AS Number to Prepend the AS Path, on page 108 • Associating a Match Rule to a Route Control Context, on page 109
Support for static route and route control profile	Provision to add a next hop address to a static route and to add a route control profile to a subnet and external network.	• Adding a Next Hop Address to a Static Route, on page 87 • Adding a Route Control Profile to a Subnet, on page 88 • Adding a Route Control Profile to an External Network, on page 87

Feature	Description	Where Documented
Support for vzAny	Provision to define labels that determine which EPG consumers and EPG providers can communicate with one another. Label matching determines which subjects of a contract are used with a given EPG provider or EPG consumer of that contract.	• vzAny , on page 170 • Creating a vzAny Contract Interface, on page 171 • Creating APIC vzAny Consumed Subject Label to VRF, on page 174 • Creating APIC vzAny Provided Subject Label to VRF, on page 173 • Creating a vzAny Consumed Contract, on page 171 • Creating a vzAny Provided Contract, on page 170 • Creating a vzAny EPG Consumed Any Labels, on page 172 • Creating a vzAny EPG Provided Any Labels, on page 173

New and Changed Information for this Release

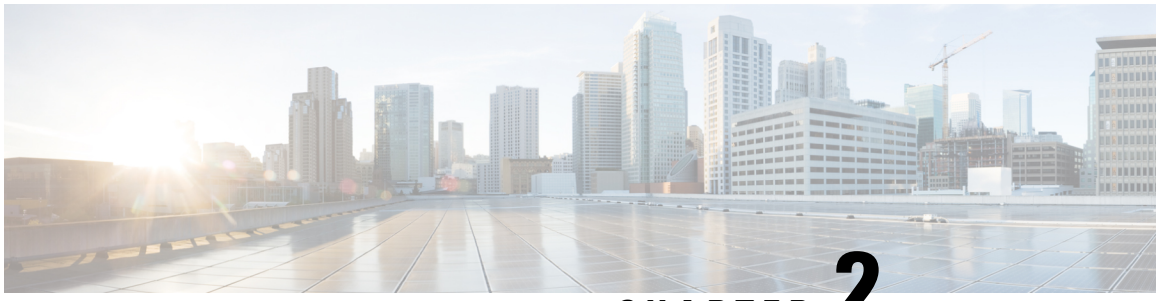

CHAPTER 2

Overview

Cisco UCS Director and Cisco Application Centric Infrastructure

Cisco UCS Director is a unified infrastructure management solution that provides management from a single interface for compute, network, storage, and virtualization layers. Cisco UCS Director uses a workflow orchestration engine with workflow tasks that support the compute, network, storage, and virtualization layers. Cisco UCS Director supports multitenancy, which enables policy-based and shared use of the infrastructure.

Cisco UCS Director also supports the ability to define contracts between different container tiers, enabling you to apply rules between tiers.

Cisco Application Centric Infrastructure (ACI) allows application requirements to define the network. This architecture simplifies, optimizes, and accelerates the entire application deployment cycle.

The combination of Cisco UCS Director and Cisco ACI enables automatic provisioning and delivery of an application-centric infrastructure.

Note To use ACI 1.1(1*), ensure that TLSv1 is enabled in Cisco Application Policy Infrastructure Controller (APIC). In APIC, choose **Fabric > Fabric Resources > Pod Polices > Communication > Default** and enable **TLSv1**.

Cisco Application Policy Infrastructure Controller

The Cisco Application Policy Infrastructure Controller (APIC) is the unified point of automation, management, monitoring, and programmability for the Cisco Application Centric Infrastructure (ACI). The APIC supports the deployment, management, and monitoring of any application anywhere, with a unified operations model for physical and virtual components of the infrastructure. It is the central control engine for the broader cloud network. The APIC programmatically automates network provisioning and control-based on user-defined application requirements and policies.

The Cisco UCS Director orchestration feature allows you to automate APIC configuration and management tasks through operational workflows. A complete list of the APIC orchestration tasks is available in the

Workflow Designer, and in the Task Library. For more information about orchestration in Cisco UCS Director, see the Cisco UCS Director Orchestration Guide.

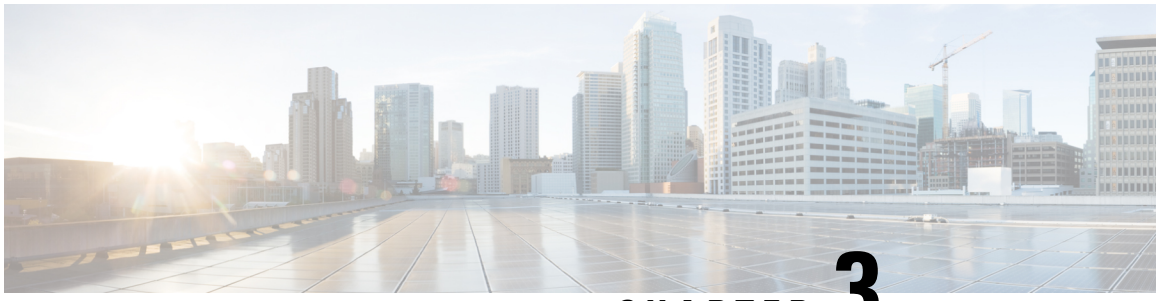

CHAPTER 3

Configuring APIC Accounts

Guidelines for APIC Accounts

Before you create an APIC account in Cisco UCS Director, consider the following guidelines and best practices.

Account Permissions on Cisco APIC

The Cisco APIC account username and password that you provide when you add the APIC account to Cisco UCS Director must have all the Cisco APIC privileges required to do the following:

- Access the supported features in Cisco APIC
- Perform actions in Cisco APIC, such as viewing and accessing reports
- Execute workflow tasks in Cisco UCS Director

Account Authentication on Cisco APIC

The Cisco APIC account username and password that you use for the APIC account in Cisco UCS Director is authenticated by Cisco APIC not by Cisco UCS Director. As a result, you can use one of the following types of accounts:

- Local authentication through a Cisco APIC account
- Remote authentication by Cisco APIC through one of the following:
  - LDAP
  - RADIUS
  - TACACS+

If you use a Cisco APIC account with remote authentication, enter the username on the **Add Account** screen in the following format: **apic:<Domain Name>\<Remote User Name>**

Note Cisco UCS Director does not support authentication through RADIUS or TACACS+ for accounts used to log in to Cisco UCS Director. Support is only available for authentication of accounts that Cisco UCS Director uses to log in to Cisco APIC.

APIC Clusters

Each APIC account in Cisco UCS Director represents an APIC cluster. When you add an APIC cluster to an APIC account, Cisco UCS Director automatically discovers the controllers in that cluster.

To view details of the controllers, choose **Physical** > **Network**, choose the APIC account, and then click **View Details**.

ACI Fabric Integration

To integrate Cisco UCS Director with the ACI fabric, ensure that TLSv1 is enabled on the ACI fabric.

You must enable TLSv1 in Cisco APIC, as follows: **Fabric Policies** > **Pod Policies** > **Policies - Communication**.

APIC Accounts and Pods

Cisco APIC accounts are multi-domain manager accounts that are not tied to a specific pod. You can assign the account to a pod, but that is optional.

APIC Accounts and Resource Groups

If you add an APIC account to a resource group and that account is associated with a pod, you cannot edit the pod.

You cannot delete an account that is part of a resource group.

Support for In-Band and Out-of-Band Management

Cisco UCS Director supports in-band and out-of-band management of Cisco ACI. You can add a Cisco APIC account to Cisco UCS Director in the following scenarios:

- Out-of-Band—An out-of-band IP address is configured and the Cisco UCS Director VM is in a domain that is not managed by Cisco APIC.
- In-Band—An in-band IP address is configured and reachable, no out-of-band IP address is configured, and the Cisco UCS Director VM is in a domain managed by Cisco APIC.

Adding an APIC Account

Before you begin

Review the guidelines and best practices in Guidelines for APIC Accounts, on page 13.

Step 1 Choose **Administration** > **Physical Accounts**.

Step 2 On the **Physical Accounts** page, click **Multi-Domain Managers**.

Step 3 Click **Add**.

Step 4 On the **Add Account** screen, choose **APIC** from the **Account Type** drop-down list and click **Submit**.

Step 5 On the **Add Account** screen, complete the fields, including the following:

a) Enter a unique account name and description.

b) From the **Pod** drop-down list, choose the pod where you want to add the APIC account.

c) In the **Server IP** field, enter the IP address of one of the APIC controllers in the APIC cluster.

Cisco UCS Director automatically discovers the IP address of the other APIC controllers in the APIC cluster.

If the IP address of the APIC controller is not reachable, Cisco UCS Director relies on the Out-of-Band IP address of another APIC controllers for managing Cisco APIC.

d) Check the **Use Credential Policy** box if you want to use a credential policy for this account rather than enter the username and password information manually.

e) If you checked the **Use Credential Policy** box, choose a policy from the **Credential Policy** drop-down list.

The APIC account in the credential policy must meet the criteria listed in Guidelines for APIC Accounts, on page 13.

Note You can only connect to Cisco APIC with HTTPS protocol. You cannot connect through SSH or Telnet protocol. If the credential policy specifies SSH or Telnet protocol, you are prompted to check the protocol defined in the credential policy.

f) If you did not check **Use Credential Policy**, enter the username and password that this account uses to access Cisco APIC.

This username must be a valid account with the required privileges in Cisco APIC. The account must also meet the criteria listed in Guidelines for APIC Accounts, on page 13.

Note For an account with remote authentication by Cisco APIC through LDAP, RADIUS, or TACACS+, enter the username in the following format: **`apic:<Domain Name>\<Remote User Name>`**.

g) If you did not check **Use Credential Policy**, do the following:

- From the **Protocol** drop-down list, choose **https**.
- In the **Port** field, enter the port used to access the APIC account. The default port is 443.

h) Enter the email address and location of the administrator or other person responsible for this account.

Step 6 Click **Submit**.

Cisco UCS Director tests the connection to the APIC server. If that test is successful, it adds the APIC account and discovers all controllers and other infrastructure elements in the APIC server. This discovery process and inventory collection takes a few minutes to complete.

Viewing APIC Resources

After creating an APIC account in Cisco UCS Director, you can view related resources of the APIC account.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click one of the following tabs to view the details of a specific component in the server:

- **Summary** tab—Displays the system overview and summary of the APIC controller.
- **Fabric Nodes** tab—Displays the list of fabric nodes with their details such as the node name, model, vendor, role, serial, and node ID with the status.

  To view more details about fabric nodes, choose a fabric node and click **View Details**. The following tabs appear:

  - **Fabric Chassis**—Displays the fabric name, ID, model, vendor, serial, revision, and operation status of the fabric chassis.
  - **Fan Slots**—Displays the fabric name, slot ID, type, operation status, and inserted-card details of the fan slots.
  - **Physical Interfaces**—Displays the interface details that include the speed, mode, CFG access VLAN, CFG native VLAN, bundle index, operational duplex mode, operational port state, and reason for the current operation state. The operational state of the port can be one of the following: Unknown, Down, Link-up, and Up.
  - **Fabric Routed Vlan Interfaces**—Displays the status and reason for the current operation status of the fabric-routed VLAN interfaces.
  - **Fabric Encapsulated Routed Interfaces**—Displays a list of the fabric-encapsulated routed interfaces.
  - **Fabric Routed Loopback Interfaces**—Displays a list of the fabric-routed loopback interfaces.
  - **Fabric Management Interfaces**—Displays a list of the fabric management interfaces.
  - **Tunnel Interfaces**—Displays the interface, operation state, reason for the current operation state, tunnel layer, tunnel type, and type of the tunnel interface.

- **System** tab—Displays the system details that include the node name, in-band management IP address, out-of-band management IP address, infrastructure IP address, fabric MAC address, ID, role, and serial number.
- **Fabric Memberships** tab—Displays the fabric membership details that include the node name, serial number, node ID, model, role, IP address, decommissioned status, and supported model.
- **Physical Domains** tab—Displays the physical domains in the APIC server. Click **Add** to add a domain.
- **Tenants Health** tab—Displays the health score of tenants.

  To view more details about a tenant's health, choose a tenant and click **View Details**. The following tabs appear:

  - **EPGs Health**—Displays the health score of endpoint groups (EPGs).

- **Application Health**—Displays the health score of applications.

- **Nodes Health** tab—Displays the health score of nodes.

  To view more details about the health of the nodes, choose a node and click **View Details**. The following tabs appear:

  - **Access Ports Health**—Displays the health score of access ports.
  - **Fabric Ports Health**—Displays the health score of fabric ports.
  - **Line Cards Health**—Displays the health score of line cards.

- **Access Entity Profile** tab—Displays the names and descriptions of the access entity profiles.

  To view more details about the access entity profile, choose an entity profile and click **View Details**. The following tabs appear:

  - **Policy Groups**—Displays the policy groups of an entity profile.
  - **Domain Associated To Interfaces**—Displays a list of domains that are associated with the interfaces.

- **Link Level Policy** tab—Displays the name, automatic negotiation, speed, link debounce interval, and description of the link level policy.

  To view more details about a link level policy, choose a policy and click **View Details**. The following tab appears:

  - **Link Level Policy Alias**—Displays the names of link level policy and alias.

- **VLAN Pool** tab—Displays the VLAN pools that are added in the APIC server. Click **Add** to add a VLAN pool.

  To view more details about a VLAN pool, choose a VLAN pool and click **View Details**. The following tab appears:

  - **VLAN Pool Range**—Displays the VLAN pool name, mode of allocation, and the pool range. Click **Add** to add a VLAN range to the VLAN pool.

- **FEX Profile** tab—Displays the FEX profiles that are added in the APIC server. Click **Add** to add a FEX profile.

  To view more details about a FEX profile, choose a FEX profile and click **View Details**. The following tab appears:

  - **FEX Profile Access Port Selectors**—Displays the access port selectors of the FEX profile. Click **Add** to add an access port selector to the FEX profile. Choose an access port selector and click **View Details** to view the access port blocks and sub port blocks of the access port selector.

- **Fabric Spine Interface Profiles** tab—Displays the name and description of the fabric spine interface profiles. To add a fabric spine interface profile, click **Add** and enter a unique name and short description for the fabric spine interface profile.

- **CDP Interface Policy** tab—Displays the name and description of the Cisco Discovery Protocol (CDP) interface policy, with the administration status.

  To view more details about a CDP interface policy, choose a policy and click **View Details**. The following tab appears:

  - **CDP Interface Policy Alias**—Displays the names of CDP interface policy and alias.

- **CoPP Spine Policy** tab—Displays the name, type of profile, and description of the Cisco Control Plane Policing (CoPP) spine policy. Click **Create** to add a CoPP spine policy. To customize the existing APIC CoPP spine policy

or create a customized APIC CoPP spine policy, click **Configure Custom Values for APIC CoPP Spine Policy** and specify the protocol with rate and burst details.

To view more details about a CoPP spine policy, choose a CoPP spine policy and click **View Details**. The following tab appears:

- **CoPP Spine Custom Values**—Displays the customized APIC CoPP spine policy details.

- **Port Security Policy** tab—Displays the port security policies defined for the APIC account. To create a port security policy, click **Create** and complete the following fields:
  - Enter a unique name and short description for the port security policy.
  - In the **Port Security Timeout** field, choose a specific timeout value after which MAC learning can be re-enabled on an interface.You can choose a value between 60 and 3600 seconds as a timeout value.
  - In the **Maximum Endpoints** field, choose a maximum number of endpoints that can be learned on an interface. You can choose a value between 0 and 12000. If the maximum endpoints value is set as zero, the port security policy is disabled on that port.
  - The **Violation Action** field displays protect. The violation action of the port security policy is always set as the protect mode.

    In the protect mode, MAC learning is disabled and MAC addresses are not added to the Content Addressable Memory (CAM) table. MAC learning is re-enabled after the configured timeout value.

- **VSPAN Sessions** tab—Displays the Virtual Switched Port Analyzer (VSPAN) session details. To create a VSPAN Session, click **Create** and complete the following fields:
  - Enter a unique name and short description for the VSPAN session.
  - From the **Admin State** drop-down list, choose **Start** or **Stop**. By default, **Start** is chosen as the admin state.
  - Click **Select** and choose a VSPAN destination group that needs to be used for the VSPAN sessions.

  Choose a VSPAN session and click **View Details** to view the following tabs:
  - **Destination Group**—Displays the destination group of the VSPAN session.
  - **Sources**—Displays the sources of the VSPAN session. To add a source to the VSAPN session, click **Add** and complete the following fields:
    - Enter a unique name and description for the source.
    - From the **Direction** drop-down list, choose both, incoming, or outcoming as the direction of the source.
    - From the **Source Type** drop-down list, choose **EPG** or **CEP** as the source type.
    - Click **Select** and choose a tenant to be associated with the VSPAN session source.
    - Click **Select** and choose an EPG to be associated with the VSPAN session source.

    To view more details of a source, choose a **Source** and click **View Details**. The following tabs appear:
    - **Source EPG**—Displays the EPG of the VSPAN session source.
    - **Source CEP**—Displays the CEP of the VSPAN session source.

- **Source Path**—Displays the source path associated with the VSPAN session. To associate a path to a VSPAN session source, click **Add**. In the **Associate Source Path to VSPAN Source** window, choose a port type and a static path.

- **VSPAN Destination Groups** tab—Displays the name and description of the VSPAN destination groups. To create a VSPAN destination group, click **Create** and enter a unique name and description for the VSPAN destination group.

  Choose a VSPAN destination group and click **View Details** to view the VSPAN destination group details. To add a destination to the VSPAN destination group, click **Add** and complete the following fields:

  - Enter a unique name and description for the destination.
  - From the **Destination Type** drop-down list, choose **ERSPAN** or **LSPAN** as the destination type.
  - In the **Destination IP** field, enter a valid IP address as the destination address. This field is displayed only when **ERSPAN** is chosen as the destination type.
  - In the **Flow ID** field, enter a number in the range of 1 to 1023 as the flow ID. The default value is 1. This field is displayed only when **ERSPAN** is chosen as the destination type.
  - In the **TTL** field, enter a number in the range of 1 to 225 as the TTL. The default value is 64. This field is displayed only when **ERSPAN** is chosen as the destination type.
  - In the **MTU** field, enter a number in the range of 64 to 9216 as the MTU value. The default value is 1518. This field is displayed only when **ERSPAN** is chosen as the destination type.
  - In the **DSCP** field, enter a number in the range of 0 to 64 as the DSCP value. This field is displayed only when **ERSPAN** is chosen as the destination type.
  - Click **Select** and choose a destination CEP that you want to use for the VSPAN destination group. This field is displayed only when **LSPAN** is chosen as the destination type.

  To view more details of a destination, choose a destination and click **View Details**. The following tabs appear:

  - **Destination ERSPAN Properties**—Displays the ERSPAN destination details of the VSPAN destination group.
  - **Destination LSPAN Properties**—Displays the LSPAN destination details of the VSPAN destination group.

- **Fibre Channel Interface Policy** tab—Displays the name, description, auto max speed, fill pattern, port mode, receive buffer credit, speed, and trunk mode of the fibre channel interface policy. To create a fibre channel interface policy, click **Create** and complete the following fields:

  - Enter a unique name and description for the fibre channel interface policy.
  - From the **Port Mode** drop-down list, choose **F** or **NP** as the port mode. The default value is **F**.
  - From the **Trunk Mode** drop-down list, choose **auto**, **trunk-off**, or **trunk-on** as the trunk mode. The default value is **trunk-off**.
  - From the **Speed** drop-down list, choose a speed for the fibre channel interface. The default value is **Auto**. On choosing **Auto**, the interface speed will be negotiated to automatically match the speed of the attached link.
  - From the **Auto Max Speed** drop-down list, choose a value as the automatic maximum interface speed. The default value is **32 Gbps**.
  - From the **Fill Pattern** drop-down list, choose **ARBFF** or **IDLE** as the fill pattern for the fibre channel interface. The default value is **IDLE**.

- In the **Receive Buffer Credit** field, enter a receive buffer credit in the range of 16 to 64 for the fibre channel interface. The default value is **64**.

- **Data Plane Policing** tab—Displays the name,burst size, excessive burst size, rate, and peak rate of the Data Plane Policing (DPP) policy. To add a DPP policy, click **Add** and complete the following fields:
  - Enter a unique name for the DPP.
  - Choose **enabled** from the drop-down list, to enable the administrative state of the DPP. The default value is disabled.
  - Choose **Bit Policer** or **Packet Policer** as the policer mode.
  - Choose one of the following as the traffic policer type:
    - **1 Rate 2 Color**—To rate-limit a traffic flow to an average bits-per-second arrival rate.
    - **2 Rate 3 Color**—To rate-limit a traffic flow to two rates and three traffic categories (green, yellow, and red).
  - Choose **Drop**, **Mark**, or **Transmit** as the action to be taken on categorizing a traffic flow as conforming (green).
  - In the **Conform mark cos** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the Class of Service (CoS) value for the traffic belonging to a specific class when the traffic flow is categorized as conforming.
  - In the **Conform mark dscp** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the differentiated services code point (DSCP) value for the traffic belonging to a specific class when traffic flow is categorized as conforming.
  - The **Exceed Action** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Choose **Drop**, **Mark**, or **Transmit** as the action to be taken when the traffic flow is exceeded.
  - The **Exceed mark cos** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the CoS value for the traffic belonging to a specific class when the traffic flow is exceeded.
  - The **Exceed mark dscp** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the DSCP value for the traffic belonging to a specific class when the traffic flow is exceeded.
  - Choose **Drop**, **Mark**, or **Transmit** as the action to be taken on categorizing a traffic flow as nonconforming (red).
  - In the **Violate mark cos** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the CoS value for the traffic belonging to a specific class on traffic violation.
  - In the **Violate mark dscp** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the DSCP value for the traffic belonging to a specific class on traffic violation.
  - Choose **Shared Policer** or **Dedicated Policer** as the policy mode. The default value is **Dedicated Policer**. The shared policer mode allows you to apply the same policing parameters to several interfaces simultaneously.
  - Enter the number of packets allowed at line rate during burst, as the burst size.
  - From the drop-down list, choose the unit at which the burst size has to be calculated.

- Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 5497555813760 to configure the excessive burst size.
- From the drop-down list, choose the unit at which the excessive burst size has to be calculated.
- Enter a number between 0 and 4398046510080 as an allowed rate. This is the committed rate at which the packets are allowed into the system (raw NTPD format).
- From the drop-down list, choose the unit at which the allowed rate has to be calculated.

- **CoPP Leaf Policy** tab—Displays the name, type of profile, and description of the Cisco Control Plane Policing (CoPP) leaf policy. Click **Create** to add a CoPP leaf policy. To customize the existing APIC CoPP leaf policy or create a customized APIC CoPP leaf policy, click **Configure Custom Values for APIC CoPP Leaf Policy** and specify the protocol with rate and burst details.

  To view more details about a CoPP leaf policy, choose a CoPP leaf policy and click **View Details**. The following tab appears:

  - **CoPP Leaf Custom Values**—Displays the customized APIC CoPP leaf policy details.

- **NetFlow Exporters** tab—Displays the name, description, source type, source IP address, destination port, destination IP address, QoS DSCP value, and version fromat of the NetFlow exporters. To add a NetFlow exporter, click **Add** and complete the following fields in the **Create External Collector Reachability** screen:
  - Enter a unique name and description for the external collector reachability.
  - From the **Source Type** drop-down list, choose the source type for the external collector reachability.
  - In the **Source IP Address** field, enter either the IPv4 or IPv6 address of the source with the subnet mask.
  - From the **Destination Port** drop-down list, choose the destination port for the external collector reachability.
  - In the **Destination IP Address** field, enter either the IPv4 or IPv6 address of the destination with the subnet mask.
  - From the **QoS DSCP Value** drop-down list, choose a QoS DSCP value. The default value is **None**.
  - From the **Netflow Explorer Version Format** drop-down list, choose a version from the list. The default value is **Cisco proprietary version 1**.
  - From the **EPG Type** drop-down list, choose an EPG type.

  To view more details about a NetFlow exporter, choose a NetFlow exporter and click **View Details**. The following tabs appears:

  - **Associated EPG**—Dsiplays the type and name of EPG associated with the external controller reachability.
  - **Associated VRF**—Dsiplays the name of VRF associated with the external controller reachability.

- **Port Channel Member Policies** tab—Displays the port channel member policies of the APIC accounts. To add a port channel member policy, click **Add** and complete the following fields:
  - Enter a unique name and description for the port channel member policy.
  - In the **Priority** field, enter any value in the range of 1 to 65535 as the priority for the port channel member policy. By default, 32768 is set as priority.
  - From the **Transmit Rate** drop-down list, choose **Normal** or **Fast** as the transmit rate.

To view more details about a port channel member policy, choose a policy and click **View Details**. The following tab appears:

- **Port Channel Member Policy Alias**—Displays the name and alias of the port channel member policies.

- **MACsec Interface Policies** tab—Displays the details of the Media Access Control Security (MACsec) interface policies of the APIC account. To create a MACsec interface policy, click **Add** and complete the following fields:
  - Enter a unique name and description for the MACsec interface policy.
  - From the **Admin State** drop-down list, choose **Disabled** to disable the admin state, if required. By default, the admin state is enabled.
  - Click **Select** and choose a MACsec parameter policy that you want to associate with the MACsec interface policy.
  - Click **Select** and choose a MACsec KeyChain policy that you want to associate with the MACsec interface policy.

  To view more details about a MACsec interface policy, choose a MACsec interface policy and click **View Details**. The following tabs appear:
  - **MACsec Parameter**—Displays the MACsec parameter policies that are associated with the MACsec interface policy.
  - **MACsec KeyChain**—Displays the MACsec KeyChain policies that are associated with the MACsec interface policy.
  - **MACsec Interface Policy Alias**—Displays the names of MACsec access interface policy and alias.

- **MACsec KeyChain Policies** tab—Displays the details of the MACsec KeyChain policies of the APIC account. To create a MACsec KeyChain policy, click **Create** and enter a unique name and description for the MACsec KeyChain policy.

  To view more details about a MACsec Keychain policy, choose a MACsec KeyChain policy and click **View Details**. The following tabs appear:
  - **MACsec KeyChain Policy Alias**—Displays the name and alias of the MACsec access interface policies.
  - **MACsec Key Policy**—Displays the defined MACsec key policy. To add a MACsec key policy, click **Add** and complete the following fields:
    - Enter a unique name and description for the MACsec key policy.
    - In the **Key Name** field, enter a key name. It allows only hexadecimal characters.
    - In the **Pre-shared Key** field, enter a pre-shared key (PSK) information.

      For 128-bit cipher suites, only 32 character PSKs are permitted.

      For 256-bit cipher suites, only 64 Character PSKs are permitted.
    - In the **Start Time** field, select a date and time from when the key is valid. By default, the current date and time is set as the start time.
    - Check the **Set End Time** check box to set the expiry date for the key. If this check box is left unchecked, the end time of the key is set as infinite.

- In the **End Time** field, select a date and time for the key to expire. This field appears only when the **Set End Time** check box is checked.

- **MACsec Access Parameters Policy** tab—Displays the details of the MACsec access parameters policies of the APIC account. To add a MACsec access parameters policy, click **Add** and complete the following fields:
  - Enter a unique name and description for the MACsec access parameters policy.
  - From the **Cipher Suite** drop-down list, choose a cipher suite that needs to be used to encrypt traffic on an Ethernet link secured with MACsec.
  - From the **Confidentiality Offset** drop-down list, choose one of the following options to configure a confidentiality offset for MACsec:
    - **Skip 0 bytes**—No octets are unencrypted. All traffic on the interface where the secure channel is applied is encrypted.
    - **Skip 30 bytes**—The first 30 octets of each Ethernet frame are unencrypted.
    - **Skip 50 bytes**—The first 30 octets of each Ethernet frame are unencrypted.
  - In the **Key Server Priority** field, enter the key server priority for the MACsec Key Agreement (MKA) server selection. By default, 16 is set as the key server priority. The switch with the lower priority number is selected as the key server.
  - In the **Window Size** field, enter the window size. By default, 64 is set as window size.
  - In the **Secure Association Key Expiry Time** field, enter the expiry time for Secure Association Key (SAK). The default value is disabled.
  - From the **Security Policy** drop-down list, choose one of the following options as the security policy:
    - **Must secure mode**—To allow encrypted traffic on the link.
    - **Should secure mode**—To allow both clear and encrypted traffic on the link.

- **Storm Control Policy** tab—Displays the details of the storm control policies of the APIC account. To create a storm control policy, click **Add** and complete the following fields:
  - Enter a unique name and description for the storm control policy.
  - From the **Configure Storm Control** drop-down list, choose **All Types** or **Unicast, Broadcast, Multicast**. Choosing the **Unicast, Broadcast, Multicast** option allows you to configure Storm Control on each traffic type separately.
  - From the **Specify Policy In** drop-down list, choose **Percentage** or **Packets Per Second**.
    - If you chose **Percentage**, perform the following steps:
      - In the **Rate** field, enter a traffic rate percentage.

        Enter a number between 0 and 100 that specifies a percentage of the total available bandwidth of the port. When the ingress traffic reaches this level during a one second interval, traffic storm control drops traffic for the remainder of the interval. A value of 100 means no traffic storm control. A value of 0 suppresses all traffic.
      - In the **Max Burst Rate** field, enter a burst traffic rate percentage.

Enter a number between 0 and 100 that specifies a percentage of the total available bandwidth of the port. When the ingress traffic reaches this level, traffic storm control begins to drop traffic.

- If you chose **Packets Per Second**, perform the following steps:
  - In the **Rate** field, enter a traffic rate in packets per second.

    During this interval, the traffic level, expressed as packets flowing per second through the port, is compared with the traffic storm control level that you configured. When the ingress traffic reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the interval ends.
  - In the **Max Burst Rate** field, enter a burst traffic rate in packets per second.

    During this interval, the traffic level, expressed as packets flowing per second through the port, is compared with the burst traffic storm control level that you configured. When the ingress traffic reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the interval ends.

- In the **Alias** field, enter the name of the alias.
- From **Storm Control Action** drop-down list, choose a storm control action type.

To view more details about a storm control policy, choose a storm control policy and click **View Details**. The following tab appears:

- **Storm Control Policy Alias**—Displays the name and alias of the storm control policies.

- **Priority Flow Control Policy** tab—Displays the details of the state of Priority Flow Control (PFC) on the interfaces to which this policy group is applied. This policy specifies under what circumstances QoS-level PFC will be applied to FCoE traffic. To add a PFC policy, click **Add** and complete the following fields:
  - **Name** and **Description** fields—Enter a unique name and description for the PFC policy.
  - **State** drop-down list—Choose one of the following states to which PFC policy is applied:
    - **Auto**—Enables PFC on local port on the no-drop CoS as configured, on the condition that values advertised by the DCBX and negotiated with the peer succeed. Failure causes PFC to be disabled on the no-drop CoS. This is the default option.
    - **On**—Enables FCoE PFC on the local port under all circumstances.
    - **OFF**—Disables FCoE PFC on the local port under all circumstances.

- **Slow Drain Policy** tab—Displays the details of the slow drain policy for handling FCoE packets that cause traffic congestion on an ACI Fabric. To create a slow drain policy, click **Add** and complete the following fields:
  - **Name** and **Description** fields—Enter a unique name and description.
  - **Congestion Clear Action** drop-down list—Choose one of the following actions to be taken during FCoE traffic congestion:
    - **Err - disable**—Disable the port. This is the default option.
    - **Log**—Record congestion in the Event Log.
    - **Disabled**—Take no action.

- **Congestion Detect Multiplier** field—Enter the number of pause frames received on a port that triggers a congestion clear action to address FCoE traffic congestion. The default value is 10.
- **Flush Admin State** drop-down list—Choose one of the following state:
  - **Enabled**—Flush the buffer.
  - **Disabled**—Don't flush the buffer. This is the default option.
- **Flush Timeout** field—Enter the threshold in milliseconds to trigger buffer flush drop during congestion. The default value is 500 milliseconds.

- **L2 Interface Policy** tab—Displays the details of the Layer 2 interface policy of the APIC account. To create a Layer 2 interface policy, click **Create** and complete the following fields:
  - Enter a unique name and description for the Layer 2 interface policy.
  - From the **QinQ** drop-down list, choose one of the following options to enable encapsulation:
    - **corePort**—To create an interface policy that enables an interface to be used as a core port in a Dot1q tunnel.
    - **disabled**—To disable Q-in-Q encapsulation.
    - **doubleQtagPort**—To create an interface policy that enables Q-in-Q encapsulation.
    - **edgePort**—To create an interface policy that enables an interface to be used as an edge port in a Dot1q tunnel.
  - From the **Reflective Relay (802.1Qbg)** drop-down list, choose **enabled** to enable reflective relay on a port, port channel, or virtual port channel as a Layer 2 interface policy on the switch. By default, reflective relay is disabled.
  - From the **VLAN Scope** drop-down list, choose **Port Local scope** for the leaf port and **Global scope** for interfaces configured with Multiple Spanning Tree (MST). This field appears only when **corePort**, **disabled**, or **doubleQtagPort** is chosen in the **QinQ** drop-down list.

- **LLDP Interface Policy** tab—Displays the name and description of the Link Layer Discovery Protocol (LLDP) interface policy, with the receive status and transmit status.

  To view more details about a LLDP interface policy, choose a policy and click **View Details**. The following tab appears:
  - **LLDP Interface Policy Alias**—Displays the names of LLDP interface policy and alias.

- **Port Channel Policy** tab—Displays the port channel member of the APIC accounts. To add a port channel policy, click **Add** and complete the following fields:
  - Enter a unique name and description for the port channel policy.
  - **Mode** drop-down list—Choose a port channel policy mode.
  - **Fast Select Hot Standby Ports** check box—Check this to enable fast select for hot standby ports.
  - **Graceful Convergence** check box—Check this to enable graceful convergence.
  - **Load Defer Member Ports** check box—Check this to enable load defer for member ports.

- **Suspended Individual Port** check box—Check this to enable suspended individual port.
- **Symmetric Hashing** check box—Check this to enable symmetric hashing.
- **Minimum Number of Links** field—Enter the minimum number of links. The valid range is from 1 to 16.
- **Maximun Number of Links** field—Enter the maximun number of links. The valid range is from 1 to 16.
- **Alias** field—Enter the alias for the port channel policy.

To view more details about a port channel policy, choose a policy and click **View Details**. The following tab appears:

- **Port Channel Policy Alias** tab—Displays the names of port channel policy and alias.

- **Leaf Policy Group** tab—Displays the name and description of the leaf policy group.
- **Monitoring Policy** tab—Displays the monitoring policy details of the APIC account. To create a monitoring policy, click **Add** and enter a unique name and short description for the monitoring policy.
- **Spanning Tree Interface Policy** tab—Displays the spanning tree interface policy details of the APIC account. The spanning tree interface policy dictates the behavior of southbound leaf port Spanning Tree features. To create a spanning tree interface policy, click **Add** and complete the following fields:
  - Enter a unique name and short description for the spanning tree interface policy.
  - Check the **BPDU filter** check box to enable bridge protocol data unit (BPDU) filter on the spanning tree interface policy.
  - Check the **BPDU guard** check box to enable BPDU guard on the spanning tree interface policy.

  To view more details about a spanning tree interface policy, choose a policy and click **View Details**. The following tab appears:

  - **Spanning Tree Interface Policy Alias**—Displays the names of spanning tree interface policy and alias.

- **PC/VPC Leaf Policy** tab—Displays the name, description, and link aggregation type of the PC/VPC leaf policies.

  To create a PC interface leaf policy, click **Create PC Interface Policy Group** or choose **Create PC Interface Policy Group** from the **More Actions** drop-down list. In the **Create PC Interface Policy Group** screen, enter a unique name and short description for the PC interface policy group. Click **Select** and choose the following policies that you want to use in the PC interface policy group: Link Level Policy, CDP Policy, LLDP Policy, Attached Entry Profile, MCP Policy, CoPP Policy, Storm Control Interface Policy, STP Interface Policy, Port-Channel Policy, Monitoring Policy, L2 Interface Policy, Port Security Policy, Egress Data Plane Policing Policy, Ingress Data Plane Policing Policy, Priority Flow Control Policy, Fibre Channel Interface Policy, Slow Drain Policy, and MACsec Policy.

  To create a VPC interface leaf policy, choose **Create VPC Interface Policy Group** from the **More Actions** drop-down list. In the **Create VPC Interface Policy Group** screen, enter a unique name and short description for the VPC interface policy group. Click **Select** and choose the following policies that you want to use in the VPC interface policy group: Link Level Policy, CDP Policy, LLDP Policy, Attached Entry Profile, Port-Channel Policy, MCP Policy, CoPP Policy, Storm Control Interface Policy, L2 Interface Policy, Port Security Policy, Priority Flow Control Policy, STP Interface Policy, Egress Data Plane Policing Policy, Ingress Data Plane Policing Policy, Fibre Channel Interface Policy, Slow Drain Policy, Monitoring Policy, and MACsec Policy.

  To view more details about a PC/VPC leaf policy, choose a PC/VPC leaf policy and click **View Details**. The following tabs appear:

- **Netflow Monitor Policy**—Displays the netflow monitoring policy and netflow IP filter type associated with the PC/VPC interface policy group. To associate a netflow monitoring policy with the PC/VPC interface policy group, click **Create**. In the screen, choose a **Netflow IP Filter Type** from the drop-down list and click **Select** and choose a netflow monitor policy that needs to be associated with the PC/VPC interface policy group.
- **VDestination Groups**—Displays the VDestination groups associated with the PC/VPC interface policy group. To associate a VDestination group with the PC/VPC interface policy group, click **Add** and choose a VDestination group.
- **VSource Groups**—Displays the VSource groups associated with the PC/VPC interface policy group. To associate a VSource group with the PC/VPC interface policy group, click **Add** and choose a VSource group.
- **Override Policy Groups**—Displays the override policy groups associated with the PC/VPC interface policy group. To create an override policy group, click **Add**. In the **Create Override Policy Group** screen, enter a unique name and short description for the override policy group. Click **Select** and choose a port channel member policy that needs to be associated with the override policy group. Choose an override policy group and click **View Details** to view the PC/VPC interface policy group name, override policy group name, and port channel member policy.

- **Netflow Monitor Policy** tab—Displays the name and description of the NetFlow monitor policy. To add a NetFlow monitor policy, click **Add** and complete the following fields:
  - Enter a unique name and short description for the NetFlow monitor policy.
  - **Associated Flow Record** field—Click **Select** and choose a flow record that needs to be associated with the NetFlow monitoring policy.

  To view more details about a NetFlow monitor policy, choose a NetFlow monitor policy and click **View Details**. The following tabs appear:
  - **Flow Record**—Displays the flow record of the NetFlow monitor policy.
  - **NetFlow Exporter**—Displays the NetFlow exporter associated with the fabric NetFlow monitor policy. To associate a NetFlow exporter with a fabric NetFlow monitor policy, click **Add** and choose a NetFlow exporter to be associated with the policy.

- **Flow Record** tab—Displays the name, collect parameters, match parameters, and description of the NetFlow records. A NetFlow record lets you define a flow and the statistics to collect for each flow. You can define match parameters to identify packets in the flow, and define collect parameters that the NetFlow gathers for the flow. To create a flow record, click **Create** and complete the following fields:
  - Enter a unique name and short description for the NetFlow record.
  - **Collect Parameters** field—Click **Select** and choose a list of parameters that need to be collected for a given flow.
  - **Match Parameters** field—Click **Select** and choose a list of parameters that are used by NetFlow to identify packets in the flow.

- **Relay Policy** tab—Displays the name, description, and mode of the relay policy. Click **Add** to add a relay policy.
- **Tenant(s)** tab—Displays the tenants in the APIC server. Click **Add** to add a tenant.

  To view more details about a tenant, choose a tenant and click **View Details**. The following tabs appear:
  - **Summary**—Displays the overview of the tenant.

- **Application Profile**—Displays the name, tenant, description, and QoS Class of the tenant application profile. Click **Add** to add a tenant application profile. Choose an application profile and click **View Details** to view the EPGs of the application profile.

  Choose an EPG and click **View Details** to view the provided contracts, consumed contracts, Layer 4 to Layer 7 EPG parameters, consumed contract interface, static node, domain, static path, and subnet of the EPG. In the **Consumed Contract Interface** tab, click **Add** to add a consumed contract interface to EPG.

- **Deployed Service Graph**—Displays the list of service graphs that are deployed in the tenant. Choose a service graph and click **View Details** to view the Layer 4 to Layer 7 deployed service graph parameters.

- **Filters**—Displays the tenant, name, and description of the filters. To view the tenant filter rules, choose a filter and click **View Details**.

- **Bridged Outside**—Displays the tenant, name, and description of the external bridge network. Choose a network and click **View Details** to view the following tabs:

  - **External Network**—Choose an external network and click **View Details** to view the provided contracts, and consumed contracts details.

  - **Node Profile**—Choose a node profile and click **View Details** to view the interface profile details.

- **Routed Outside**—Displays the tenant, name, and description of the external routed network. Choose a network and click **View Details** to view the following tabs:

  - **Route Map or Profile**—Choose a route profile and click **View Details** to view the context details.

  - **Logical Node Profile**—Choose a logical node profile and click **View Details**. The following tabs appear:

    - **Logical Nodes** tab—Displays the logical nodes. Click **Add** to add a logical node to the logical node profile of the external routed network. Choose a logical node and click **View Details** to view the static routes to the logical node.

    - **Logical Interface Profile** tab—Choose a logical interface profile and click **View Details** to view the logical interface and logical OSPF interface. Click **Add** in the Logical OSPF Interface tab to create an interface profile with the OSPF profile data.

    - **BGP Peer Connectivity** tab—Displays the BGP peer connectivity of the logical node profile. Click **Add** to add a peer connection to a node profile.

  - **External Network**—Choose an external network and click **View Details** to view the subnet, provided contracts, and consumed contracts details. You can tag an external network and consumed contract using the **Add Tags** option. The tag is used to identify the network and contract that you want to use in the application container deployment.

- **Bridge Domains**—Displays the tenant, name, description, segment ID, unicast traffic, ARP flooding, multicast IP address, customer MAC address, unicast route, and Layer 2 unknown unicast value.

  To view more details about a bridge domain, choose a bridge domain and click **View Details**. The following tabs appear:

  - **DHCP Relay Label**—Displays the tenant, name, description, and scope of the DHCP relay.

  - **Subnet**—Displays the tenant, bridge domain, description, subnet control, and gateway address of the tenant.

- **VRF**—Displays the tenant name, name, description, policy control, and segment of the VRFs. Click **Add** to add a create VRF.
- **BGP Timers**—Displays the tenant, name, graceful restart control, hold interval, keepalive interval, and stale interval of the Border Gateway Protocol (BGP) timer.
- **Contracts**—Displays the tenant, name, description, type, QoS, and scope of the contracts.

  To view more details about a contract, choose a contract and click **View Details**. The following tabs appear:

  - **Contract Subject**—Choose a contract subject and click **View Details** to view the filter chain, filter chain for consumer to provider, filter chain for provider to consumer, provided label, and consumed label. Each tab has the **Add** option to add a filter, in term filter, out term filter, provided label, and consumed label to a contract subject.
  - **Exported Tenants**—Displays the contracts of the exported tenants.

- **Taboo Contracts**—Displays the tenant, name, description, and scope of the taboo contracts.
- **Relay Policy**—Displays a list of the relay policies.
- **Option Policy**—Displays a list of the option policies.
- **End Point Retention**—Displays the tenant, name, description, hold interval, bounce trigger, bounce entry aging interval, local endpoint aging interval, remote endpoint aging interval, and move frequency of the tenant.
- **OSPF Interface**—Displays the tenant, name, description, network type, priority, cost of interface, interface controls, hello interval, dead interval, retransmit interval, and transmit delay of the Open Shortest Path First (OSPF) interface. Click **Create** to create an OSPF interface policy.
- **EIGRP Interface**—Displays the EIGRP Interface details.
- **OSPF Timers**—Displays the OSPF timer details.
- **IGMP Snoop**—Displays the IGMP snoop details.
- **Custom QOS**—Displays the custom QoS details.
- **Set Rules for Route Map**—Displays the set rules for route map of the tenant. Click **Create** to create a set rules for route map. In the **Create Set Rules For A Route Map** dialog box, enter the name and description of the action rule profile. To set action rules, check the required check boxes and fill the additional field that appears on selecting the check box.
- **L4-L7 Service Graph**—Displays the Layer 4 to Layer 7 service graph details. Choose a service graph and click **View Details** to view the following tabs:
  - **Consumer EPG**—Displays the list of EPGs that are labeled as consumer in tenants. When an EPG consumes a contact, the endpoints in the consuming EPG may start communication with any endpoint in an EPG that is providing that contract.
  - **Provider EPG**—Displays the list of EPGs that are labeled as provider in tenants. When an EPG provides a contract, communication with that EPG can be initiated from other EPGs as long as the communication complies with the provided contract.
  - **Nodes**—Displays the list of nodes in the tenant. Choose a node and click **View Details** to view the node functions and connectors of the node. Choose a node function and click **View Details** to view the Layer 4 to Layer 7 function node parameters.

- **Connections**—Displays the list of connections in the tenant. Choose a connection and click **View Details** to view the connection terminals in the tenant.

- **Function Profile Group——**—Displays the function profile groups of tenants. Choose a function profile group and click **View Details** to view the function profiles of the group. Click **Add** to add a function profile. To view more details about a function profile, choose a function profile and click **View Details**. The following tabs appear:
  - **Function Profile Parameter**—Displays the function profile parameters. In the **Function Profile Parameter** tab, you can add an ACL, an interface, and add a bridge group interface to a function profile, and add a network object to a function profile. Choose a function profile parameter and click **View Details** to view the function profile parameter configuration and function profile parameter level-one folder.
  - **L4-L7 Function Profile Parameters**—Displays the list of Layer 4 to Layer 7 function profile parameters.
  - **Function Profile Function Parameter**—Displays the list of function profile function parameters. Click **View Details** to view the function profile function parameter Rel details.

- **L4-L7 Devices**—Displays the device cluster details. To view more details about a device cluster, choose a device cluster and click **View Details**. The following tabs appear:
  - **Device Cluster State**—Displays the cluster name, device state, and configured status of the device.
  - **Concrete Device**—Displays the list of concrete devices. Choose a concrete device and click **View Details** to view the virtual network interface card (vNIC) to concrete interface and the path to concrete interface.
  - **Cluster Interface**—Displays the list of cluster interfaces in the device cluster. Choose a cluster interface and click **View Details** to view the cluster interface details.

- **Deployed Device Cluster**—Displays the device clusters that are deployed in the tenant.
- **Imported Device Cluster**—Displays the device clusters that are imported in the tenant.
- **Router Configurations**—Displays the router configurations of the tenant. Click **Add** to add a router configuration.
- **Logical Device Context**—Displays the logical device context details. Choose the logical device context and click **View Details** to view the cluster interface context.

- **Layer 3 Domain** tab—Displays a list of Layer 3 domains in the APIC accounts. To create a Layer 3 domain, click **Create (+)**.

  On the **Create Layer 3 Domain** screen, complete the following fields:
  - **L3 Domain** field—Name of the Layer 3 domain.
  - **Associated Attachable Entity Profile** field—Click **Select** and check an attachable access entry profile that you want to associate with the Layer 3 domain.
  - **VLAN Pool** field—Click **Select** and check a VLAN pool.
  - Click **Submit**.

  Choose a layer 3 domain and click **View Details** to view the security domain of the layer 3 domain. Click **Add** to select and associate a security domain to the layer 3 domain.

- **Layer2 Domain** tab—Displays a list of Layer 2 domains in the APIC accounts. To create a Layer 2 domain, click **Create(+)**.

  On the **Create Layer 2 Domain** screen, complete the following fields:

  - **L2 Domain** field—Name of the Layer 2 domain.
  - **Associated Attachable Entity Profile** field—Click **Select** and check an attachable access entry profile that you want to associate with the Layer 2 domain.
  - **VLAN Pool** field—Click **Select** and check a VLAN pool.
  - Click **Submit**.

- **BGP Route Reflector Policy** tab—Displays the BGP route reflector policy configured for the APIC server. To add a BGP route reflector policy, click **Configure** and provide the short description and autonomous system (AS) number for the BGP route reflector policy. Choose the BGP route reflector policy and click **View Details** to the view the following:

  - **Autonomous System Number**—Displays the AS number defined for the BGP route reflector policy.
  - **External Route Reflector Node**—Displays the external route reflector node of the BGP route reflector policy. To add an external route reflector node, click **Add** and complete the following fields:
    - Choose a spine node from the drop-down list. Choose **Enter Customized Value** from the drop-down list, to set any value between 101 and 4001 as the spine node in the **Spine Node** field.
    - Enter a short description for the external route reflector node.
    - Click **Submit**.
  - **Route Reflector Node**—Displays the route reflector node of the BGP route reflector policy. To add a route reflector node, click **Add** and complete the following fields:
    - Choose a spine node from the drop-down list. Choose **Enter Customized Value** from the drop-down list, to set any value between 101 and 4001 as the spine node in the **Spine Node** field.
    - Enter a short description for the route reflector node.
    - Click **Submit**.

- **VM Networking** tab—Displays the virtual machine (VM) networks with the vendor detail.

  To view more details about a VM network, choose a VM and click **View Details**. The following tabs appear:

  - **Domains**—Displays a list of VMware domains with the vendor details. To add a new domain, click **Add** and complete the following fields:
    - **Domain Name** field—Enter a unique name for the VMM domain.
    - **Vendor** drop-down list—Choose a vendor of the VMM domain from the drop-down list.
    - **Virtual Switch** drop-down list—Choose a virtual switch from the drop-down list.
    - **Delimiter** field—Enter ~, !, +, @, ^, or | as a delimiter character.
    - **ARP Learning** check box—This field appears only on choosing **Cisco AVS** as the virtual switch. Check this box to enable the ARP learning

- **Enable Tag Collection** check box—Check this box to enable tag collection for the domain.
- **Access Mode**—This field appears only on choosing **VMware vSphere Distributed Switch** as the virtual switch. Choose **Read Only Mode** or **Read Write Mode** as the domain access mode.
- **AVE Time Out Time Interval (seconds)** field—This field appears only on choosing **Cisco AVE** as the virtual switch. Enter a value in the range of 10 to 300 as the AVE time out interval in seconds.
- **Host Availability Assurance** check box—This field appears only on choosing **Cisco AVS** as the virtual switch. Check this box to enable host availability assurance in the domain.
- **Endpoint Retention Time** field—Enter an endpoint retention time in seconds to delay the deletion of an endpoint. The valid range is from 0 to 600.
- **Switching Preference** drop-down list—This field appears only on choosing **Cisco AVS** as the virtual switch. Choose **No Local Switching** or **Local Switching** as the switching preference form the drop-down list.
- **Multicast Address** field—This field appears only on choosing **Cisco AVS** as the virtual switch. Enter a multicast IP address. The valid range is from 224.0.0.0 to 239.255.255.255.
- **Multicast Address Pool** field—This field appears only on choosing **Cisco AVS** as the virtual switch. Click **Select** and choose a multicast address pool that you want to use for the VMM domain.
- **Associated Attachable Entity Profile** field—Click **Select** and choose an attachable access entity profile that needs to be associated with the domain.
- **VLAN Pool** field—This field appears on choosing **VMware vSphere Distributed Switch** or **Cisco AVE** as the virtual switch. Click **Select** and choose a VLAN pool that needs to be associated with the domain.
- **Security Domains** field—Click **Select** and choose one or more security domains that need to be associated with the domain.
- **vCenter Credentials** check box—Check this box to set the credentials for vCenter. On choosing this box, the additional fields appear to provide name, description, user name, and password of the vCenter account profile.
- **Controller Type** drop-down list—Choose **vCenter**, **vShield**, or **None** as the controller type. On choosing **vCenter** or **vShield**, the additional fields appear to provide host name or host IP address, enable statistics collection, choose DVS version, and fields to provide data center, management EPG, and associated credential.
- **Port Channel Mode** drop-down list—Choose a mode for the port channel policy.
- **vSwitch Policy** drop-down list—Choose a vSwitch policy from the drop-down list.
- **BPDU Guard** check box—This field appears on choosing **Cisco AVS** or **Cisco AVE** as the virtual switch. Check this box to enable BPDU guard in the interface policy creation.
- **BPDU Filter** check box—This field appears on choosing **Cisco AVS** or **Cisco AVE** as the virtual switch. Check this box to enable BPDU filter in the interface policy creation.
- **Firewall Mode** drop-down list—This field appears on choosing **Cisco AVS** or **Cisco AVE** as the virtual switch. Choose **Enabled**, **Disabled**, or **Learning** as the firewall mode for the domain.
- **Netflow Exporter Policy** field—Click **Select** and choose a netflow exporter policy that needs to be associated with the domain. On choosing a netflow exporter policy, the additional fields appear to provide active flow timeout, idle flow timeout, and sampling rate.

Choose a VMware domain and click **View Details** to view the VMware domain controllers, vCenter credential, and vCenter/vShield. Choose a VMware domain controller and click **View Details** to view the distributed virtual switch (DVS), hypervisors, and virtual machine. Choose a DVS and click **View Details** to view the DVS port groups.

- **L4-L7 Service Device Types** tab—Displays the Layer 4 to Layer 7 service device types with their model, vendor, version, and capabilities.

  To view more details about the Layer 4 to Layer 7 service device type, choose a Layer 4 to Layer 7 service device type and click **View Details**. The following tabs appear:

  - **L4-L7 Service Device Properties**—Displays the vendor, package name, package version, and logging level of Layer 4 to Layer 7 service device types.
  - **L4-L7 Service Device Interface Labels**—Displays a list of interface labels.
  - **L4-L7 Service Functions**—Displays a list of service functions. Choose a service function and click **View Details** to view the details of the Layer 4 to Layer 7 service function connectors.

- **Fabric Nodes Topology** tab—Displays the topology details of fabric nodes.
- **L2 Neighbors** tab—Displays the Layer 2 neighbor details that include the protocol, fabric name, device ID, capability, port ID, local interface, hold time, and platform.
- **Deployed Service Graph** tab—Displays the tenant, contract, state, service graph, context name, node function, and description of the APIC account.
- **EPG to Contract Association** tab—Displays the details of the contract association with EPGs.
- **Access Port Policy Groups** tab—Displays the access port policy group name, link level policy, Cisco Discovery Protocol (CDP) policy, Link Aggregation Control Protocol (LACP) policy, Link Layer Discovery Protocol (LLDP) policy, link aggregation type, and attached entity profile of the accounts in the APIC server.
- **Fabric Interface Profiles** tab—Displays the fabric interface profiles of the APIC server.

  To view more details about a fabric interface profile, choose a profile and click **View Details**. The following tab appears:

  - **Access Port Selector**—Displays the access port selectors of the fabric interface profile. To add an access port selector to the fabric interface profile, click **Add** and complete the following fields in the **Create Access Port Selector** screen:
    - **Name** and **Description** field—Enter a unique name and description for the access port selector.
    - **Interface IDs** field—Enter comma separated interface IDs. You can enter **All** or range of interface IDs.
    - **Interface Description** field—Enter comma separated description for each interface ID.
    - **Connected to FEX** check box—If the port is connected to a Cisco Fabric Extender (FEX), check this check box.
    - **Interface Policy Group** field—Click **Select** and choose an interface policy group that needs to be associated to these ports.

  Choose an access port selector and click **View Details** to view the port blocks and sub port blocks of the access port selector.

- **Fabric Switch Profiles** tab—Displays the fabric switch profiles of the APIC account.

- **Spine Access Port Policy Group** tab—Displays the details of the spine access port policy group of the APIC account. To add a spine access port policy group, click **Add**.

  On the **Create Spine Access Port Policy Group** screen, complete the following fields:

  - **Spine Access Port Policy Group Name** field—Enter a unique name for the spine access port policy group.
  - **Description** field—Enter a short description for the spine access port policy group.
  - **Link Level Policy** field—Click **Select** and check a link level policy that you want to associate with the spine access port policy group.
  - **CDP Policy** field—Click **Select** and check a CDP policy that you want to associate with the spine access port policy group.
  - **MACsec Policy** field—Click **Select** and check a MACsec policy that you want to associate with the spine access port policy group.
  - **Attached Entity Profile** field—Click **Select** and check an attachable access entity profile that you want to associate with the spine access port policy group.
  - Click **Submit**.

- **System** tab—Displays the system details that includes node name, in-band management IP address, out-of-band management IP address, infra IP address, fabric MAC address, ID, role, and serial number.
- **Fabric Memberships** tab—Displays the details of the fabric membership of the APIC account. To create a fabric membership, click **Create** and complete the following fields:
  - **Pod ID** field—Enter a unique identifier of the pod where the node is located.
  - **Serial Number** field—Enter a serial number of the switch.
  - **Node ID** field—Enter a number greater than 100 as node ID, as the first 100 IDs are reserved for APIC appliance nodes.
  - **Switch Name Number** field—Enter the name of the switch added to the pod.
  - **Node Type** field—Choose **leaf** or **spine** as the node type. The default value is **unspecified**.
- **Access Entity Profile** tab—Displays the entity profile details of the APIC server. To add an access entity profile, click **Add** and provide the unique name and description for the access entity profile, enable or disable infrastructure VLAN, and then choose domain profile and interface policy group for the access entity profile.
- **Multicast Address** tab—Displays the multicast address of the APIC server. To add a multicast address, click **Add** and provide a unique name and description for the multicast address.

  Choose a multicast address and click **View Details** to view the range of addresses assigned to the multicast address. To add an address block to the multicast address, click **Add** and enter the starting and ending IP addresses of the multicast address block. The valid range of multicast address is from 224.0.0.0 to 239.255.255.255.
- **Fabric Spine Profile** tab—Displays the fabric spine profile of the APIC server. To add a fabric spine profile, click **Add** and provide a unique name and description for the fabric spine profile.
- **Fabric TEP Pool Physical Pod** tab—Displays the fabric TEP pool physical pod of the APIC server. To add a fabric TEP pool physical or virtual pod, click **Add** and complete the following fields:
  - **Pod ID**—Enter defaultValue or a number in the range of 1 to 254 as the pod ID. The default value is 1.

- **TEP Pool**—Enter an IPv4 address with mask as the Tunnel Endpoint (TEP) pool for the pod. For physical pod, the mask must be in the range of 1 to 23. For virtual pod, the mask must be in the range of 22 to 28.
- **Pod Type**—Choose **physical** or **virtual** as the pod type. The default value is **physical**.

- **Fabric TEP Pool Virtual Pod** tab—Displays the fabric TEP pool virtual pod of the APIC server.

  To view more details about the fabric TEP pool virtual pod, choose a fabric TEP pool virtual pod and click **View Details**. The following tabs appear:

  - **Fabric TEP Pool Virtual Pod Properties**—Displays the virtual pod ID, pod type, reserved address and gateway address of the fabric TEP pool virtual pod.
  - **Additional Subnets**—Displays the associated subnets of the fabric TEP pool virtual pod. To add an additional subnet to fabric virtual pod TEP pool, click **Add** and complete the following fields:
    - **Network Address** field—Enter a valid IPv4 address with subnet mask as the network address for the subnet of fabric virtual pod. The valid format is IPv4/subnet mask. The valid range of subnet mask is from 22 to 28.
    - **Reserved Subnet** field—The pool consists of IP addresses that are not sent by the DHCP server and which you can reserve for a specific use. Enter a valid IPv4 address with subnet mask as the reserved subnet. The valid format is IPv4/subnet mask. The value of subnet mask must be greater than the subnet mask of TEP pool by two.
    - **Gateway IP Address** field—Enter a valid IPv4 address as the gateway IP address of the subnet.

    Choose an additional subnet and click **View Details** to view the properties of the additional subnet.
  - **DHCP Servers**—Displays the DHCP servers of the fabric TEP pool virtual pod. To add a DHCP server, click **Add** and complete the following fields:
    - **Node ID** field—Enter a unique node ID for the DHCP server. The valid range of node ID is from 101 to 4000.
    - **Type** drop-down list—Choose **Primary** or **Secondary** as the DHCP server type.

- **Fabric Configured Switch Interfaces** tab—Displays the fabric configured switch interfaces of the APIC server.
- **Leaf Profiles** tab—Displays the leaf profiles of the APIC server.

Assigning an APIC Account to a Pod

In the **Converged** menu of the user interface (UI), Cisco UCS Director displays the converged stack of devices for a data center. To display the APIC account in the converged UI, assign the APIC account to a pod.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account that you want to assign to a pod.

Step 4 From the **More Actions** drop-down list, choose **Assign to Pod**.

The **Assign to Pod** screen appears.

Step 5 Click **Select** and check a pod to which you want to assign the APIC account.

Step 6 Click **Submit**.

The APIC account appears in the converged UI.

Handling APIC Failover

APIC controllers are deployed in an APIC cluster. The recommendation is to have a minimum of three APIC controllers per cluster to ensure high availability. When you create an APIC account in Cisco UCS Director, provide the IP address of one of the APIC controllers in the APIC cluster. Cisco UCS Director discovers the other APIC controllers in the APIC cluster and their respective IP addresses.

If the IP address of the controller which was used to manage the APIC device goes down or is not reachable for 45 seconds, Cisco UCS Director tries to use any of the reachable controller IP addresses to interact with the APIC device.

If you have multiple ACI fabrics and each fabric with multiple controllers, one of the controllers of the ACI fabric is used to manage the APIC device. If the controller goes down or is not reachable for 45 seconds, Cisco UCS Director uses the next reachable controller within the ACI fabric.

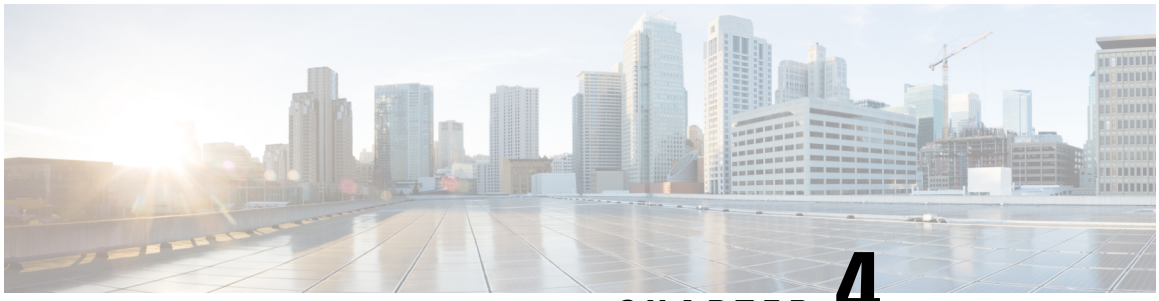

CHAPTER 4

Managing Tenants

Tenants

A tenant is a logical container for application policies that enables you to exercise domain-based access control by isolating the resources such as applications, databases, web servers, network-attached storage, virtual machines, firewalls, Layer 4 to Layer 7 services, and so on. Tenants can represent a customer in a service provider setting, an organization or domain in an enterprise setting, or just a convenient grouping of policies.

A fabric can contain anywhere from one tenant, which may be useful for a small commercial environment, to 64,000+ tenants, for a cloud service provider in which case you assign each company their own tenant. Another use case would be to have a Dev tenant and a Production tenant. In this case, you create network constructs, EPGs, and policies in Dev tenant first and then simply copy it to the Production tenant. It ensures that the dev and prod are the exact same and takes away the human error that comes along with manual copying of these objects.

Note Configure a tenant before you can deploy any Layer 4 to Layer 7 services.

Tenant Types

The system provides the following four kinds of tenants:

- User tenant—Defined by the administrator according to the needs of users. It contains policies that govern the operation of resources such as applications, databases, web servers, network-attached storage, virtual machines, and so on.
- Common tenant—Provided by the system but can be configured by the fabric administrator. It contains policies that govern the operation of resources accessible to all tenants, such as firewalls, load balancers, Layer 4 to Layer 7 services, intrusion detection appliances, and so on.
- Infrastructure tenant—It contains policies that govern the operation of infrastructure resources such as the fabric VXLAN overlay.
- Management tenant—It contains policies that govern the operation of fabric management functions used for in-band and out-of-band configuration of fabric nodes.

Tenant Features

- Tenants can be isolated from one another or can share resources.
- Tenants do not represent a private network.
- Entities in the tenant inherit its policies.
- The primary elements that the tenant contains are filters, contracts, outside networks, bridge domains, Virtual Routing and Forwarding (VRF) instances, and application profiles that contain endpoint groups (EPGs).

Note In the APIC GUI under the tenant navigation path, a VRF (context) is called a private network.

Setting up a Tenant

This procedure provides an overview of how to set up a tenant for an APIC account in Cisco UCS Director. You can also use the workflows provided in Cisco UCS Director Orchestration to complete a guided setup of tenants for various use cases. For more information, see Cisco UCS Director Orchestration Guide.

This procedure assumes that you have already completed the following prior to creating tenants:

- The Day 0 setup of ACI fabric.
- The nodes in ACI fabric are connected and discovered.
- The APIC controller cluster has been configured.
- Cisco UCS Director is configured and the ACI pod has been set up.

Step 1 Create a Tenant.

See Creating a Tenant, on page 40.

Step 2 Create a Virtual Routing and Forwarding (VRF) (also known as Private Network).

See Creating a VRF, on page 43.

Step 3 Add Bridge Domain to the VRF.

See Adding a Bridge Domain to VRF, on page 49.

Step 4 Create Application Profiles.

See Creating an Application Profile for the Tenant, on page 54.

Step 5 Create EPGs.

See Adding an EPG, on page 55.

Step 6 Add domain to EPGs.

See Adding a Domain to an EPG, on page 56.

Step 7 Add Static path to EPGs.

See Adding a Static Path to EPG, on page 58.

Step 8 Create Contracts.

See Creating Contracts, on page 63.

Step 9 Add contracts to EPGs.

See Adding a Consumed Contract to an EPG, on page 66.

See Adding a Provided Contract to an EPG, on page 65.

Creating a Tenant

Before you begin

Verify that Tags, monitoring policy, and security domains for the objects in the APIC account are configured before adding a tenant.

Create users in ACI and assign a security domain to the users or user groups. See User Access, Authentication, and Accounting chapter in Cisco APIC Basic Configuration Guide.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click **Add**.

Step 6 On the **Add APIC Tenant** screen, complete the fields, including the following:

a) Add a unique name and description for the Tenant.

b) Enter an alias name for the tenant. While the tenant name cannot be changed after creation, the alias name of the tenant can be changed as required.

c) (Optional) Click **Select** and check the tag that you want to use.

Tags are used to assign a descriptive name to a group of objects. For example, to enable easy searchable access to all web server EPGs, assign a web server tag to all such EPGs. Web server EPGs throughout the fabric can be located by referencing the web server tag.

d) (Optional) Click **Select** and check the monitoring policy that you want to use.

When you apply a monitoring policy, it overrides the default monitoring policy.

e) Click **Select** and check the security domain that you want to use.

It is necessary to also assign the user to one or more security domains. By default, the ACI fabric includes two special pre-created domains:

- **All**—Allows access to the entire management information tree (MIT).
- **Infra**—Allows access to fabric infrastructure objects/subtrees, such as fabric access policies.

For example, if you have created a security domain for Production, given users roles, and attached them to that security domain, then choose the Production security domain instead of **All**.

f) Enter the unique name for VRF within the tenant.

g) Click **Submit**.

What to do next

After creating a tenant, create a VRF (also known as a private network) for the tenant.

Adding a GUID to a Tenant

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **GUID**.

Step 7 Click **Add**.

Step 8 On the **Add GUID to Tenant** screen, complete the following fields:

a) Enter the globally unique identifier (GUID) for an SCVMM provider. GUID must be in the format : xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.

b) Enter the unique name for the SCVMM account.

Step 9 Click **Submit**.

Viewing Tenants

You can view a list of tenants that are onboarded in Cisco UCS Director and its details.

Step 1 Choose **Policies** > **Resource Groups**.

Step 2 On the **Resource Groups** page, click **Tenant**.

Step 3 Click the row with the tenant for which you want to view details.

Step 4 Click **View Details** to view the service offerings of the tenant.

Step 5 Click the row with the service offering and click **View details** to view the resource groups of a tenant.

Note If the disaster recovery support is enabled for the tenant, the resource groups of the primary site and the disaster recovery site are displayed.

Step 6 Click the row with the resource group and click **View details** to view the following information:

- **Resource Entity**—Displays a list of available resources, such as, VMWare cluster, resource pool, and data store, in a vPOD. During tenant onboarding, the resources matching the capacity, capability and tag of the tenant requirement are filtered from resource group and matched resources are added to the vPOD. With the capacity expansion support, the vPOD can store more than one resources for each resource type such as VMware cluster, resource pool, and storage pool. As multiple resources of same resource type is available in vPOD, the tenant expansion is possible after consumption of allocated resources.

  The tenant-specific and container-specific resource limits assist in provisioning VMs and BMs. During provisioning, all the available resources in vPOD are referred to find out the matching resources for resource allocation. After the resource filtration and selection, the matching resources from the same account are allocated for VM deployment.

  When a resource is no longer consumed by the container, you can delete the resource. To delete the resource, click the row with the resource and click **Delete**.

- **Tenant Details**—Displays more details of the tenant.

- **Tenant Resource Limits**—Displays availability of both virtual and physical resources in a tenant. The resources reserved during tenant onboarding are displayed along with the used and available resource values. The VDCs Limit column specifies the maximum number of containers that are reserved for the tenant. The Available Number of VDCs column represents the number of containers that are available for provisioning. The physical resource limits display the blades that are reserved as part of tenant onboarding, along with the number of blades used for bare metal provisioning.
- **Container Resource Limits**—Displays availability of both virtual and physical resources in a container. The resource limits that are set during container creation are displayed along with the used and available resources.

  Note If a container is created without a resource limit, the value of the virtual resources is displayed as Not Set.

- **Private Network**—Displays the private networks created for the tenant. Click the row of a private network and click **View Details** to view the supernet and subnet pools of the private network. The **Supernets** screen lists the supernets available for the tiers. The **Subnets** screen displays the sub-network pool that is used for load balancer configuration during the container deployment.

Virtual Routing and Forwarding (VRF)

A Virtual Routing and Forwarding (VRF) is similar to a virtual router that defines a Layer 3 address domain. It is an IP technology that allows multiple instances of a routing table to coexist on the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflict. For example, a production VRF could be on the same network as the development VRF but the two have different default gateways.

A tenant can have multiple VRFs (also known as private network). One or more bridge domains are associated with a VRF. There are several policies you can associate with a private network, including OSPF and BGP timers, as well as how long end points should be retained.

Virtual Routing and Forwarding (VRF) Guidelines

The following guidelines and limitations apply for virtual routing and forwarding (VRF) instances:

- Within a single VRF instance, IP addresses must be unique. Between different VRF instances, you can have overlapping IP addresses.
- If shared services are used between VRF instances or tenants, make sure that there are no overlapping IP addresses.
- Any VRF instances that are created in common tenant is seen in other user-configured tenants.
- VRF supports enforced mode or unenforced mode. By default, a VRF instance is in enforced mode, which means all endpoint groups within the same VRF instance cannot communicate to each other unless there is a contract in place.
- Switching from enforced to unenforced mode (or the opposite way) is disruptive.

For more in-depth information, see the Cisco Application Centric Infrastructure Fundamentals Guide.

Creating a VRF

A Virtual Routing and Forwarding (VRF) object (also known as private layer 3 network in ACI) contains the Layer 2 and Layer 3 forwarding configuration, and IP address space isolation for tenants. Each tenant can have one or more VRFs, or share one default VRF with other tenants as long as there is no overlapping IP addressing being used in the ACI fabric.

Before you begin

Verify that you have configured the BGP Timers Policies and OSPF Timers

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click **Add**.

Step 8 On the **Create VRF** screen, complete the following fields:

a) Enter a unique name for the private network and alias.

b) Click **Tags** and check the tag that you want to use.

c) From the **Policy Control Enforcement Preference** drop-down list, choose from the following options:

- **Enforced**—Security rules (contracts) are enforced.
- **Unenforced**—Security rules (contracts) are not enforced.

The default is **enforced**.

d) From the **Policy Control Enforcement Direction** drop-down list, choose **Egress** or **Ingress**. The default is **Ingress**.

VRF enforcement must be set to **Egress** when the QoS classification is done in the contract.

e) Check the **BD Enforcement Status** check box to enable BD enforcement status.

f) From the **IP Data Plane Learning** drop-down list, choose **Enabled** or **Disabled** to enable or disable the data-plane IP learning on the VRF. The default is **Enabled**.

g) Enter the description for the private network.

h) Click **Select** and check the BGP timer that you want to use.

The Border Gateway Protocol (BGP) timer policy enables you to specify the intervals for the periodic activities and supplies two options for graceful restart control.

i) Click **Select** and check the OSPF timer that you want to use.

The context-level OSPF timer policy provides the Hello timer and Dead timer intervals configuration. OSPF timers control the behavior of protocol messages and shortest path first (SPF) calculations.

j) Click **Select** and check the monitoring policy that you want to associate with the tenant.

When you apply a monitoring policy, it overrides the default monitoring policy.

k) Enter the comma separated DNS labels for private network.

l) Click **Submit**.

m) Click **Select** and check the endpoint retention policy that you want to associate with the tenant.
n) Click **Select** and check the transit route tag policy that you want to associate with the tenant.
o) Check the **Enable GOLF-OPFLEX Mode** check box to enable GOLF-OPFLEX mode.
p) Enter the DCI VRF name. This field is displayed only when the GOLF-OPFLEX mode is enabled.

Step 9 Click **Submit**.

What to do next

After creating a VRF, you create a bridge domain and link it to this VRF.

Adding an EIGRP to the VRF

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with the VRF to which you want to add an EIGRP and click **View Details**.
Step 8 Click **EIGRP**.
Step 9 Click **Add**.
Step 10 On the **Add APIC EIGRP to VRF** screen, complete the following fields:
a) Choose **IPv4 unicast address family** or **IPv6 unicast address family** as the EIGRP address family type, to accordingly configure an Enhanced Interior Gateway Routing Protocol (EIGRP) routing instance.
b) Click **Select** and choose an EIGRP that you want to add to the VRF.
Step 11 Click **Submit**.

Adding a BGP Route Target Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with VRF where you want to add a BGP route target profile and click **View Details**.
Step 8 Click **BGP Route Target Profile**.
Step 9 Click **Add**.

Step 10 On the **Add BGP Route Target Profile** screen, choose IPv4 or IPv6 as the BGP route target profile type. The default value is **IPv4 unicast address family**.

Step 11 Click **Submit**.

Adding a BGP Route Target to BGP Route Target Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with VRF where you want to add a BGP route target profile and click **View Details**.

Step 8 Click **BGP Route Target Profile**.

Step 9 Click the row with the BGP route target profile where you want to add a BGP route target and click **View Details**.

Step 10 Click **Add**.

Step 11 On the **Add Route Target to BGP Route Target Profile** screen, complete the following fields:

a) From the **Type** drop-down list, choose **Import** or **Export**.

b) Specify the route target.

Step 12 Click **Submit**.

Adding a BGP Context per Address Family

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with VRF where you want to add a BGP context per address family and click **View Details**.

Step 8 Click **BGP Context Per Address Family**.

Step 9 Click **Add**.

Step 10 On the **Add BGP Context Per Address Family** screen, complete the following:

a) From the **BGP Address Family Type** drop-down list, choose IPv4 or IPv6.

b) Click **Select** and check the BGP address family context that you want to use for the VRF.

Step 11 Click **Submit**.

Adding a OSPF Context per Address Family

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with VRF where you want to add an OSPF context per address family and click **View Details**.
Step 8 Click **OSPF Context Per Address Family**.
Step 9 Click **Add**.
Step 10 On the **Add OSPF Context Per Address Family** screen, complete the following:
a) From the **OSPF Address Family Type** drop-down list, choose IPv4 or IPv6.
b) Click **Select** and check the OSPF timer that you want to use for the VRF.
Step 11 Click **Submit**.

Adding a SNMP Context

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with VRF where you want to add a SNMP context and click **View Details**.
Step 8 Click **SNMP Context**.
Step 9 Click **Add**.
Step 10 On the **Add SNMP Context** screen, specify a SNMP context name.
Step 11 Click **Submit**.

Creating a Community Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with VRF where you want to add a community profile and click **View Details**.
Step 8 Click **Community Profile**.
Step 9 Click **Create**.
Step 10 On the **Create Community Profile** screen, add a unique name and description for the community profile.
Step 11 Click **Submit**.

BGP Timers

Border Gateway Protocol (BGP) Timers can be defined and associated on a per VRF per node basis. A node can have multiple VRFs, each corresponding to a fvCtx. A node configuration (l3extLNodeP) can now contain configuration for BGP Protocol Profile (bgpProtP) which in turn refers to the desired BGP Context Policy (bgpCtxPol). This makes it possible to have a different node within the same VRF contain different BGP timer values.

For each VRF, a node has a bgpDom concrete MO. Its name (primary key) is the VRF, <fvTenant>:<fvCtx>. It contains the BGP timer values as attributes (for example, holdIntvl, kaIntvl, maxAsLimit). All the steps necessary to create a valid Layer 3 Out configuration are required to successfully apply a per VRF per node BGP timer. For example, MOs such as the following are required: fvTenant, fvCtx, l3extOut, l3extInstP, LNodeP, bgpRR.

When a BGP timer is configured on a specific node, then the BGP Timer policy on the node is used and the BGP policy timer associated with the VRF is ignored.

Adding a BGP Timer Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **BGP Timers**.
Step 7 Click **Add**.
Step 8 On the **Create APIC BGP Timers Policy** screen, complete the fields including the following:

a) Enter the name and description for the BGP Timer policy.
b) In the **Keepalive Interval (sec)** field, enter the keepalive value to retain the route information learned from BGP in the routing table. The keepalive interval must be in the range of 0 to 3600. The default value is 60.
c) In the **Hold Interval (sec)** field, enter the hold-time value to use when negotiating a connection with the peer. The hold interval must be in the range of 0 to 3600. The default value is 180.
d) In the **Stale Interval (sec)** field, enter the period of time for which stale routes must be preserved by using the long-lived graceful restart capability for BGP sessions on the restarting router. The stale interval must be in the range of 1 to 3600. The default value is 300.

e) Check the **Graceful Restart Controls** check box to enable or turn on the helper mode to assist a neighboring router attempting a graceful restart.
f) In the **Maximum AS Limit** field, enter the maximum allowed number of autonomous system (AS) in the range of 0 to 2000. The default value is 0.

Step 9 Click **Submit**.

Bridge Domains

A bridge domain represents a Layer 2 forwarding construct within the fabric. It helps you to constrain broadcast and multicast traffic. It is a logical container for subnets.

A bridge domain must have at least one subnet associated with it but can contain multiple subnets. When you configure a bridge domain with multiple subnets, the first subnet added becomes the primary IP address on the SVI interface. Subsequent subnets are configured as secondary IP addresses. When the switch reloads, the primary IP address can change unless it is marked explicitly.

One or more EPGs can be associated with each bridge domain. EPGs within the same bridge domain may be configured to talk to each other, but they do not have layer 2 adjacency enabled by default.

Bridge domains in Cisco Application Centric Infrastructure (ACI) have several configuration options to allow the administrator to tune the operation in various ways. To learn more about the various options, see Cisco Application Centric Infrastructure Fundamentals Guide.

Note Once a bridge domain is configured, its mode cannot be switched.

A bridge domain must be linked to a Virtual Routing and Forwarding (VRF).

Subnets

A subnet defines the IP address range that can be used within the bridge domain. A bridge domain can contain multiple subnets, but a subnet is contained within a single bridge domain. The scope of a subnet can be public, private, or shared under a bridge domain or an EPG. See Adding a Subnet to a Bridge Domain, on page 50.

DHCP Relay Labels

DHCP Relay is required only when the DHCP server is in a different EPG or private network than the clients. DHCP label associates the provider DHCP server with the bridge domain. The DHCP label object also specifies the owner. If your infrastructure requires DHCP relay labels, see Adding a DHCP Relay Label to a Bridge Domain, on page 51.

Note The bridge domain DHCP label must match the DHCP Relay name. Label matching enables the bridge domain to consume the DHCP Relay.

Adding a Bridge Domain to VRF

A bridge domain is a unique Layer 2 forwarding domain that contains one or more subnets. Each bridge domain must be linked to a VRF.

Before you begin

Create a Tenant for your customer, organization, or domain and configure your private network.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Bridge Domains**.

Step 7 Click **Add**.

Step 8 On the **Add Tenant Bridge Domain** screen, complete the following fields:

a) Add a unique name and description for the Bridge Domain.

b) Click **Select** and check the tag you want to use for the account.

c) From the **Type** drop-down list, choose regular or FC based on your requirement.

d) Check **Advertise Host Routes** to enable the host route.

e) Click **Select** and check the network you want to use for the account.

This is the virtual routing and forwarding (VRF) object associated with the tenant for which this bridge domain is created. It is also known as context or private network.

f) Click **Select** and check the VRF you want to use for the account.

g) From the **Forwarding** drop-down list, choose the forwarding parameter from the following options:

This sets the forwarding capacity between Layer 2 and Layer 3 networks. The values can be:

- **Optimize**—Automatically sets the Unicast and ARP parameters. Selects options: Hardware Proxy for L2 Unknown Unicast and Flood for Unknown Multicast Flooding with Unicast Routing enabled.
- **Custom**—Reveals the Unicast and ARP selections for custom configuration. If you choose custom forwarding, then complete the following additional parameters:

  1. From the **L2 Unknown Unicast** drop-down list, select the unicast parameter. The values can be **Flood** or **Hardware Proxy**.

     The default is Hardware Proxy. If enabled, unicast traffic flooding is blocked at a specific port, only permitting egress traffic with MAC addresses that are known to exist on the port. If you chose Flood, it floods the unicast traffic to all Layer 2 ports.

  2. From the **L3 Unknown Multicast Flooding** drop-down list, select the multicast parameter. The values can be **Flood** or **Optimized Flood**.

  3. From the **Multicast Destination Flooding** drop-down list, select the multicast destination flooding. The values can be **Flood in BD**, **Drop**, or **Flood in Encapsulation**.

  4. Check **ARP Flooding** to configure the flooding for the bridge domain.

This enables ARP flooding, so that the Layer 2 broadcast domain maps IP addresses to the MAC addresses. If flooding is disabled, unicast routing is performed on the target IP address.

EP Move Detection Mode check box is displayed only when ARP Flooding is enabled.

5. Check **Unicast Routing** to configure the bridge domain routing. This forwarding method is based on predefined forwarding criteria (IP or MAC address). The default is layer 3 forwarding (IP address).
6. Enter the virtual address in the **Virtual MAC Address** field.

h) Click **Select** and choose an endpoint retention policy that you want to use.
i) Check **Custom BD MAC Address** to configure the bridge domain MAC address and enter the address in the **MAC Address** field.

By default, a bridge domain takes the fabric wide default MAC address of 00:22:BD:F8:19:FF. Configure this property to override the default address.

j) By default, the **IP Data-plane Learning** drop-down list is set to true to enable data-plane IP learning on remote and local leaf switches.
k) The **Limit IP Learning to Subnet** drop-down list appears only when the **IP Data-plane Learning** drop-down list is set to true. By default, the value of the **Limit IP Learning to Subnet** drop-down list is set to true. If this option is set to true, the fabric will learn only IP addresses for subnets configured on the bridge domain.
l) Click **Select** and check the IGMP snoop policy you want to use for this tenant.

This policy inspects the IGMP membership report messages from interested hosts. It limits the multicast traffic to the subset of VLAN interfaces on which the hosts reside.

m) Click **Select** and check the MLDP snoop policy you want to use for this tenant.
n) Click **Select** and check the L3 out interface that you want to assign to this tenant.

This is the name of the Layer 3 outside interface associated with this object.

o) Click **Select** and check the route profile of L3 out network.

L3 Out is the network outside the fabric, configured for the tenant consuming this bridge domain, that is reachable by a specific route to external networks of a tenant application. The route profile specifies policies for external networks.

p) Click **Select** and check the ND Policy.
q) Click **Select** and check the monitoring policy associated with the tenant.
r) Click **Select** and check the First Hop Security (FHS) policy associated with the tenant.
s) Check **Optimize WAN Bandwidth** to optimize the WAN bandwidth .

Step 9 Click **Submit**.

Adding a Subnet to a Bridge Domain

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Bridge Domains**.

Step 7 On the **Bridge Domains** page, choose the row with the domain to which you want to add the subnet and click **View Details**.

Step 8 Click **Subnet**.

Step 9 On the **Subnet** page, click **Add**.

Step 10 On the **Add Subnet to Tenant Bridge Domain** screen, complete the following fields:

a) From the **IP Type** drop-down list, choose **IPv4** or **IPV6**.
b) In the **Gateway IP (Address)** field, enter the IP address of the default gateway.
c) In the **Gateway IP (Prefix)** field, enter a prefix in the range of 1-32 that starts with "/x" for IPV4 and 1-128that starts with "/x" for IPV6.
d) Check the **Shared Subnet** check box to share the subnet with multiple contexts (VRFs) in the same tenant or across tenants as part of a shared service.

   Shared subnets must be unique across the contexts (VRF) involved in the communication. When a subnet under an EPG provides a Layer 3 external network shared service, such a subnet must be globally unique within the entire ACI fabric.

e) Check the **Treat as virtual IP address** check box to treat the given IP as virtual address.
f) Check the **Make this IP address as primary** check box to make this IP address as primary for the subnet.
g) Check the **Public Subnet** check box to export it to a routed connection.
h) Check the **Private Subnet** check box to apply the subnet only within its tenant.
i) Check the **Scope (Shared between VRFs)** check box to enable subnet control for a shared subnet.
j) Check the **Scope (Advertised Externally)** check box to enable subnet control for a public subnet.
k) Check the **Scope (Private to VRF)** check box to enable subnet control for a private subnet.
l) Enter a description for the subnet.
m) Check the **Subnet Control (No Default SVI Gateway)** check box to enable the No Default SVI Gateway.
n) Check the **Subnet Control (Querier IP)** check box to apply specific protocol to the subnet. Querier IP enables IGMP Snooping on the subnet.
o) Click **Select** and check the L3 out for route profile that you want to use for the bridge domain.

   This is the Layer 3 Outside Network (L3extOut) configured for the tenant consuming this bridge domain.

p) Click **Select** and check the route profile that you want to use for this bridge domain.

   The route profile specifies policies for external networks.

q) Click **Submit**.

Adding a DHCP Relay Label to a Bridge Domain

DHCP Relay is required when the DHCP server is in a different EPG or private network than the clients. A DHCP relay label contains a name for the label, the scope, and a DHCP option policy. The scope is the owner of the relay server and the DHCP option policy supplies DHCP clients with configuration parameters such as domain, nameserver, and subnet router addresses.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Bridge Domains**.

Step 7 Click the row with the domain to which you want to add the DHCP label and click **View Details**.

Step 8 Click **DHCP Relay Label**.

Step 9 On the **DHCP Relay Label** page, click **Add**.

Step 10 On the **Add DHCP Label To Tenant Bridge Domain** screen, complete the following fields:

a) From the **Scope** drop-down list, choose the scope. Options are:

- **Infra**—The owner is the infrastructure.
- **Tenant**—The owner is the tenant.

The default is **Infra**.

b) Click **Select** and check the DHCP relay policy that you want to use for the tenant bridge domain.

c) Click **Select** and check the DHCP option policy that you want to use.

d) Click **Submit**.

Creating an ND RA Prefix Policy

This section provides a procedure to create Neighbor Discovery Router Advertisement (ND RA) Prefixes for Layer 3 interfaces. ND RA prefix policies are deployed for individual subnets. Every bridge domain can have multiple subnets, and each subnet can have a different ND RA prefix policy.

An ND RA configuration applies only to IPv6 prefixes.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **ND RA Prefix Policy**.

Step 7 Click **Add**.

Step 8 On the **Create ND RA Prefix Policy** screen, complete the following fields:

a) Enter a unique name and description for the ND RA prefix policy. The name can be up to 64 alphanumeric characters. The description can be up to 128 alphanumeric characters.

b) Check the **Auto Configuration** check box to enable the auto-configuration and use the configured prefix in stateless address allocation.

c) Check the **On Link** check box to enable on link and indicate that the prefix carried in the RA message received by the host on the local link can be allocated to the local link.

d) Check the **Router Address** check box to enable router address. By default, this check box is unchecked.

e) Enter any value in the range of 0 to 4294967295 seconds, as the valid lifetime of the prefix. The default value is 2592000.
f) Enter any value in the range of 0 to 4294967295 seconds, as the preferred lifetime of the prefix. The preferred lifetime cannot be bigger than the valid lifetime. The default value is 604800.

Step 9 Click **Submit**.

Creating an Endpoint Retention Policy

The endpoint retention policy is configured for the bridge domain to send ARP requests (for IPv4) and neighbor solicitations (for IPv6) at 75 percent of the local endpoint aging interval. The policy is used for tracking IP addresses on an endpoint.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **End Point Retention**.
Step 7 Click **Add**.
Step 8 On the **Create Endpoint Retention Policy** screen, complete the following fields:
a) Enter a unique name and description for the endpoint retention policy.
b) Enter the hold interval for endpoints in seconds. The hold interval must be in the range of 5 to 65535. The default value is 300.
c) Enter the bounce entry aging interval for endpoints in seconds. The aging interval must be in the range of 150 to 65535. The default value is 630.
d) Enter the local endpoint aging interval for endpoints in seconds. The aging interval must be in the range of 120 to 65535. The default value is 900.
e) Enter the remote endpoint aging interval for endpoints in seconds. The aging interval must be in the range of 120 to 65535. The default value is 300.
f) Enter the move frequency for endpoints in seconds. The move frequency must be in the range of 0 to 65535. The default value is 256.

Step 9 Click **Submit**.

Application Profiles

Application profiles are logical containers that define the policies, services, and relationships between End Point Groups (EPGs). Each application profile contains one or more EPG that can communicate with the other EPGs in the same application profile, and with EPGs in other application profiles according to the contract rules. At minimum, associate one application profile with one EPG.

Modern applications contain multiple components. An application profile models the requirements of an application. For example, an e-commerce application could require a web server, a database server, data

located in a storage area network, and access to outside resources that enable financial transactions. The application profile contains as many (or as few) EPGs as necessary that are logically related for the e-commerce application.

EPGs can be organized according to one of the following:

- The application they provide (such as sap in the example in Appendix A).
- The function they provide (such as infrastructure).
- Where they are in the structure of the data center (such as DMZ).
- Whatever organizing principle that a fabric or tenant administrator chooses to use.

Creating an Application Profile for the Tenant

The application profile is a set of requirements that an application instance has on the virtualized fabric. The policy regulates connectivity and visibility among endpoints within the scope of the policy.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click **Add**.

Step 8 On the **Add Tenant Application Profile** screen, complete the following fields:

a) Add a unique name, description, and an alias for the Application Profile.

b) Click **Select** and check the tag name that you want to use for the APIC account.

c) From the **QoS Class** drop-down list, choose from the following options for the priority class:

- **Unspecified**—This is the default value.
- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value

d) Click **Select** and check the monitoring policy associated with the tenant.

When you apply a monitoring policy, it overrides the default monitoring policy.

Step 9 Click **Submit**.

Endpoint Groups

An Endpoint Group (EPG) is a logical container of endpoints that have common policy requirements such as security, virtual machine mobility (VMM), QoS, or Layer 4 to Layer 7 services. They have an address (identity), a location, attributes (such as version or patch level), and can be physical or virtual. Knowing the address of an endpoint enables access to all its other identity details. Rather than configure and manage endpoints individually, they are placed in an EPG and are managed as a group.

The ACI fabric can contain the following types of EPGs:

- Application endpoint group
- Layer 2 external outside network instance
- Layer 3 external outside network instance
- Management endpoint groups for out-of-band or in-band access

By default, all endpoints in the same endpoint group can talk to each other without requiring a contract. Intra-endpoint group (intra-EPG) isolation prevents all endpoints in an EPG from talking to each other but inter-EPG communication is still permitted if there is a contract. This is similar to a private VLAN. For example, assume that you have three endpoints: two are in the client endpoint group, while the other endpoint is in the Web endpoint group. If there is a contract between endpoint groups, they can talk to each other.

Regardless of how an EPG is configured, EPG policies are applied only to the endpoints they contain. For example, to configure a WAN router connectivity to the fabric, you configure an EPG that includes any endpoints within the associated WAN subnet. The fabric learns of the endpoints through a discovery process and applies the policies accordingly.

After creating an EPG, add a static path to the EPG to determine the port and leaf/node for the traffic. See Adding a Static Path to EPG, on page 58.

You can also add static nodes (leaf, spine, or APIC), and domains (physical, VMM, L3, or L3 external - see examples) to EPGs and define how and when they are deployed. See Adding a Static Node to EPG, on page 60 and Adding a Domain to an EPG, on page 56.

Adding an EPG

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the profile that you want to update and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click **Add**.

Step 10 On the **Add Tenant EPG** screen, complete the required fields including the following:

a) Add a unique name, description, and alias for the EPG.

b) (Optional) Click **Select** and check the tag you want to use.
c) From the **QoS** drop-down list, choose one of the following options for the priority class:

- **Unspecified**—This is the default value.
- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value
- **Custom**—Complete the additional parameter

d) If you chose custom QoS, click **Select** and check the customized quality of service (QoS) class that you want to use for the EPG.
e) Click **Select** and check the bridge domain for the EPG.
f) Click **Select** and check the monitoring policy associated with the tenant.

When you apply a monitoring policy, it overrides the default monitoring policy.

g) Click **Select** and check the data plane policy that you want to apply for the EPG.
h) Choose **Enforced** or **Unenforced** to prevent communication between endpoint devices within the same base EPG, accordingly. On choosing **Enforced**, the **Forwarding Control** drop-down list appears. Choose **True** from the **Forwarding Control** drop-down list to enable proxy ARP.
i) Choose **Include** to mark the EPG as the preferred group member.
j) Choose **Enable** to enable the flood on encapsulation, so that the EPG flooding traffic does not reach the other EPG.
k) Click **Select** and check the First Hop Security (FHS) trust control policy that you want to apply for the EPG.
l) Click **Submit**.

Adding a Domain to an EPG

An EPG is associated with domains by being linked to a domain profile, which can be a VMM, physical, Layer 2 external, or Layer 3 external domain.

Before you begin

Create a physical, VMM, Layer 3, or Layer 2 domain for the APIC account.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Application Profile**.
Step 7 Click the row with the profile that you want to update and click **View Details**.
Step 8 Click **EPG**.
Step 9 Click the row with the EPG that you want to update and click **View Details**.
Step 10 Click **Domain**.

Step 11 Click **Add**.

Step 12 On the **Add Domain To EPG** screen, complete the required fields, including the following:

a) Click **Select** and check the domain profile that you want to add to the EPG.

b) From the **Deploy Immediacy** drop-down list, choose one of the following options:

- **Immediate**—Specifies that the policy is programmed in the hardware policy CAM as soon as it is downloaded in the leaf software.
- **On Demand**—Specifies that the policy is programmed in the hardware policy CAM only when the first packet is received through the data path. This process helps to optimize the hardware space.

c) From the **Resolution Immediacy** drop-down list, choose one of the following options:

It specifies whether policies are resolved immediately or when needed.

- **Immediate**—Specifies that EPG policies (including contracts and filters) are downloaded to the associated leaf switch software upon hypervisor attachment to VDS. LLDP or OpFlex permissions are used to resolve the hypervisor to leaf node attachments.
- **On Demand**—Specifies that a policy (for example, VLAN, VXLAN bindings, contracts, or filters) is pushed to the leaf node only when a hypervisor is attached to VDS and a VM is placed in the port group (EPG).
- **Pre-provision**—Specifies that a policy (for example, VLAN, VXLAN binding, contracts, or filters) is downloaded to a leaf switch even before a hypervisor is attached to the VDS. Therefore, this option pre-provisions the configuration on the switch.

d) **Port Binding**—Choose one of the following options to configure a default port binding type to be applied automatically to all new vEthernet port profiles unless explicitly configured:

- **Dynamic Binding**—To assign a DVPortID to a virtual machine only when the virtual machine is powered on and its NIC is in a connected state. In the **Number of Ports** field that appears on choosing **Dynamic Binding**, enter the number of ports available in the environment. Dynamic binding can be used in environments where you have more virtual machines than available ports but do not plan to have a greater number of virtual machines active than available ports.
- **Ephemeral**—To assign a new DVPortID to the port every time the VM is powered on. The port keeps this same DVPortID while the VM is up. All available distributed virtual switch (DVS) ports are shared. Ports are not allocated from the port group pool.
- **Default**
- **Static Binding**—To assign a DVPortID from the port group pool when you first assign the port group to the port. The DVPortID persists for the life of the network adapter. The port group has a fixed number of ports. From the **Port Allocation** drop-down list that appears on choosing **Static Binding**, choose **Fixed** to allocate port at the time of creation only or **Elastic** to elastically increase or decrease port numbers depending on the number of ports needed. In the **Number of Ports** field that appears on choosing **Static Binding**, enter the number of ports available in the environment.

e) From the **Allow Promiscuous** drop-down list, choose from the following options:

It enables all packets to pass to the VMM domain, which is often used to monitor network activity.

- **Reject**—Packets that do not include the network address are dropped.
- **Accept**—All traffic is received within the VMM domain.

f) From the **Forged Transmits** drop-down list, choose one of the following options:

- **Reject**—All non-matching frames are dropped.
- **Accept**—Non-matching frames are received.

It specifies whether to allow forged transmits. A forged transmit occurs when a network adapter starts sending out traffic that identifies itself as something else. This security policy compares the effective address of the virtual network adapter and the source address inside an 802.3 Ethernet frame generated by the virtual machine to ensure that they match.

g) From the **MAC Changes** drop-down list, choose one of the following options:

- **Reject**—Does not allow new MAC addresses.
- **Accept**—Allows new MAC addresses.

It enables you to define new MAC addresses for the network adapter within the virtual machine (VM).

h) The following fields appear when the VMM type domain profile is chosen:

- In the **Delimiter** field, enter one of the following: **|, ~, !, @, ^, +,** or **=**. If you do not enter a symbol, the system default **|** delimiter will appear in the domain name.
- Choose **True** from the drop-down list to allow Micro-Segmentation which enables automatic assignment of endpoints to EPGs. The default value is unspecified.
- Choose **Static** or **Dynamic** as the VLAN mode. On choosing **Static**, enter the encapsulation for the port. For example, vlan-1. Packets sent out of the port are tagged in accordance with the specified encapsulation.
- From the **Switching Mode** drop-down list, choose **AVE** or **native** as the switching mode. On choosing **native**, the **Enhanced Lag Policy** field appears. Click **Select** and choose an enhanced lag policy that you want to add to the VMM domain.
- Choose **Auto**, **VLAN**, or **VXLAN** as the encapsulation mode. The **Encap Mode** field is disabled when you choose **native** as the switching mode type.
- Choose **Enable** from the **NetFlow** drop-down list, to monitor IP packets that are passing through the ports.

Step 13 Click **Submit**.

Adding a Static Path to EPG

Static path policies provide a summary of the configured properties of the policy, fault counts, and history for the static path. Configure the static path to the destination EPG.

Note When an EPG uses a static binding path, the encapsulation VLAN associated with this EPG must be part of a static VLAN pool.

Before you begin

Contracts must be configured.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the profile that you want to update and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click the row with the EPG that you want to update and click **View Details**.

Step 10 Click **Static Path**.

Step 11 Click **Add**.

Step 12 On the **Add Static Path To EPG** screen, complete the following fields:

a) From the **Path Type** drop-down list, choose from the following options:

- **Port**—Is the default value
- **Direct Port Channel**—Class 1 Differentiated Services Code Point (DSCP) value
- **Virtual Port Channel**—Class 2 DSCP value

b) Click **Select** and check the static path that you want to add to the EPG. The static paths are displayed according to the path type.

c) In the **Encapsulation** field, enter a VLAN value that is part of one of the static VLAN blocks associated with the domain. The encapsulation supports VXLAN (virtual extensible LAN) and QinQ technologies. For example, the value of encapsulation can be vlan-1, vxlan-50102, or qinq-123-213.

d) From the **Deployment Immediacy** drop-down list, choose from the following options:

- **Immediate**—Specifies that the policy is programmed in the hardware policy CAM as soon as it is downloaded in the leaf software.
- **On Demand**—Specifies that the policy is programmed in the hardware policy CAM only when the first packet is received through the data path. This process helps to optimize the hardware space.

Once EPG policies are downloaded to the leaf software, you can choose the deployment preferences to specify when a policy is pushed into leaf switches.

e) From the **Mode** drop-down list, choose the static association from the following options:

EPG tagging refers to configuring a static path under an EPG.

- **Tagged**—Select this mode if the traffic from the host is tagged with a VLAN ID.
- **Untagged**—Select this mode if the traffic from the host is untagged (without VLAN ID).

  When a leaf switch is configured for an EPG to be untagged, for every port this EPG uses, the packets exit the switch untagged.

Note When an EPG is deployed as untagged, do not deploy that EPG as tagged on other ports of the same switch.

- **802.1P Untagged**—Select this mode if the traffic from the host is tagged with a 802.1P tag. When an access port is configured with a single EPG in native 802.1p mode, its packets exit that port untagged. When an access port is configured with multiple EPGs, one in native 802.1p mode, and some with VLAN tags, all packets exiting that access port are tagged VLAN 0 for EPG configured in native 802.1p mode and for all other EPGs packets exit with their respective VLAN tags.

  Note Only one native 802.1p EPG is allowed per access port.

f) In the **Primary VLAN for Micro-Seg** field, enter the encapsulation for a primary VLAN.
g) Click **Submit**.

Adding a Static Node to EPG

Before you begin

Create nodes in the APIC system.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Application Profile**.
Step 7 Click the row with the profile that you want to update and click **View Details**.
Step 8 Click **EPG**.
Step 9 Click the row with the EPG that you want to update and click **View Details**.
Step 10 Click **Static Node**.
Step 11 Click **Add**.
Step 12 On the **Add Static Node To EPG** screen, complete the following fields:

a) Click **Select** and check the node that you want to add to the EPG.
b) In the **Encapsulation** field, enter a VLAN value that is part of one of the static VLAN blocks associated with the domain. The encapsulation supports VXLAN (virtual extensible LAN) and QinQ technologies. For example, the value of encapsulation can be vlan-1, vxlan-50102, or qinq-123-213.
c) From the **Mode** drop-down list, choose the static association from the following options:

- **Trunk**—Traffic for the EPG is sourced by the leaf switch with the specified VLAN tag.
- **Access (802.1p)**—If only one EPG is bound to that interface, the behavior is identical as in the untagged case. If other EPGs are associated with the same interface, traffic for the EPG is sourced with an IEEE 802.1q tag using VLAN 0 (IEEE 802.1p tag), or is sourced as untagged in the case of EX switches.

  Access (Untagged)—Traffic for the EPG is sourced by the leaf as untagged. Traffic received by the leaf switch as untagged or with the tag specified during the static binding configuration is associated with the EPG.

d) From the **Deployment Immediacy** drop-down list, choose the policy from the following options:.

- **Immediate**—Specifies that the policy is programmed in the hardware policy CAM as soon as the policy is downloaded in the leaf software.
- **Lazy**—Specifies that the policy is programmed in the hardware policy CAM only when the first packet is received through the data path. This process helps to optimize the hardware space.

Once EPG policies are downloaded to the leaf software, you can choose the deployment preferences to specify when a policy is pushed into leaf switches.

e) Click **Submit**.

Adding an IGMP Snoop Static Group to Static Path

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Application Profile**.
Step 7 Click the row with the application profile that you want to update and click **View Details**.
Step 8 Click **EPG**.
Step 9 Click the row with the EPG that you want to update and click **View Details**.
Step 10 Click **Static Path**.
Step 11 Click the row with the static path to which you want to add an IGMP snoop static group and click **View Details**.
Step 12 Click **IGMP Snoop Static Group**.
Step 13 Click **Add**.
Step 14 On the **Add IGMP Snoop Static Group to Static Path** screen, complete the following fields:
a) Enter the multicast IP address as group address that has to be associated with an IGMP static group class map.
b) Enter wildcard or a valid IP address as source address that has to be associated with an IGMP static group class map.
Step 15 Click **Submit**.

Adding an EPG Contract Master

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the application profile that you want to update and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click the row with the EPG that you want to update and click **View Details**.

Step 10 Click **EPG Contract Master**.

Step 11 On the **Add EPG Contract Master** screen, complete the following fields:

a) Click **Select** and check the application profile that you want to use in the EPG contract.

b) Click **Select** and check the EPG name that you want to use in the EPG contract.

Step 12 Click **Submit**.

Contracts

EPGs can only communicate with other EPGs according to contract rules. Contracts determine the type of traffic that can pass between EPGs, including the protocols and ports allowed. If there is no contract, inter-EPG communication is disabled by default. There is no contract required for intra-EPG communication. A contract contains one or more subjects.

An EPG can both provide and consume the same contract. An EPG can also provide and consume multiple contracts simultaneously.

Contract Subjects

A subject is a sub-application running behind an endpoint group. A contract can encapsulate multiple subjects. An EPG associates with a subject and defines rules under the association for consuming, providing, or for peer-to-peer communications to that subject. Subjects contain filters and optional labels.

Export Contract feature enables you to export the XML or JSON code for later use with the REST API.

Provider and Consumer Contracts

Contracts can contain multiple communication rules and multiple endpoint groups. The relationship between an EPG and a contract can be either a provider or consumer. When an EPG provides a contract, communication with that EPG can be initiated from other EPGs as long as the communication complies with the provided contract. When an EPG consumes a contract, the endpoints in the consuming EPG may initiate communication with any endpoint in an EPG that is providing that contract.

Filters

Filters enable you to specify the protocols you want to permit for traffic management between two EPGs. For example, you may want to permit only `https` traffic. There are several types of filters.

- Permit—It allows traffic.
- Deny (Taboo)—For specific use cases. You may specify to allow all traffic in a contract, but set up taboos to deny certain traffic.
- Redirect—Useful to send traffic from an EPG to a layer 4-7 device such as a firewall, load balancer, or IPS/IDS.

- Mark—To mark traffic for Quality of Service reasons.

You can add filters to a contract by adding filter chains (consumer or provider) to contract subjects.

Contract Labels

Labels are optional advanced identifiers. When you use labels, you can specify more complex relationships between EPGs. Labels allow for control over which subjects and filters to apply when communicating between a specific pair of endpoint groups. Without labels, a contract applies every subject and filter between consumer and provider endpoint groups. You can use labels to represent a complex communication scenario, within the scope of a single contract, then reuse this contract while specifying only a subset of its policies across multiple endpoint groups.

Taboo Contracts

A Taboo contract provides a way for an EPG to specify the subjects on which communication is not allowed.

Creating Contracts

Without a contract, the default forwarding policy is to not allow any communication between EPGs but all communication within an EPG is allowed.

Note If two tenants are participating in same contract, ensure that they are not able to see each other and that their endpoint groups are not able to communicate.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Contracts**.

Step 7 Click **Add**.

Step 8 On the **Add Tenant Contract** page, complete the following fields:

a) Add a unique name for the contract.

If you create contracts under the common and user tenants, that are consumed by the same tenant, they must have different names.

b) Enter an alias name for the contract. While the contract name cannot be changed after creation, the alias name of the contract can be changed as required.

c) From the **Scope** drop-down list, choose from the following options:

- **Application Profile**—The contract is applied to endpoint groups in the application profile.
- **Context**—The contract is applied to endpoint groups in the same Virtual Routing and Forwarding (VRF).
- **Global**—This contract is applied to endpoint groups throughout the fabric.

- **Tenant**—This contract is applied to endpoint groups within the same tenant.

d) From the **Priority** drop-down list, choose the priority level of the service contract. Each level represents the respective class of Differentiated Services Code Point (DSCP) value. The default option is **Unspecified**.
e) Choose a DSCP target to which the DSCP range has to be mapped. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.
f) Add a description for the contract.
g) Click **Select** and check the tag you want to use.

Step 9 Click **Submit**.

What to do next

Create contract subjects to specify the information that can be communicated and the mechanism of communication.

Creating a Contract Subject

A subject is a sub-application running behind an endpoint group. A contract can encapsulate multiple subjects. An endpoint group always associates with a subject and defines rules under the association for consuming, providing, or for peer-to-peer communications to that subject.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Contracts**.
Step 7 Click the row with the contract that you want to update and click **View Details**.
Step 8 Click **Contract Subjects**.
Step 9 Click **Add**.
Step 10 On the **Add Tenant Contract Subject** page, complete the required fields, including the following:

a) Add a unique name and description for the contract subject.
b) Check **Reverse Filter Ports** to apply the same subject rule to the reverse filter ports when the contract applies in both directions. This field is disabled when you uncheck the **Apply Both Directions** check box.
c) Check **Apply Both Directions** to apply the contract to both inbound and outbound traffic. If the selected contract does not apply to both, then the filter chain must be configured for consumer to provider and provider to consumer separately.

If you uncheck the **Apply Both Directions** check box, complete the following fields for the in term and out term properties:

- **Service Graph**—Click **Select** to choose the service graph that you want to use.
- **QoS Priority**—Choose the traffic priority for the associated EPG.

- **Target DSCP**—Choose a target that changes the DSCP (Differentiated Services Field) marks inside a packet. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.

d) Click **Select** and check the box for the service graph that you want to add to the contract.

The service graph is an image that shows the relationship between contracts and subjects.

e) Click **Select** and choose a service graph for the filter. This field appears only when the **Apply Both Directions** check box is checked.

f) From the **QoS** drop-down list, choose the traffic priority for the associated EPG from the following options:

Each system class manages one lane of traffic.

- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value
- **Unspecified**—This is the default value.

The default option is **Unspecified**.

Step 11 Click **Submit**.

What to do next

Create consumer and provider contracts.

Adding Contracts to EPGs

Provided Contracts

EPGs can only communicate with other EPGs according to contract rules. Contracts determine the types of traffic that can pass between EPGs, including the protocols and ports allowed.

The relationship between an EPG and a contract can be either a provider or consumer. When an EPG provides a contract, communication with that EPG can be initiated from other EPGs as long as the communication complies with the provided contract. When an EPG consumes a contract, the endpoints in the consuming EPG may initiate communication with any endpoint in an EPG that is providing that contract.

An EPG can both provide and consume the same contract. An EPG can also provide and consume multiple contracts simultaneously.

Adding a Provided Contract to an EPG

You can associate contracts that were created earlier to create policy relationships between the EPGs. A provided contract is a contract for which the EPG is a provider.

Note Verify that both provided and consumed contracts have the same name.

Before you begin

Contracts must be configured.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the profile that you want to update and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click the row with the EPG that you want to update and click **View Details**.

Step 10 Click **Consumed Contract**.

Step 11 Click **Add**.

Step 12 On the **Add Provided Contract To EPG** screen, complete the fields including the following:

a) Click **Select** and check the contract that you want to add to the EPG.

b) From the **QoS** drop-down list, choose the traffic priority for the associated EPG from the following options:

Each system class manages one lane of traffic.

- **Unspecified**—This is the default value.
- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value

c) Enter the label for the provided contract.

d) Enter the subject label for the provided contract.

e) Click **Submit**.

Consumed Contracts

Also need to look into Taboo Contract and Filters.

Adding a Consumed Contract to an EPG

You can associate contracts that were created earlier to create policy relationships between the EPGs. A consumed contract is a contract for which the EPG is a consumer.

Note Verify that both provided and consumed contracts have the same name.

Before you begin

Contracts must be configured.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the profile that you want to update and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click the row with the EPG that you want to update and click **View Details**.

Step 10 Click **Consumed Contract**.

Step 11 Click **Add**.

Step 12 On the **Add Consumed Contract To EPG** screen, complete the fields including the following:

a) Click **Select** and check the contract that you want to add to the EPG.

b) From the **QoS** drop-down list, choose the traffic priority for the associated EPG from the following options:

Each system class manages one lane of traffic.

- **Unspecified**—This is the default value.
- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value

c) Enter the label for the consumed contract.

d) Enter the subject label for the consumed contract.

e) Click **Submit**.

Adding a Consumed Contract Interface

Before you begin

Contracts must be configured.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Application Profile**.
Step 7 Click the row with the profile that you want to update and click **View Details**.
Step 8 Click **EPG**.
Step 9 Click the row with the EPG that you want to update and click **View Details**.
Step 10 Click **Consumed Contract Interface**.
Step 11 Click **Add**.
Step 12 On the **Add Consumed Contract Interface To EPG** screen, complete the fields including the following:

a) Click **Select** and check the contract interface that you want to add to the EPG.
b) From the **Priority** drop-down list, choose a priority for the selected EPG from the following options:

- **Level1**—Class 1 Differentiated Services Code Point (DSCP) value.
- **Level2**—Class 2 DSCP value
- **Level3**—Class 3 DSCP value
- **Unspecified**—Is the default value

c) Click **Submit**.

Contract Labels

Adding a Consumed Label to a Contract Subject

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Contracts**.
Step 7 Click the row with the contract that you want to update and click **View Details**.
Step 8 Click **Contract Subjects**.
Step 9 Click the row with the contract subject that you want to update and click **View Details**.
Step 10 Click **Consumed Label**.
Step 11 Click **Add**.
Step 12 On the **Add Consumed Label to Contract To Contract Subject** screen, complete the following fields:

a) From the **Match Type** drop-down list, choose the subject match criteria across providers, from the following options:

- **Atleast One**
- **Atmost One**
- **None**
- **All**

b) Enter a **Label Name**.

A subject label is used as classification criteria for subjects being consumed by the EPGs participating in the contract.

c) From the **Label Tag** drop-down list, choose a tag.

It is the search keyword or term that is assigned to the application profile. A tag allows you to group multiple objects by a descriptive name. You can assign the same tag name to multiple objects and you can assign one or more tag names to an object.

d) Check **Complement** for the contract to take effect if the labels do not match.

e) Click **Submit**.

Adding a Provided Label to a Contract Subject

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Contracts**.

Step 7 Click the row with the contract that you want to update and click **View Details**.

Step 8 Click **Contract Subjects**.

Step 9 Click the row with the contract subject that you want to update and click **View Details**.

Step 10 Click **Provided Label**.

Step 11 Click **Add**.

Step 12 On the **Add Provided Label to Contract To Contract Subject** screen, complete the following fields:

a) From the **Match Type** drop-down list, choose the subject match criteria across providers, from the following options:

- **Atleast One**
- **Atmost One**
- **None**
- **All**

b) Enter a **Label Name**.

A subject label is used as classification criteria for subjects being consumed or provided by the EPGs participating in the contract.

c) From the **Label Tag** drop-down list, choose a tag.

It is the search keyword or term that is assigned to the application profile. A tag allows you to group multiple objects by a descriptive name. You can assign the same tag name to multiple objects and you can assign one or more tag names to an object.

d) Check **Complement** for the contract to take effect if the labels do not match.
e) Click **Submit**.

Fabric Extender (FEX)

Fabric Extender (FEX) behave as a remote line card for a parent switch. The FEX is an extension of the parent switch fabric, with the FEX and the parent switch together form a distributed modular system. This means that the FEXs are completely managed from the parent switch and appear as physical ports on that switch.

Adding a FEX Profile

A FEX profile enables you to define policy for the ports facing hosts on the FEX and configure FEX interfaces.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **FEX Profile**.
Step 5 Click **Add**.
Step 6 Enter the name and description of the profile used for configuring FEX.
Step 7 Click **Submit**.

What to do next

You have to add the interface port selector to the FEX profile to identify the interfaces between the node and the host.

Adding an Access Port Selector to the FEX Profile

Access port selector is used for identifying the interfaces between the node and the host (such as hypervisor), which consume the policies in the interface policy group.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **FEX Profile**.

Step 5 Click the row with the FEX profile to which you want to add an access port selector and click **View Details**.

Step 6 In the **FEX Profile Access Port Selectors** tab, click **Add**.

Step 7 On the **Create Fex Profile Access Port Selector** screen, complete the following fields:

- Enter the name and description of the interface selector. We recommend that you include information about where and when the policy must be used, in the description field.
- Enter the ID of the interfaces that consume the policies in the interface policy group. You can enter a single FEX interface, one or more interface ranges, or All.
- Click **Select** to view the list of available interface policy group. Check the interface policy group that you want the interfaces to consume and click **Select**.

Step 8 Click **Submit**.

What to do next

You can add access port blocks and sub-port blocks to the access port selector of the FEX profile. Choose an access port selector and click **View Details** to view the access port blocks and sub port blocks of the access port selector. Click **Add** under respective tabs to add the access port block and sub port block.

Fabric Ext Connection Policies

Before connecting a Cisco APIC cluster (fabric) in a Cisco ACI Multi-Site topology, you must configure the Dataplane Tunnel Endpoint (TEP) in the Fabric Ext Connection Policy for each fabric.

You can use the **Create Intrasite/Intersite Profile** screen to add connection details for APIC multipod, remote leaf switches connecting to the ACI fabric, and APIC sites managed by Cisco ACI Multi-Site. When the Multi-Site infrastructure has been configured, the Multi-Site system adds the Intersite Dataplane TEP to this APIC policy.

Adding a Fabric External Routing Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the infra tenant that you want to update and click **View Details**.

Step 6 Click **Fabric Ext Connection Policies**.

Step 7 Click the row with the fabric external connection policies that you want to update and click **View Details**.

Step 8 Click **Fabric External Routing Profile**.

Step 9 Click **Add**.

Step 10 On the **Add Fabric External Routing Profile** screen, complete the following fields:
a) Enter a unique name and description for the fabric external routing profile.
b) Enter a subnet or comma-separated subnets for the fabric external routing.

Step 11 Click **Submit**.

You can view the subnets of the fabric external routing profile under the **Subnets** tab (**APIC Account** > **Tenant(s)** > **Fabric Ext Connection Policies** > **Fabric External Routing Profile**).

Adding a Pod Connection Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the infra tenant that you want to update and click **View Details**.
Step 6 Click **Fabric Ext Connection Policies**.
Step 7 Click the row with the fabric ext connection policies that you want to update and click **View Details**.
Step 8 Click **Pod Connection Profile**.
Step 9 Click **Add**.
Step 10 On the **Add Pod Connection Profile** screen, complete the following fields:
a) Click **Select** and choose a pod to which you want to add the connection profile.
b) (Optional) The **Password** field appears only when you choose a virtual pod. Enter the BGP password for the connection profile.
c) Re-enter the password for confirmation.
d) The **Unicast TEP** field appears only when you choose a physical pod. Enter the unicast TEP IP address for the physical pod. The format of the IP address is IPv4/mask. The valid range of mask is from 1 to 32.
e) Enter the data plane TEP IP address. The format of the IP address is IPv4/mask. The subnet mask must be 32.

Step 11 Click **Submit**.

You can view the data plane TEP IP addresses of the pod connection profile under the **Data Plane TEP** tab (**APIC Account** > **Tenant(s)** > **Fabric Ext Connection Policies** > **Pod Connection Profile**).

Creating an Intrasite or Intersite Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the infra tenant that you want to update and click **View Details**.

Step 6 Click **Fabric Ext Connection Policies**.

Step 7 Click **Add**.

Step 8 On the **Create Intrasite/Intersite Profile** screen, complete the following fields:

a) A unique fabric ID is displayed.
b) Enter a unique name for the intrasite/intersite profile.
c) Enter the community name. The example of a community name format is:extended:as2-nn4:4:15. The numbers are variables.
d) Choose **Full Mesh** or **Route Reflector** as the peering type of the intrasite/intersite profile.
e) Enter the administrative password to access the site or pod peering profile.
f) Re-enter the password for confirmation.

Step 9 Click **Submit**.

Filter Chain

Adding a Filter Chain

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Contracts**.

Step 7 Click the row with the contract that you want to update and click **View Details**.

Step 8 Click **Contract Subject**.

Step 9 Click the row with the contract subject that you want to update and click **View Details**.

Step 10 Click **Filter Chain**.

Step 11 Click **Add**.

Step 12 On the **Add Filter to Contract Subject** screen, complete the following fields:

a) Click **Select** and check the filter that you want to use for the contract subject.
b) Check **Apply Both Directions** to apply the filter to both inbound and outbound traffic. If the selected filter does not apply to both, then the filter chain must be configured for consumer to provider and provider to consumer separately.

   On checking the **Apply Both Directions** check box, you have to define the following fields once for both inbound and outbound traffic. If you uncheck the **Apply Both Directions** check box, you have to define the filter and following fields for consumer to provider and provider to consumer separately:

   1. Check the **Log** check box to use log directive on filters in contract subject.
   2. Check the **No Stats** check box to prevent generation of statistics on a path.

3. Choose **Permit** or **Deny** from the drop-down list to accordingly apply the filter settings on traffic. When the **Deny** option is selected, the **Priority** drop-down list appears to set the priority to deny filtered traffic.

Step 13 Click **Submit**.

Adding a Filter Chain for Consumer to Provider

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Contracts**.
Step 7 Click the row with the contract that you want to update and click **View Details**.
Step 8 Click **Contract Subject**.
Step 9 Click the row with the contract subject that you want to update and click **View Details**.
Step 10 Click **Filter Chain For Consumer to Provider**.
Step 11 Click **Add**.
Step 12 On the **Add In Term Filter to Contract Subject** screen, complete the following fields:
a) Click **Select** and choose the filter that need to be applied to traffic from consumer to provider.
b) Check the **Log** check box to use log directive on in term filter.
c) Check the **No Stats** check box to prevent generation of statistics on the in term path.
d) Choose **Permit** or **Deny** from the drop-down list to accordingly apply the filter settings on in term traffic.
Step 13 Click **Submit**.

Adding a Filter Chain for Provider to Consumer

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Contracts**.
Step 7 Click the row with the contract that you want to update and click **View Details**.
Step 8 Click **Contract Subject**.
Step 9 Click the row with the contract subject that you want to update and click **View Details**.
Step 10 Click **Filter Chain For Provider to Consumer**.

Step 11 Click **Add**.

Step 12 On the **Add Out Term Filter to Contract Subject** screen, complete the following fields:

a) Click **Select** and choose the filter that need to be applied to traffic from provider to consumer.
b) Check the **Log** check box to use log directive on out term filter.
c) Check the **No Stats** check box to prevent generation of statistics on the out term path.
d) Choose **Permit** or **Deny** from the drop-down list to accordingly apply the filter settings on out term traffic.

Step 13 Click **Submit**.

Data Plane Policing

Use data plane policing (DPP) to manage bandwidth consumption on ACI fabric access interfaces. DPP policies can apply to egress traffic, ingress traffic, or both. DPP monitors the data rates for a particular interface. When the data rate exceeds user-configured values, marking or dropping of packets occurs immediately.

Policing does not buffer the traffic; therefore, the transmission delay is not affected. When traffic exceeds the data rate, the ACI fabric can either drop the packets or mark QoS fields in them.

Creating a Data Plane Policing

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Data Plane Policing**.

Step 7 Click **Add**.

Step 8 On the **Create Data Plane Policing** screen, complete the following fields:

a) Enter a unique name for the DPP.
b) Choose **enabled** from the drop-down list, to enable the administrative state of the DPP. The default value is disabled.
c) Choose **Bit Policer** or **Packet Policer** as the policer mode.
d) Choose one of the following as the traffic policer type:

- **1 Rate 2 Color**—To rate-limit a traffic flow to an average bits-per-second arrival rate.
- **2 Rate 3 Color**—To rate-limit a traffic flow to two rates and three traffic categories (green, yellow, and red).

e) Choose **Drop**, **Mark**, or **Transmit** as the action to be taken on categorizing a traffic flow as conforming (green).
f) In the **Conform mark cos** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the class of service (CoS) value for the traffic belonging to a specific class when the traffic flow is categorized as conforming.
g) In the **Conform mark dscp** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the differentiated services code point (DSCP) value for the traffic belonging to a specific class when traffic flow is categorized as conforming.

h) The **Exceed Action** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Choose **Drop**, **Mark**, or **Transmit** as the action to be taken when the traffic flow is exceeded.
i) The **Exceed mark cos** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the CoS value for the traffic belonging to a specific class when the traffic flow is exceeded.
j) The **Exceed mark dscp** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the DSCP value for the traffic belonging to a specific class when the traffic flow is exceeded.
k) Choose **Drop**, **Mark**, or **Transmit** as the action to be taken on categorizing a traffic flow as nonconforming (red).
l) In the **Violate mark cos** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 6 to set the CoS value for the traffic belonging to a specific class on traffic violation.
m) In the **Violate mark dscp** field, enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 63 to set the DSCP value for the traffic belonging to a specific class on traffic violation.
n) Choose **Shared Policer** or **Dedicated Policer** as the policy mode. The default value is **Dedicated Policer**. The shared policer mode allows you to apply the same policing parameters to several interfaces simultaneously.
o) Enter the number of packets allowed at line rate during burst, as the burst size.
p) From the drop-down list, choose the unit at which the burst size has to be calculated.
q) Enter **unspecified**, **default value**, **0xffff**, or enter a number between 0 and 5497555813760 to configure the excessive burst size.
r) From the drop-down list, choose the unit at which the excessive burst size has to be calculated.
s) Enter a number between 0 and 4398046510080 as an allowed rate. This is the committed rate at which the packets are allowed into the system (raw NTPD format).
t) From the drop-down list, choose the unit at which the allowed rate has to be calculated.
u) The **Peak Rate** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. Enter a number between 0 and 4398046510080 as the peak rate. This is the higher rate configured in a dual rate policer.
v) The **Peak Rate Unit** field appears only when **2 Rate 3 Color** is chosen as the traffic policer type. From the drop-down list, choose the unit at which the peak rate has to be calculated.

Step 9 Click **Submit**.

FHS Trust Policy

First-Hop Security (FHS) features enable a better IPv4 and IPv6 link security and management over the layer 2 links. In a service provider environment, these features closely control address assignment and derived operations, such as Duplicate Address Detection (DAD) and Address Resolution (AR).

Creating a FHS Trust Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click **Tenant(s)**.
Step 4 Click the row with the tenant that you want to update and click **View Details**.
Step 5 Click **FHS Trust Policy**.
Step 6 Click **Add**.

Step 7 On the **Create FHS Trust Policy** screen, complete the following fields:

a) Enter a unique name and description for the FHS trust policy.
b) Check the **Trust DHCP v4 Server** check box to allow the port to trust the DHCPv4 servers to get IP addresses and other information. Uncheck this box to drop traffic from DHCPv4 servers to prevent unauthorized servers from providing any configuration information.
c) Check the **Trust DHCP v6 Server** check box to allow the port to trust the DHCPv6 servers to get IP addresses and other information. Uncheck this box to drop traffic from DHCPv6 servers to prevent unauthorized servers from providing any configuration information.
d) Check the **Trust IP v6 Server** check box to allow the port to trust the IP v6 servers to get IP addresses and other information. Uncheck this box to drop traffic from IP v6 servers to prevent unauthorized servers from providing any configuration information.
e) Check the **Trust ARP** check box to allow the port to trust ARP packets.
f) Check the **Trust ND** check box to allow the port to trust neighbor discovery (ND) protocol packets.
g) Check the **Trust RA** check box to allow the port to forward all router advertisement (RA) messages without being validated against the policy.

Step 8 Click **Submit**.

Creating a BGP Address Family Context Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **BGP Address Family Context Policy**.

Step 7 Click **Create**.

Step 8 On the **Create APIC BGP Address Family Context Policy** screen, complete the following fields:

a) Enter a unique name and short description for the BGP Address Family Context policy.
b) In the **eBPG Distance** field, enter the administrative distance for external route. The external distance must be in the range of 1 to 255. The default value is 20.
c) In the **iBGP Distance** field, enter the administrative distance for internal route. The internal distance must be in the range of 1 to 225. The default value is 200.
d) In the **Local Distance** field, enter the administrative distance for the routes added with the network command. The internal distance must be in the range of 1 to 225. The default value is 220.
e) In the **eBGP Max ECMP** field, enter a value in the range of 1 to 16 as the maximum allowed equal-cost multipath (ECMP) limit for external route. The default value is 16.
f) In the **iBGP Max ECMP** field, enter a value in the range of 1 to 16 as the maximum allowed ECMP limit for internal route. The default value is 16.
g) Check the **Enable Host Route Leak** check box to enable host route leak.

Step 9 Click **Submit**.

NetFlow Monitor Policy

NetFlow policies can be deployed on a per-interface basis. Depending on the traffic-type or address family to be monitored (IPv4, IPv6, or Layer 2), you can enable different NetFlow monitor policies. A monitor policy acts as a container to hold relationships to the record policy and exporter policy. A monitor policy identifies packet flows for ingress IP packets and provides statistics based on these packet flows.

Adding a Logical NetFlow Monitoring Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **NetFlow Monitor Policy**.

Step 7 Click **Add**.

Step 8 On the **Create NetFlow Monitor Policy** screen, complete the following fields:

a) Enter the name and description for the NetFlow monitor policy.

b) Click **Select** and choose a flow record that you want to associate to the NetFlow monitor policy.

Step 9 Click **Submit**.

Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding-path failure detection times for media types, encapsulations, topologies, and routing protocols. You can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different protocol hello mechanisms. BFD makes network profiling and planning easier and reconvergence time consistent and predictable.

Use BFD to provide sub-second failure detection times in the forwarding path between ACI fabric border leaf switches configured to support peering router connections.

Creating a BFD Interface Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Cisco UCS Director

Step 7 Click **BFD Interface Policy**.

Step 8 Click **Add**.

Step 9 On the **Create APIC BFD Interface Policy** screen, complete the following fields:

a) Enter the name and description for the BFD interface policy.
b) From the **Admin State** drop-down list, choose **Enabled** to enable the admin state for the BFD interface policy.
c) Check the **Control State** check box to enable the control state for the BFD interface policy.
d) In the **Detection Multiplier** field, enter the value for the detection multiplier. The value must be in the range of 1 to 50. The default value is 3.
e) In the **Minimum Transmit Interval(msec)** field, enter the minimum transmit interval in milliseconds. The value must be in the range of 50 to 999. The default value is 50.
f) In the **Minimum Receive Interval(msec)** field, enter the minimum receive interval in milliseconds. The value must be in the range of 50 to 999. The default value is 50.
g) In the **Echo Receive Interval(msec)** field, enter the echo receive interval in milliseconds. The value must be in the range of 50 to 999. The default value is 50.
h) From the **Echo Admin State** drop-down list, choose **Enabled** to enable the echo admin state for the BFD interface policy.

Step 10 Click **Submit**.

Creating a BFD Interface Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile that you want to update and click **View Details**.

Step 10 Click **Logical Interface Profile**.

Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.

Step 12 Click **BFD Interface Profile**.

Step 13 Click **Add**.

Step 14 On the **Add Logical BFD Interface Profile** screen, complete the following fields:

a) From the **Authentication Type** drop-down list, choose **No authentication** or **Keyed SHA1**. If you choose authenticate (by selecting Keyed SHA1), enter the **Authentication Key ID**, enter the **Authentication Key** (password), then confirm the password by re-entering it next to **Confirm Authentication Key**.
b) Click **Select** and choose a BFD interface policy to which you want to add a logical BFD interface profile.

Step 15 Click **Submit**.

Hot Standby Router Protocol

Hot Standby Router Protocol (HSRP) is a first-hop redundancy protocol (FHRP) that allows a transparent failover of the first-hop IP router. HSRP provides a redundant first hop gateway for the hosts on the LAN. In the Cisco ACI fabric, HSRP can be configured on a Layer 3 physical interface or a Layer 3 sub-interface at leaf switches connected to the Layer 2 switches. The protocol acts as the gateway for the endpoints behind the Layer 2 switches.

Adding a HSRP Interface Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **HSRP Interface Policy**.

Step 7 Click **Add**.

Step 8 On the **Add HSRP Interface Policy** screen, complete the following fields:

a) Enter the name and description for the HSRP interface policy.

b) Check the **Enable Bidirectional Forwarding Detection (BFD) protocol** check box to enable BFD protocol.

c) Check the **Use Burnt-In MAC Address of the interface** check box to source hellos from the burned-in MAC address (BIA) so that HSRP routers can correctly identify each other devices.

d) Enter the minimum time (in seconds) to delay HSRP group initialization after an interface comes up. This period applies to all subsequent interface events. The default value is 0.

e) Enter the time period to delay HSRP group initialization after the router has reloaded. This period applies only to the first interface-up event after the router has reloaded. The default value is 0.

Step 9 Click **Submit**.

Creating a HSRP Group Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **HSRP Group Policy**.

Step 7 Click **Create**.

Step 8 On the **Create HSRP Group Policy** screen, complete the following fields:

a) Enter the name and description for the HSRP Group policy.

b) Enter a key or password for authentication. The default key is *cisco*.
c) Check the **Enable preemption for the group** check box to enable the preemption so that if a router fails and then comes back up, preemption restores load sharing.
d) Enter a value to set the priority in HSRP to define the active router and the standby router. This is used to exchange HSRP hello messages. The priority must be in the range of 0 to 255. The default value is 100.
e) Choose one of the following as an authentication method type. The default value is **Simple authentication**.

- **MD5 authentication**—To use the checksum of the key created, for authentication.
- **Simple authentication**—To use the clear text password, for authentication.

f) Enter the hello interval in milliseconds as the interval between hello packets that HSRP sends on the interface. If the hello interval is smaller, the topological changes will be detected at faster phase but more routing traffic will ensue. The interval must be in the range of 250 to 254000. The default value is 3000.
g) Enter the hold interval in milliseconds as the duration set for the HSRP extended hold timer for both IPv4 and IPv6 groups. The interval must be in the range of 750 to 255000. The default value is 10000.
h) Enter the minimum preemption delay that can be taken by the router to take over as the active router for an HSRP group if it has a higher priority than the current active router. The delay must be in the range of 0 to 3600. The default value is 0.
i) Enter the delay time for resuming the preemptive action after reloading the active HSRP leaf. The delay time must be in the range of 0 to 3600. The default value is 0.
j) Enter the maximum amount of time allowed for the HSRP client to prevent preemption. The valid range is from 0 to 3600. The default value is 0.
k) Enter the timeout value for authentication after which HSRP will only accept a new key for authentication. The timeout must be in the range of 0 to 32767. The default value is 0.

Step 9 Click **Submit**.

Creating a HSRP Interface Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside that you want to update and click **View Details**.
Step 8 Click **Logical Node Profile**.
Step 9 Click the row with the logical node profile that you want to update and click **View Details**.
Step 10 Click **Logical Interface Profile**.
Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.
Step 12 Click **HSRP Interface Profile**.
Step 13 Click **Add**.
Step 14 On the **Add HSRP Interface Profile** screen, complete the following fields:

a) Click **Select** and choose a HSRP interface policy that you want to use for the logical interface profile.
b) From the **Version** drop-down list, choose a version. The default is **Version 1**.

Step 15 Click **Submit**.

Note On the **Logical Interface Profile** screen, click **HSRP Interface Profile Version** to view the version details.

Creating HSRP Interface Group for HSRP Interface Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside that you want to update and click **View Details**.
Step 8 Click **Logical Node Profile**.
Step 9 Click the row with the logical node profile that you want to update and click **View Details**.
Step 10 Click **Logical Interface Profile**.
Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.
Step 12 Click **HSRP Interface Profile**.
Step 13 Click the row with the HSRP interface profile that you want to update and click **View Details**.
Step 14 Click **HSRP Interface Group**.
Step 15 Click **Add**.
Step 16 On the **Create HSRP Interface Group** screen, complete the following fields:
a) Enter the name and description of the HSRP interface group.
b) In the **Group ID** field, enter the group ID. If you have chosen HSRP version 1 in the HSRP interface profile, the maximum allowed value is 255. If you have chosen HSRP version 2 in the HSRP interface profile, the maximum allowed value is 4095.
c) In the **IP** field, enter an IP address. The IP address must be in the same subnet as the interface.
d) In the **MAC** field, enter a MAC address.
e) In the **Group Name** field, enter a group name. This is the name used in the protocol by HSRP for the HSRP MGO feature.
f) From the **Group Type** drop-down list, choose the desired type.
g) From the **IP Obtain Mode** drop-down list, choose the desired mode.
h) Click **Select** and choose a HSRP group policy.
i) In the **Secondary Virtual IPs** field, enter comma separated multiple IP addresses that can be used as secondary virtual IP address. You can provide either IPv4 or IPv6 addresses.

Step 17 Click **Submit**.

Routed Outside

Layer 3 connectivity between the devices that reside in the fabric and the rest of the enterprise network, is referred as external routed network.

Creating a Routed Outside

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click **Add**.

Step 8 On the **Create Routed Outside** screen, complete the following fields:

a) Enter a unique name, description, and alias for the external routed network.

b) Click **Select** and check the tag that you want to use for the external routed network.

c) Click **Select** and check the network that you want to use for the external routed network.

d) The **Do you want to create Routed Outside for Multipod?** check box appears only for Infra tenant. Check this check box to enable creation of routed outside for multipod. On checking this check box, the **Route Target** drop-down list appears. From the **Route Target** drop-down list, choose one of the following:

- **Automatic**—To implement automatic BGP route-target filtering on VRFs associated with this routed outside configuration.
- **Explicit**—To implement route-target filtering through use of explicitly configured BGP route-target policies on VRFs associated with this routed outside configuration.

Note You can view the route targets set for the Infra tenant under the **Route Target** tab (**APIC Account** > **Tenant(s)** > **Routed Outside**).

e) Click **Select** and check the external routed domain that you want to use for the external routed network.

f) Check the **BGP** check box to configure BGP as the routing protocol.

g) Check the **OSPF** check box to configure OSPF as the routing protocol. On checking the **OSPF** check box, enter the OSPF area ID and choose the OSPF area type.

h) Check the **EIGRP** check box to configure EIGRP as the routing protocol. This check box is disabled when the **BGP** check box or **OSPF** check box is selected.

i) Check the **PIM** check box to configure protocol independent multicast (PIM) protocol as the routing protocol.

j) Choose a target that changes the DSCP (Differentiated Services Field) marks inside a packet. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.

k) Click **Select** and check the route profile for interleak.

l) Enter the label for the provider.

m) Enter the label for the consumer.

n) By default, the export mode is disabled.

o) Check the **Import** check box to enable import mode for the external routed network.

Step 9 Click **Submit**.

Adding a Route Map or Profile to an External Routed Network

The route profile is a logical policy that defines an ordered set of logical match action rules with associated set action rules. The route profile is the logical abstract of a route map.

Any protocol enabled on Layer 3 Out, can use the export and import route map for route filtering.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside to which you want to add a route profile and click **View Details**.

Step 8 Click **Route Map or Profile**.

Step 9 Click **Add**.

Step 10 On the **Add Route Map or Profile to External Routed Network** screen, complete the following fields:

a) Choose one of the following to create route profile:

- **default-import**—To create route map for import route control.
- **default-export**—To create route map for export route control.
- **Enter Customized Value**—To create other route maps (not named default-export or default-import), choose **Enter Customized Value**. Enter route control profile name in the **Enter Route Control Profile Name** field.

b) Choose one of the following as the match type of the route profile.

- **Match Prefix AND Route Policy**—Pervasive subnets (fvSubnet) and external subnets (l3extSubnet) are combined with a route profile and merged into a single route map (or route map entry). This option is the default value.
- **Match Route Policy Only**—The route profile is the only source of information to generate a route map, and it will overwrite other policy attributes.

c) Enter a short description for the route control profile.

Step 11 Click **Submit**.

Adding a Logical Node Profile to an External Routed Network

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside to which you want to add a logical node profile and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click **Add**.

Step 10 On the **Add Logical Node Profile to External Routed Network** screen, complete the following fields:

a) Enter a unique name and description for the logical node profile.

b) Choose a DSCP target to which the DSCP range has to be mapped. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.

Step 11 Click **Submit**.

Adding a Logical Node to a Logical Node Profile of an External Routed Network

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the external routed network to which you want to add a logical node and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile to which you want to add a logical node and click **View Details**.

Step 10 Click **Logical Nodes**.

Step 11 Click **Add**.

Step 12 On the **Add Logical Node to Logical Node Profile of External Routed Network** screen, complete the following fields:

a) By default, the logical node profile name is displayed. If you want to add a logical node to a different logical node profile, click **Select** and choose a logical node profile.

b) Click **Select** and choose a logical node that needs to be added to the logical node profile.

c) Enter the IP address of the local routing device.

d) Choose **true** from the **Use Router ID as Loop Back Address** drop-down list, to use the router ID as the loopback address. The default value is **Unspecified**.

e) The **External Control Peering** checkbox is enabled only for infra tenant and when you have an external routed network created with OSPF input in the infra tenant. Check the **External Control Peering** checkbox to enable external control peering only when the remote leaf switches are being added to a pod in a multipod fabric.

Step 13 Click **Submit**.

Adding a Static Route to a Logical Node

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside to which you want to add a static route and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile to which you want to add a static route and click **View Details**.

Step 10 Click **Logical Nodes**.

Step 11 Click the row with the logical node to which you want to add a static route and click **View Details**.

Step 12 Click **Static Routes**.

Step 13 Click **Add**.

Step 14 On the **Add Static Route to Logical Node** screen, complete the following fields:

a) Enter an IP address for the static route with mask. For example, 10.10.10.0/24.

b) Enter a value between 1 and 255 as the static route preference to control the routes based on next hop. The default value is 1.

c) Check the **Route Control BFD** check box to enable BFD on static route to provide fast forwarding path failure detection.

d) Enter the IP address of the next hop.

Step 15 Click **Submit**.

Adding an External Network to an External Routed Network

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside to which you want to add an external network and click **View Details**.

Step 8 Click **External Network**.

Step 9 Click **Add**.

Step 10 On the **Add External Network to External Routed Network** screen, complete the following fields:

a) Enter a unique name, description, and alias for the external network.
b) Click **Select** and choose the tag that you want to use for the external network.
c) Choose **Level1**, **Level2**, **Level3**, **Level4**, **Level5**, **Level6**, **Custom**, or **Unspecified** as the QoS class for the external network. By default, **Unspecified** is set as the QoS class.
d) Enter the contract exception tag.
e) Choose a DSCP target to which the DSCP range has to be mapped. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.
f) Choose **Include** to mark the external network as the preferred group member.

Step 11 Click **Submit**.

Adding a Next Hop Address to a Static Route

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile that you want to update and click **View Details**.

Step 10 Click **Logical Node**.

Step 11 Click the row with the logical node that you want to update and click **View Details**.

Step 12 Click **Static Routes**.

Step 13 Click the row with the static route that you want to update and click **View Details**.

Step 14 Click **Next HOP Addresses**.

Step 15 Click **Add**.

Step 16 On the **Add Next Hop Address to Static Route** screen, complete the following fields:

a) Enter a valid IP address for the next hop.
b) Enter a value in the range of 1 to 255 to set the preference for the static route.

Step 17 Click **Submit**.

Adding a Route Control Profile to an External Network

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **External Network**.

Step 9 Click the row with the external network that you want to update and click **View Details**.

Step 10 Click **Route Control profile**.

Step 11 Click **Add**.

Step 12 On the **Add Route Control Profile to External Network** screen, complete the following fields:

a) Check the **Customize Route Control Profile Name?** check box to enter a unique name for the route control profile. If you want to use the existing profile, click **Select** and choose a route control profile name that you want to use.

b) Choose **Route Import Policy** or **Route Export Policy** as the direction for the route control profile.

Step 13 Click **Submit**.

Adding a Route Control Profile to a Subnet

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **External Network**.

Step 9 Click the row with the external network that you want to update and click **View Details**.

Step 10 Click **Subnet**.

Step 11 Click the row with the subnet that you want to update and click **View Details**.

Step 12 Click **Route Control Profile**.

Step 13 Click **Add**.

Step 14 On the **Add Route Control Profile to Subnet** screen, complete the following fields:

a) Check the **Customize Route Control Profile Name?** check box to enter a unique name for the route control profile. If you want to use the existing profile, click **Select** and choose a route control profile name that you want to use.

b) Choose **Route Import Policy** or **Route Export Policy** as the direction for the route control profile.

Step 15 Click **Submit**.

Adding a Logical NetFlow Monitoring Policy

NetFlow policies can be deployed on a per-interface basis. Depending on the traffic-type or address family to be monitored (IPv4, IPv6, or Layer 2 (CE type)), you can enable different NetFlow monitor policies. A monitor policy acts as a container to hold relationships to the record policy and exporter policy. A monitor policy identifies packet flows for ingress IP packets and provides statistics based on these packet flows.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile that you want to update and click **View Details**.

Step 10 Click **Logical Interface Profile**.

Step 11 Click the row with the logical interface profile to which you want to add logical NetFlow monitor policy and click **View Details**.

Step 12 Click **Logical NetFlow Monitor Policy**.

Step 13 Click **Add**.

Step 14 On the **Add Logical NetFlow Monitoring Policy** screen, complete the following fields:

a) Click **Select** and choose a logical interface profile to which you want to add logical NetFlow monitor policy.

b) Choose **CE Type**, **IPv4**, or **IPv6** as the NetFlow IP filter type.

c) Click **Select** and choose a NetFlow monitor policy that needs to be associated with the logical interface profile.

Step 15 Click **Submit**.

Adding a Secondary IP Address to an Interface

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile that you want to update and click **View Details**.

Step 10 Click **Logical Interface Profile**.

Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.

Step 12 Click **Routed Interface**.

Step 13 Click the row with the interface that you want to configure according to the following port type:

- **l3-port**—To configure the routed interface.
- **sub-interface**—To configure the routed sub-interface interface.
- **ext-svi**—To configure the SVI interface.

Step 14 Click **Secondary Addresses**.

Step 15 Click **Add**.

Step 16 On the **Add Secondary IP Address to Interface** screen, complete the following fields that vary according to the chosen interface configuration and primary IP address:

a) Enter the IPv4 or IPv6 address with subnet mask that needs to be configured as the secondary IP address of the routed interface.

b) On entering IPv6 address, the following additional fields appear:

1. **IPv6 DAD** drop-down list—Choose **Enabled** to enable duplicate address detection to verify if IPv6 is unique.
2. **ND RA Prefix** check box—Check this box to append neighbor discovery router advertisement (ND RA) prefix for the interface.
3. **ND RA Prefix Policy** field—The **ND RA Prefix Policy** field appears on checking the **ND RA Prefix** check box. Click **Select** and choose the ND RA Prefix policy that you want to use for the interface.

c) If you are configuring this setting for the SVI interface, you have to set the secondary IP address configuration for side A and side B.

Step 17 Click **Submit**.

Adding a BGP Peer Connectivity Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Routed Outside**.

Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.

Step 9 Click the row with the logical node profile that you want to update and click **View Details**.

Step 10 Click **Logical Interface Profile**.

Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.

Step 12 Click **Routed Interface**.

Step 13 Click the row with the interface that you want to configure according to the following port type:

- **l3-port**—To configure the routed interface.

- **sub-interface**—To configure the routed sub-interface interface.
- **ext-svi**—To configure the SVI interface.

Step 14 Click **BGP Connectivity Profile**.

Step 15 Click **Add**.

Step 16 On the **Add BGP Peer Connectivity Profile** screen, complete the following fields:

a) In the **Peer Address** field, enter the peer IP address in the IPv4, IPv6, IPv4/prefix, or IPv6/prefix format.
b) Enter a short description for the BGP peer connectivity profile.
c) Check the **Allow Self AS** check box to allow the readvertisement of all prefixes that contain duplicate autonomous system numbers (ASNs).
d) Check the **AS Override** check box to replace the autonomous system (AS) number of originating router with the AS number of the sending BGP router.
e) Check the **Disable Peer AS Check** check box to disable checking of the peer AS number when advertising.
f) Check the **Next-hop Self** check box to always set the next hop attribute to the local peering address.
g) Check the **Send Community** check box to send the community attribute to the neighbor.
h) Check the **Send Extended Community** check box to send the extended community attribute to the neighbor.
i) In the **Password** field, enter the password for BGP MD5 authentication.
j) In the **Confirm Password** field, re-enter the password for BGP MD5 authentication.
k) The **Allow Self AS Count** field displays the allowed AS number count.
l) Check the **Bidirectional Forwarding Detection** check box to enable BFD in the BGP peer connectivity profile.
m) Check the **Disable Connected Check** check box to skip the connected check and attempt the connection to the peering address.
n) In the **EBGP Multihop TTL** field, enter a time-to-live (TTL) value in the BGP packets. The default value is 1. You can enter either **defaultValue** or any value in the range of 1 to 255.
o) In the **Weight for routes from this neighbour** field, enter a value in the range of 0 to 65535 as the weight value for routes from the neighbour.
p) The **Remove all private AS** check box is active only when the **Remove Private AS** check box is checked. Check **Remove all private AS** check box to remove all private AS numbers from outbound route updates to an eBGP peer.
q) Check **Remove private AS** check box to remove private AS numbers from outbound route updates to an eBGP peer.
r) The **Replace private AS with local AS** check box is active only when the **Remove Private AS** check box is checked. Check **Replace private AS with local AS** check box to replace all private AS numbers with local AS numbers.
s) In the **BGP Peer Prefix Policy** field, click **Select** and chosee a BGP peer prefix policy to be associated with the BGP peer connectivity profile.
t) In the **Remote Autonomous System Number** field, enter a value in the range of 1 to 429467925 as the remote ASN.
u) From the **Local-AS Number Config** drop-down list, choose any one of the option as the method for configuring local ASN.
v) In the **Local-AS Number** field, enter a value in the range of 1 to 429467925 as the local ASN.

Step 17 Click **Submit**.

Adding a Loopback Address to a Logical Node

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside to which you want to add a static route and click **View Details**.
Step 8 Click **Logical Node Profile**.
Step 9 Click the row with the logical node profile to which you want to add a static route and click **View Details**.
Step 10 Click **Logical Nodes**.
Step 11 Click the row with the logical node to which you want to add a static route and click **View Details**.
Step 12 Click **Loopback Address**.
Step 13 Click **Add**.
Step 14 On the **Add Loopback Address to Logical Node** screen, enter the loopback IP address of the logical node.
Step 15 Click **Submit**.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is a network management protocol used to dynamically assign an Internet Protocol (IP) address to any device, or node, on a network so they can communicate using IP. DHCP automates and centrally manages these configurations rather than requiring network administrators to manually assign IP addresses to all network devices. DHCP can be implemented on small local networks as well as large enterprise networks.

A DHCP relay policy may be used when the DHCP client and server are in different subnets. The DHCP relay profile contains one or more providers. The consumer bridge domain contains the DHCP label that associates the provider DHCP server with the bridge domain. Label matching enables the bridge domain to consume the DHCP Relay policy.

Creating a DHCP Relay Label

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside that you want to update and click **View Details**.

Step 8 Click **Logical Node Profile**.
Step 9 Click the row with the logical node profile that you want to update and click **View Details**.
Step 10 Click **Logical Interface Profile**.
Step 11 Click the row with the logical interface profile that you want to update and click **View Details**.
Step 12 Click **DHCP Relay Label**.
Step 13 Click **Add**.
Step 14 On the **Add DHCP Relay Label To Logical Interface Profile** screen, complete the following fields:
a) From the **Scope** drop-down list, choose **infra** or **tenant**.
b) Click **Select** and choose a DHCP relay policy that you want to use for the logical interface profile.
Step 15 Click **Submit**.

Creating a DHCP Relay Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **DHCP Relay Policy**.
Step 7 Click **Add**.
Step 8 On the **Create DHCP Relay Policy** screen, enter a unique name and description for the relay policy.
Step 9 Click **Submit**.

Adding a Provider to the DHCP Relay Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **DHCP Relay Policy**.
Step 7 Click the row with the DHCP relay policy to which you want to add a provider and click **View Details**.
Step 8 Click **Providers**.
Step 9 Click **Add**.
Step 10 On the **Add providers to DHCP Relay Policy** screen, complete the following fields:
a) From the **EPG Type** drop-down list, click the appropriate option depending upon where the DHCP server is connected.

b) This field varies according to the option selected in the **EPG Type** drop-down list. Click **select** and check the appropriate EPG.
c) In the **DHCP Server Address** field, enter the IP address of the DHCP server.

Step 11 Click **Submit**.

To add provider at the infrastructure level, navigate to **Physical** > **Network** > **APIC Acccount** > **Relay Policy** > **Provider** and then click **Add**. For more details on the fields, refer to the Step 10.

Creating a DHCP Option Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **DHCP Option Policy**.
Step 7 Click **Add**.
Step 8 On the **Create DHCP Option Policy** screen, enter a unique name and description for the DHCP option policy.
Step 9 Click **Submit**.

Adding a DHCP Option to a DHCP Option Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **DHCP Option Policy**.
Step 7 Click the row with the DHCP option policy to which you want to add a DHCP option and click **View Details**.
Step 8 On the **Add DHCP Option** screen, click **Add** and enter a unique name, ID, and description for the DHCP option.
Step 9 Click **Submit**.

IGMP Interface Policy

Adding an IGMP Interface Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **IGMP Interface Policy**.

Step 7 Click **Add**.

Step 8 On the **Add IGMP Interface Policy** screen, complete the fields including the following:

a) Enter the name and description for the IGMP interface policy.
b) Check the **Allow v3 ASM** check box to accept IGMPv3 reports for non SSM groups.
c) Check the **Fast Leave** check box to configure fast leave for the VLAN/BD.
d) Enter the group timeout in seconds to configure group membership timeout in all VLAN/BDs. The default value is 260 seconds.
e) Enter the query interval in seconds to configure interval between group-specific query transmissions. The default value is 125 seconds.
f) Enter the query response interval in seconds. The default value is 10 seconds.
g) Enter the last member count. The default value is 2.
h) Enter the response time of last member in seconds. The default value is 1 second.
i) Enter the start-up query count to configure the number of queries to be sent at start-up. The default value is 2.
j) Enter the start-up query interval in seconds to configure the query interval at start-up. The default value is 31 seconds.
k) Enter the querier timeout for IGMP in seconds. The default value is 255 seconds.
l) Enter the robustness variable that you can tune to reflect expected packet loss on a congested network. You can increase the robustness variable to increase the number of times that packets are resent. Values range from 1 to 7. The default value is 2.
m) Choose the IGMP version that is enabled on the interface. The IGMP version can be **Version 2** or **Version 3**. The default value is **Version 2**.
n) Click **Select** and choose a report policy route map for the IGMP.
o) Click **Select** and choose a static report route map for the IGMP.
p) Enter the maximum number of allowed multicast entries.
q) Enter the reserved multicast entries.
r) Click **Select** and choose a state limit route map for the IGMP.

Step 9 Click **Submit**.

Route Tag Policy

A route tag is a 32-bit value attached to routes. Route tags are used to filter routes and apply administrative policies,such as redistribution and route summarization, to tagged routes.

When a transit route is redistributed into OSPF or EIGRP, the route is tagged with the tag value specified in the route tag policy to prevent routing loops. If a route is received on an OSPF or EIGRP L3Out with this tag value, the route is dropped. The default route tag policy tag value is 4294967295.

Creating a Route Tag Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Route Tag Policy**.

Step 7 Click **Add**.

Step 8 On the **Create Route Tag Policy** screen, complete the following fields:

a) Enter a unique name and description for the route tag policy.

b) Enter a value in the range of 0 to 4294967295 as the route tag policy tag value. The default route tag policy tag value is 4294967295.

Step 9 Click **Submit**.

EIGRP Address Family Context Policy

EIGRP Address Family Context Policy (eigrpCtxAfPol) contains the configuration for a specific address family in a given VRF. An eigrpCtxAfPol policy is configured within multiple tenant protocol policies and can be applied to one or more VRFs under the tenant.

You can enable this policy on a VRF with a relation in the VRF-per-address family. If there is no relation to a given address family or the specified eigrpCtxAfPol in the relation does not exist, the default VRF policy created within the common tenant is used for that address family.

Creating an EIGRP Address Family Context Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **EIGRP Address Family Context Policy**.

Step 7 Click **Create**.

Step 8 On the **Create APIC EIGRP Address Family Context Policy** screen, complete the following fields:

a) Enter a unique name and short description for the EIGRP Address Family Context policy.

b) In the **Active Interval (min)** field, enter a value in the range of 1 to 65535 as the active timer interval. The default value is 3.

c) In the **External Distance** field, enter the administrative distance for external route. The external distance must be in the range of 1 to 225. The default value is 170.

d) In the **Internal Distance** field, enter the administrative distance for internal route. The internal distance must be in the range of 1 to 225. The default value is 90.

e) In the **Maximum Path Limit** field, enter a value in the range of 1 to 16 as the maximum allowed equal-cost multipath (ECMP) limit. The default value is 8.

f) From the **Metric Style** drop-down list, choose **narrow metric** to use 32-bit metrics version or **wide metric** to use 64-bit metrics version.

Step 9 Click **Submit**.

OSPF Timers

Open Shortest Path First (OSPF) routing devices constantly track the status of their neighbors, sending and receiving hello packets that indicate whether each neighbor still is functioning, and sending and receiving link-state advertisement (LSA) and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You configure OSPF timers on the interface of the routing device participating in OSPF.

Creating a OSPF Timer Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **OSPF Timers**.

Step 7 Click **Create**.

Step 8 On the **Create OSPF Timers** screen, complete the following fields:

a) Enter a unique name and short description for the OSPF timer policy.

b) In the **Bandwidth Reference (MBPS)** field, enter a value in the range of 0 to 4000000. The default value is 40,000.

c) In the **Admin Distance Preference** field, enter the administrative distance. The distance must be in the range of 1 to 255. The default value is 110.

d) In the **Maximum ECMP** field, enter a value the range of 1 to 16 as the maximum allowed equal-cost multipath (ECMP) limit. The default value is 8.

e) Check the **Enable Name Lookup for Router IDs** check box to enable name lookup for the routers.
f) Check the **Prefix Suppression** check box to hide IPv4 and IPv6 prefixes that are configured on interfaces running OSPF.
g) In the **Initial SPF Schedule Delay Interval (MS)** field, enter a value in the range of 1 to 600000 as the initial SPF schedule in milliseconds. The default value is 200.
h) In the **Minimum Hold Time Between SPF Calculations (MS)** field, enter a value in the range of 1 to 600000 as the minimum hold time in milliseconds. The default value is 1000.
i) In the **Maximum Hold Time Between SPF Calculations (MS)** field, enter a value in the range of 1 to 600000 as the maximum wait time in milliseconds. The default value is 5000.
j) In the **LSA Group Pacing Interval (SECS)** field, enter a value in the range of 1 to 1800 in seconds. The default value is 5000.
k) In the **Minimum Interval Between Arrival of a LSA (MS)** field, enter a value in the range of 1 to 600000 in milliseconds. The default value is 1000.
l) In the **LSA Generation Throttle Start Wait Interval (MS)** field, enter a value in the range of 1 to 5000 in milliseconds. The default value is 0.
m) In the **LSA Generation Throttle Hold Interval (MS)** field, enter a value in the range of 50 to 30000 in milliseconds. The default value is 5000.
n) In the **LSA Generation Throttle Maximum Interval (MS)** field, enter a value in the range of 50 to 30000 in milliseconds. The default value is 5000.
o) Check the **Graceful Restart Helper** check box to enable graceful restart helper for timer.
p) In the **Maximum Number of Not Self-Generated LSAs** field, enter a value in the range of 1 to 4294967295. The default value is 20000.
q) In the **LSA Threshold (percentage)** field, enter a value in the range of 1 to 100. The default value is 75.
r) From the **LAS Maximum Action** drop-down list, choose **Log** or **Reject** as the maximum action for the timer.

Step 9 Click **Submit**.

IGMP Snoop Policy

IGMP snooping is the process of listening to Internet Group Management Protocol (IGMP) network traffic. The feature allows a network switch to listen in on the IGMP conversation between hosts and routers and filter multicasts links that do not need them. So with this policy controlling which ports receive specific multicast traffic.

Adding an IGMP Snoop Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **IGMP Snoop**.
Step 7 Click **Add**.

Step 8 On the **Create APIC IGMP Snoop Policy** screen, complete the fields including the following:

a) Enter the name and description for the IGMP snoop policy.
b) From the **Admin State** drop-down, choose state for the IGMP snoop policy. The default is **enabled**.
c) Check the **Fast Leave** check box to enable IGMP V2 immediate dropping of queries through this policy..
d) Check the **Enable querier** check box to enable the IGMP querier activity through this policy.
e) Enter the query interval in seconds to define the amount of time the IGMP function will store a particular IGMP state if it does not hear any reports on the group. The default value is 125 seconds.
f) Enter the query response interval in seconds. When a host receives the query packet, it starts counting to a random value, less that the maximum response time. When this timer expires, host replies with a report. The default value is 10 seconds.
g) Enter the last member query interval count. IGMP uses this value when it receives an IGMPv2 Leave report. This means that at least one host wants to leave the group. After it receives the Leave report, it checks that the interface is not configured for IGMP Fast Leave and if not, it sends out an out-of-sequence query. The default value is 1.
h) Enter the start-up query count to configure the number of queries to be sent at start-up. The default value is 2.
i) Enter the start-up query interval in seconds to configure the query interval at start-up. The default value is 31 seconds.

Step 9 Click **Submit**.

MLD Snoop Policy

Multicast Listener Discovery (MLD) snooping enables the efficient distribution of IPv6 multicast traffic between hosts and routers. It is a Layer 2 feature that restricts IPv6 multicast traffic within a bridge domain to a subset of ports that have transmitted or received MLD queries or reports. In this way, MLD snooping provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving multicast traffic. This reduces bandwidth usage instead of flooding the bridge domain, and also helps hosts and routers save unwanted packet processing.

Adding a MLD Snoop Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **MLD Snoop**.
Step 7 Click **Add**.
Step 8 On the **Add APIC MLD Snoop Policy** screen, complete the fields including the following:

a) Enter the name and description for the MLD Snoop policy.
b) From the **Admin State** drop-down list, choose **Enabled** or **Disabled** to enable or disable this entire policy. The default option is **Enabled**.
c) Check **Fast leave** to enable MLD v1 immediate dropping of queries through this policy.
d) Check **Enable querier** to enable the MLD querier activity through this policy.

Note For this option to be effectively enabled, the Subnet Control: Querier IP setting must also be enabled in the subnets assigned to the bridge domains to which this policy is applied.

e) In the **Query Interval (sec)** field, enter the query interval value between general queries sent by the querier. The default value is 125.

f) In the **Query Response Interval (sec)** field, enter the query response interval value. When a host receives the query packet, it starts counting to a random value, less than the maximum response time. When this timer expires, the host replies with a report. The default value is 10.

This is used to control the maximum response time for hosts to answer an MLD query message. Configuring a value less than 10 seconds enables the router to prune groups much faster, but this action results in network burstiness because hosts are restricted to a shorter response time period.

g) In the **Last Member Query Interval (sec)** field, enter the last member query interval value. The default value is 1.

MLD uses this value when it receives an MLD Leave report. This means that at least one host wants to leave the group. After it receives the Leave report, it checks that the interface is not configured for MLD Fast Leave and, if not, it sends out an out-of-sequence query.

If no reports are received in the interval, the group state is deleted. The software can detect the loss of the last member of a group or source more quickly when the values are smaller. Values range from 1 to 25 seconds.

h) In the **Start Query Count** field, enter the number of queries sent at startup that are separated by the startup query interval. The default value is 2.

i) In the **Start Query Interval (sec)** field, enter the value to configures the MLD snooping query interval at startup.

By default, this interval is shorter than the query interval so that the software can establish the group state as quickly as possible. Values range from 1 to 18,000 seconds. The default is 31 seconds.

Step 9 Click **Submit**.

Monitoring Policy

As an administrator, you can create monitoring policy at tenant level to monitor EPGs, application profiles, services, and so on. The default policy can be overridden by a specific policy. For example, a monitoring policy applied to the Solar tenant (*tn-solar*) would override the default one for the Solar tenant while other tenants would still be monitored by the default policy.

Adding a Monitoring Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Monitoring Policy**.

Step 7 Click **Add**.

Step 8 On the **Create Monitoring Policy** screen, enter a unique name and description for the monitoring policy.

Step 9 Click **Submit**.

NetFlow Monitor Policy

NetFlow policies can be deployed on a per-interface basis. Depending on the traffic-type or address family to be monitored (IPv4, IPv6, or Layer 2), you can enable different NetFlow monitor policies. A monitor policy acts as a container to hold relationships to the record policy and exporter policy. A monitor policy identifies packet flows for ingress IP packets and provides statistics based on these packet flows.

Adding a Logical NetFlow Monitoring Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **NetFlow Monitor Policy**.

Step 7 Click **Add**.

Step 8 On the **Create NetFlow Monitor Policy** screen, complete the following fields:

a) Enter the name and description for the NetFlow monitor policy.

b) Click **Select** and choose a flow record that you want to associate to the NetFlow monitor policy.

Step 9 Click **Submit**.

Associating a NetFlow Monitor Policy to a Bridge Domain

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Bridge Domains**.

Step 7 Click the row with the bridge domain that you want to update and click **View Details**.

Step 8 Click **NetFlow Monitor Policies**.

Step 9 Click **Add**.

Step 10 On the **Add NetFlow Monitor Policy** screen, complete the following fields:

a) From the **NetFlow IP Filter Type** drop-down list, choose **CE Type**, **IPv4 Type**, or **IPv6 Type** as the NetFlow IP filter type. The default value is **IPv4 Type**.
b) Click **Select** and choose a NetFlow monitor policy that you want to associate with the bridge domain.

Step 11 Click **Submit**.

Flow Record

A NetFlow record lets you define a flow and the statistics to collect for each flow. You can define match parameters to identify packets in the flow, and define collect parameters that the NetFlow gathers for the flow.

Create a Tenant Flow Record

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Flow Records**.
Step 7 Click **Create**.
Step 8 On the **Create Tenant Flow Record** screen, complete the following fields:
a) Enter a unique name and short description for the flow record.
b) Click **Select** and choose a list of parameters that need to be collected for a given flow. The default value is **Source Interface**.
c) Click **Select** and choose a list of parameters that are used by NetFlow to identify packets in the flow.
Step 9 Click **Submit**.

Route Maps

You can use route maps for route redistribution or policy-based routing. Route map entries consist of a list of match and set criteria. The match criteria specify match conditions for incoming routes or packets, and the set criteria specify the action taken if the match criteria are met.

You can configure multiple entries in the same route map. These entries contain the same route map name and are differentiated by a sequence number.

The route map is a combination of match action rules and set action rules. After the match and set profiles are defined, the route map must be created in the Layer 3 Out.

Creating a Match Rule for a Route Map

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Match Rules for Route Maps**.
Step 7 Click **Add**.
Step 8 On the **Create Match Rule for a Route Map** screen, enter a unique name and description for the route map match rule.
Step 9 Click **Submit**.

Adding a Match Regex Community Term to a Route Map Match Rule

Before you begin

The route map match rule must be created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Match Rules for Route Maps**.
Step 7 Click the row with the route map match rule to which you want to add match regex community term and click **View Details**.
Step 8 Click **Match Regex Community Terms**.
Step 9 Click **Add**.
Step 10 On the **Add Match Regex Community Term** screen, complete the following fields:
a) Enter a unique name for the match regex community.
b) Enter the match community regular expression term and match community terms as desired.
c) Choose **Regular** or **Extended** as the community type.
d) Enter a description for the match regex community term.
Step 11 Click **Submit**.

Adding a Match Prefix to a Match Rule

Before you begin

The route map match rule must be created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Match Rules for Route Maps**.
Step 7 Click the row with the route map match rule to which you want to add match prefix and click **View Details**.
Step 8 Click **Match Prefix**.
Step 9 Click **Add**.
Step 10 On the **Add Prefix to Match Rule** screen, complete the following fields:
a) Enter the IPv4 or IPv6 address as the specific prefix for match rule.
b) Click to enable aggregate and allow prefix matches with multiple masks starting with the mask mentioned in the configuration till the maximum mask allowed for the address family of the prefix.
c) Enter a description for the match prefix.
Step 11 Click **Submit**.

Adding a Match Community Term to a Route Map Match Rule

Before you begin

The route map match rule must be created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Match Rules for Route Maps**.
Step 7 Click the row with the route map match rule to which you want to add match community term and click **View Details**.
Step 8 Click **Match Community Terms**.
Step 9 Click **Add**.
Step 10 On the **Add Match Community Term to Match Rule for Route Map** screen, enter a unique name and description for the match community term.
Step 11 Click **Submit**.

Adding a Match Community Factor to a Match Community Term

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Match Rules for Route Maps**.
Step 7 Click the row with the route map match rule that you want to update and click **View Details**.
Step 8 Click **Match Community Terms**.
Step 9 Click the row with the match community term to which you want to add a match community factor and click **View Details**.
Step 10 Click **Add**.
Step 11 On the **Add Match Community Factor to Match Community Term** screen, complete the following fields:
a) Enter a unique name for the match community factor.
b) Choose **Transitive** or **Non transitive** as the scope of the community.
c) Enter a description for the match community factor.
Step 12 Click **Submit**.

Creating a Route Map Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **RouteMap Policy**.
Step 7 Click **Add**.
Step 8 On the **Create RouteMap Policy** screen, enter a unique name and description for the route map policy.
Step 9 Click **Submit**.

What to do next

You can add entries to the route map.

Adding a Route Map Entry

Before you begin

The route map policy must be created.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **RouteMap Policy**.

Step 7 Click the row with the route map policy to which you want to add a route map entry and click **View Details**.

Step 8 Click **RouteMap Entry**.

Step 9 Click **Add**.

Step 10 On the **Add Route Map Entry** screen, complete the following fields:

a) Enter a value in the range of 0 to 65535 as the preferred order for the route map entry.
b) Enter a valid group IP address with prefix to match an IP multicast packet based on group. For example: 10.2.3.1/22.
c) Enter a valid source IP address with prefix to match an IP multicast packet based on source. For example: 10.1.2.3/22.
d) Enter a valid IP address of the rendezvous point (RP) to match an IP multicast packet based on RP. The RP is a router that you select in a multicast network domain that acts as a shared root for a multicast shared tree. For example: 10.3.2.3/22.
e) Choose **Permit** or **Deny** as the route map entry action.

Step 11 Click **Submit**.

Creating a Set Rules for Route Map

The action rule profile is used to define the route-map set clauses for the L3Out. The supported set clauses are the BGP communities (standard and extended), route tags, dampening, weight, next hop, preference, metric, and metric type.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Set Rules for Route Map**.

Step 7 Click **Create**.

Step 8 On the **Create Set Rules For A Route Map** screen, complete the following fields:

a) Enter a unique name and description for the action rule profile.
b) Click to set the community for the action rule profile. From the **Criteria** drop-down list, choose **Append Community**, **No Community**, or **Replace Community**. On choosing **Append Community** or **Replace Community**, enter a value for the community.
c) Click to set the route tag for the action rule profile. Enter the tag value in the range of 0 to 4294967295.
d) Click to set and configure route flap dampening behavior. The parameters are:

- Half Life—Decay half life, which is the time in minutes after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half life period. The valid range is 1 to 60 minutes.

- Reuse Limit—A route is unsuppressed (reused) if the penalty for a flapping route decreases enough to fall below this value. The valid range is 1 to 20000.
- Suppress Limit—A route is suppressed when its penalty exceeds this limit. The valid range is 1 to 20000.
- Max Suppress Time—The maximum time in minutes that a stable route can be suppressed. The valid range is 1 to 255.

e) Click to set the BGP weight for the routing table. The valid range is 0 to 65535.
f) Click to set the next hop IP address.
g) Click to set the BGP local preference value. The valid range is from 0 to 4294967295.
h) Click to set the metric for the destination routing protocol. The valid range is from 0 to 4294967295.
i) Click to set the **OSPF type1 metric** or **OSPF type2 metric** as the metric type.

Step 9 Click **Submit**.

Adding an Additional Community

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Action Rule Profile**.
Step 7 Click the row with the action rule profile that you want to update and click **View Details**
Step 8 Click **Additional Communities**.
Step 9 Click Add.
Step 10 On the **Add Additional Community** screen, complete the following fields:
a) Enter a unique name for the community in the following format: regular:as2-nn2:4:15, extended:as4-nn2:5:16, no-export, and no-advertise.
b) Enter a short description for the community.
Step 11 Click **Submit**.

Adding a Set AS Path to the Action Rule Profile

Before you begin

The action rule profile must be created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Action Rule Profile**.

Step 7 Click the row with the action rule profile to which you want to set autonomous system (AS) path and click **View Details**.

Step 8 Click **Set As Path**.

Step 9 Click **Add**.

Step 10 On the **Add Set As Path** screen, complete the following fields:

a) Choose one of the following as AS path criteria to append the specified AS number to the AS path of the route matched by the route map:

- Prepend AS—To prepend the specified AS number.
- Prepend Last-As—To prepend the last AS number to the AS path.

b) Enter the AS number in the range of 1 to 4294967295.

c) Enter the order in which the AS number is inserted into the AS path attribute. The valid range is 0 to 31.

Step 11 Click **Submit**.

Adding an AS Number to Prepend the AS Path

You can append the specified AS number to the autonomous system path of the route that is matched by the route map. This applies to both inbound and outbound BGP route maps.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Set Rules for Route Map**.

Step 7 Click the row with the route map rules that you want to update and click **View Details**.

Step 8 Click **AS Number to Prepend AS**.

Step 9 Click **Add**.

Step 10 On the **Add AS Number** screen, complete the following fields:

a) Enter the autonomous system number that need to be prepended to the matching AS path. This must be an integer ranging from 1 to 4294967295.

b) Enter the order in which the AS number need to be prepended.

Step 11 Click **Submit**.

Adding a Context to a Route Map or Profile

Before you begin

The action rule profile must be created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside that you want to update and click **View Details**.
Step 8 Click **Route Map or Profile**.
Step 9 Click the row with the route map or profile that you want to update and click **View Details**.
Step 10 Click **Context**.
Step 11 Click **Add**.
Step 12 On the **Add Context to Route Maps or Profiles** screen, complete the following fields:
a) Enter a unique name for the context.
b) Enter the order in which the context has to be inserted in to the route map or profile.
c) Choose **Permit** to permit routes that match criteria defined in match rule.
d) Enter a short description for the context.
e) Click **Select** to choose an action rule profile for the context.
Step 13 Click **Submit**.

Associating a Match Rule to a Route Control Context

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Routed Outside**.
Step 7 Click the row with the routed outside that you want to update and click **View Details**.
Step 8 Click **Route Map or Profile**.
Step 9 Click the row with the route map or profile that you want to update and click **View Details**.
Step 10 Click **Context**.
Step 11 Click the row with the context that you want to update and click **View Details**.

Step 12 Click **Match Rule**.

Step 13 Click **Add**.

Step 14 On the **Add Match Rule** screen, click **Select** and choose a match rule that you want to add to the route control context in the routed outside.

Step 15 Click **Submit**.

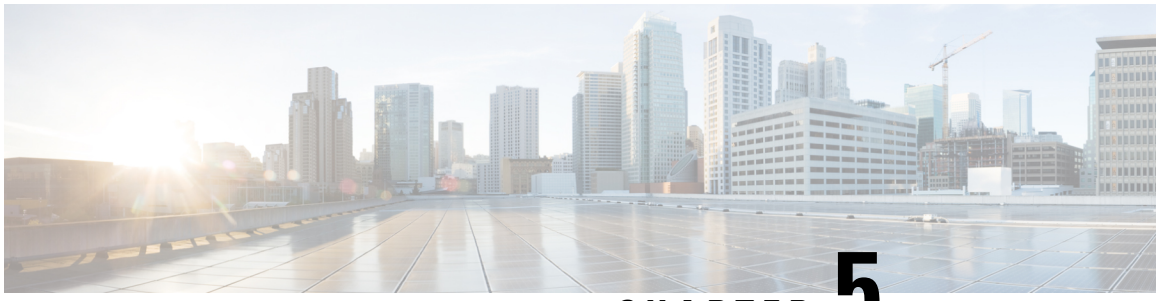

CHAPTER 5

Configuring Multi-Site Controller Accounts

Adding an ACI Multi-Site Controller Account

As an administrator, you can add an ACI Multi-Site controller account.

Step 1 Choose **Administration** > **Physical Accounts**.

Step 2 On the **Physical Accounts** page, click **Multi-Domain Managers**.

Step 3 Click **Add**.

Step 4 On the **Add Account** screen, choose **ACI Multi-Site** from the **Account Type** drop-down list and click **Submit**.

Step 5 On the **Add Account** screen, complete the fields, including the following:

a) Enter a unique account name and description for the ACI Multi-Site controller account.

b) From the **Pod** drop-down list, choose the pod where you want to add the ACI Multi-Site controller account.

c) In the **Server IP** field, enter the IP address of the ACI Multi-Site controller account.

d) Check the **Use Credential Policy** box if you want to use a credential policy for this account rather than enter the user name and password information manually.

e) If you checked the **Use Credential Policy** box, choose a policy from the **Credential Policy** drop-down list.

You can also add a new credential policy by clicking the **Add** option.

f) If you did not check **Use Credential Policy**, enter the user name and password that this account uses to access ACI Multi-Site controller.

g) If you did not check **Use Credential Policy**, do the following:

- From the **Protocol** drop-down list, choose **https** or **http**.

- In the **Port** field, enter the port used to access the ACI Multi-Site controller account. The default port is 443.

h) Enter the contact details and location of the administrator or other person responsible for this account.

Step 6 Click **Submit**.

What to do next

Cisco UCS Director tests the connection to the ACI Multi-Site controller. If that test is successful, it adds the ACI Multi-Site controller account and discovers all the elements in the ACI Multi-Site controller. This discovery process and inventory collection takes a few minutes to complete.

When you add the ACI Multi-Site controller account, system tasks related to the ACI Multi-Site account are populated on the **System Task** page (**Administration > System > System Task**).

Assigning an ACI Multi-Site Controller Account to Multiple Pods

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the ACI Multi-Site controller account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site controller account that you want to assign to one or more pods.

Step 4 Click **Assign to Pod**.
The **Assign to Pod** screen appears.

Step 5 Click **Select**.

The list of pods available in Cisco UCS Director is displayed.

Step 6 Check the pods to which you want to assign the ACI Multi-Site controller account to and click **Select**.

Step 7 Click **Submit**.

What to do next

If you want to unassign an account from the pod, choose the account and click **UnAssign from Pod**. In the **Unassign from Pod** screen, click **Select** and choose the pod from which you want to unassign the ACI Multi-Site controller account from. Click **Select** and then click **Submit**.

Managing Users

The Cisco ACI Multi-Site provides access according to a user's role through role-based access control (RBAC).

The following user roles are available in Cisco ACI Multi-Site.

- Power User—A power user can perform all the operations as an admin user.
- Site and Tenant Manager—A site and tenant manager can manage sites, tenants, and associations.

- Schema Manager—A schema manager can manage all schemas regardless of tenant associations.
- Schema Manager - Restricted —A restricted schema manager can manage schemas that contain at least one tenant to which the user is explicitly associated.
- User and Role Manager—A user and role manager can manage all the users, their roles, and passwords.

Creating a User

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **User**.

Step 5 Click **Create User**.

Step 6 On the **Create User** screen, complete the fields, including the following:

- Enter a unique name for the user.
- Enter the password in the **Password** field and **Confirm Password** field.

  The password must at least be six characters in length, and must contain at least one letter, one number, and a special character. Spaces and * are not allowed.
- Enter the first name and last name of the user in the respective fields.
- Enter e-mail address and phone number of the user.
- From the **Account Status** drop-down list, choose **Active** or **Inactive** as the user account status. Only the Active users are authenticated to access ACI Multi-Site.
- Click **Select** and choose the roles to be assigned to the user.

  You must associate at least one role with every user. A user may be associated with more than one role. Associating a user to multiple roles offers a combination of features that the user may access.

Managing Sites

Adding a Site to an ACI Multi-Site Controller Account

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with ACI Multi-Site Controller account to which you want to add a site, and click **View Details**.

Step 4 Click **Site**.

Step 5 Click **Add**.

Step 6 On the **Add Site to ACI Multi-Site** screen, complete the fields, including the following:

a) Add a unique name for the site.

b) Click **Select** and check the site labels that you want to use for the site.

You can choose a maximum of three site labels for the site.

c) Enter the URL of the APIC controller that has to be added as the site object. If you want to add multiple APIC controllers, enter comma separated APIC controller URLs.

Note While adding multiple APIC controllers, ensure that the credentials used for accessing the APIC controllers is same for all the APIC controllers.

d) Enter the user name and password that is used to access the APIC controller.

e) Enter a unique site ID.

Note Once saved, you cannot edit the site ID.

Step 7 Click **Submit**.

Associating a Template to the Site

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site Controller account and click **View Details**.

Step 4 Click **Site**.

Step 5 Click the row with the site to which you want to associate the template and click **Associate Template**.

Step 6 On the **Associate Template to MSC Site** screen, click **Select** and check the template that you want to associate with the site.

Step 7 Click **Submit**.

Managing Tenants

Creating a Tenant

Before you begin

- The site to which the tenant has to be associated must be added.
- The tenant user account must be created.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account to which you want to add a tenant and click **View Details**.

Step 4 Click **Tenant**.

Step 5 Click **Create Tenant**.

Step 6 On the **Create Tenant** screen, complete the fields, including the following:

- Enter a unique account name and description for the tenant.
- Click **Select** and choose one or more sites to which you want to associate the tenant.
- Click **Select** and choose the security domains.
- Click **Select** and choose one or more users who can access the tenant.

  By default, the Admin user is selected.

Step 7 Click **Submit**.

Managing Schemas

Adding a Schema

Schema includes the site-configuration objects that will be pushed to sites. Schemas are the containers for single or multiple templates that are used for defining the policies. Templates are the framework for defining and deploying the policies to the sites.

Before you begin

A tenant must be available in Cisco UCS Director.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account to which you want to add a schema and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click **Add**.

Step 6 On the **Add Schema to ACI Multi-Site** screen, complete the fields, including the following:

- Enter a unique name for the schema.
- Enter a unique name for the template.
- Click **Select** and check the tenant that you want to use.

Step 7 Click **Submit**.

What to do next

You can see the schema in the ACI Multi-Site.

Adding a Template to a Schema

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account to which you want to add a schema template and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a template and click **View Details**.
Step 6 Click **Template**.
Step 7 Click **Add**.
Step 8 On the **Add Template to ACI Multi-Site Schema** screen, complete the fields, including the following:

- Click **Select** and check the tenant that you want to use.
- Enter a display name for the template.

Step 9 Click **Submit**.

What to do next

You can see the template in the ACI Multi-Site.

Deploying a Schema Template to the Site

Before you begin

- The schema template must be assigned to a tenant and the template must be associated to the site. See Associating a Template to the Site, on page 114.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema and click **View Details**.
Step 6 Click the row with the template that you want to deploy to the site and click **Deploy Template**.
Step 7 Click **Submit** to confirm the deployment of template to the site.

You can view the status of the template deployment status under the **Deployed Status** tab (**ACLI Multi-site Account** > **Schema** > **Sites**).

To undeploy a template from the site, navigate to the **Sites** tab (**ACLI Multi-site Account** > **Schema**), choose the row with the site from which you want to undeploy the template from, and then choose **Undeploy Template**.

Adding an ACI Multi-Site Service Graph

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account to which you want to add a service graph and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema to which you want to add a service graph and click **View Details**.

Step 6 Click **Template**.

Step 7 Click the row with the template to which you want to add a service graph and click **View Details**.

Step 8 Click **Service Graph**.

Step 9 Click **Add**.

Step 10 On the **Create ACI Multi-Site Service Graph** screen, complete the following fields:

a) Enter a unique name and description for the service graph.

b) Choose **1**, **2**, or **3** as the number of service nodes for the service graph.

Complete the following fields to add the configurations to the service graph. You have to specify the following details for the number of nodes chosen.

1. Click **Select** and choose a service node that you want to use for the service graph.
2. Click **Select** and choose sites and L4-L7 devices for node that you want to use for the service graph. Ensure that you choose one L4-L7 device per site.

Step 11 Click **Submit**.

Adding a Service Graph to a Contract

Terminal nodes connect a service graph with the contracts. You can insert a service graph for the traffic between two application endpoint groups (EPGs) by connecting the terminal node to a contract. Once connected, traffic between the consumer EPG and provider EPG of the contract is redirected to the service graph.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema to which you want to add the service graph contract and click **View Details**.

Step 6 Click **Template**.

Step 7 Click the row with the template to which you want to add the service graph contract and click **View Details**.

Step 8 Click **Contract**.

Step 9 Click the row with the contract to which you want to add the service graph contract and click **View Details**.

Step 10 Click **Service Graph**.

Step 11 Click **Add**.

Step 12 On the **Add Service Graph to Contract** screen, complete the following fields:

a) Click **Select** and choose a service graph that you want to add to the contract.

b) Complete the following fields to add the node configuration for template. You have to specify the following details for number of nodes in chosen service. For example, if you have chosen service graph with two nodes, you have to specify the following details for two nodes.

1. Click **Select** and choose a bridge domain for consumer connector.
2. The **Route Peering** check box appears for consumer connector of node 2 and node 3 when you choose service graph with two and three nodes. Check this check box to enable route peering on a service appliance such as a load balancer or a firewall to advertise it's reachability through the ACI fabric.
3. The **Route Peering** check box appears for provider connector of node 1 and node 2 when you choose service graph with two and three nodes. Check this check box to enable route peering on a service appliance such as a load balancer or a firewall to advertise it's reachability through the ACI fabric.
4. Click **Select** and choose a bridge domain for provider connector.

c) Complete the following fields to add the node configuration for site. You have to specify the following details for number of nodes in chosen service. If you have chosen two nodes service graph, you have to specify the following details for two nodes.

1. Click **Select** and choose a cluster interface for consumer connector.
2. Click **Select** and choose a redirect policy for consumer connector.
3. Click **Select** and choose a cluster interface for provider connector.
4. (Optional) Click **Select** and choose a redirect policy for provider connector.

Step 13 Click **Submit**.

Adding an Application Profile to a Schema Template

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema and click **View Details**.

Step 6 Click **Template**.

Step 7 Click the row with the schema template to which you want to add an application profile and click **View Details**.

Step 8 Click **Application Profile**.

Step 9 Click **Add**.

Step 10 Enter a display name for the application profile in the **Display Name** field.

Step 11 Click **Submit**.

What to do next

You can see the application profile in the ACI Multi-Site.

Adding a VRF to a Schema Template

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the schema template to which you want to add a VRF and click **View Details**.
Step 8 Click **VRFs**.
Step 9 Click **Add**.
Step 10 Enter a display name for the VRF in the **Display Name** field.
Step 11 Click **Submit**.

What to do next

You can see the VRF in the ACI Multi-Site.

Adding a Contract to the Template

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a contract and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template to which you want to add a contract and click **View Details**.
Step 8 Click **Contract**.
Step 9 Click **Add**.
Step 10 On the **Add Contract to ACI Multi-Site Schema** screen, complete the fields, including the following:

- Enter a unique display name for the contract.
- From the **Scope** drop-down list, choose application profile, VRF, tenant or global as the scope of the contract.

Step 11 Click **Submit**.

Adding a Contract to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a contract and click **View Details**.
Step 6 Click the row with the template to which you want to add a contract and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add a contract and click **View Details**.

The EPGs available in the application profile is displayed.

Step 9 Click the row with the EPG to which you want to add a contract.
Step 10 From the **More Actions** drop-down list, choose **Add Contract to EPG**.
Step 11 In the **Add Contract to ACI Multi-Site EPG** screen, complete the fields, including the following:

- Click **Select** and choose the contract that you want to add to the EPG.
- From the **Type** drop-down list, choose **Consumer** or **Provider**.

Step 12 Click **Submit**.

Adding Multiple Contracts to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a contract and click **View Details**.
Step 6 Click the row with the template to which you want to add a contract and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add a contract and click **View Details**.

The EPGs available in the application profile is displayed.

Step 9 Click the row with the EPG to which you want to add a contract.
Step 10 From the **More Actions** drop-down list, choose **Add Multiple Contracts to EPG**.
Step 11 In the **Add Multiple Contracts to ACI Multi-Site EPG** screen, complete the fields, including the following:

- Click **Select** and choose multiple contracts that you want to add to the EPG.
- From the **Type** drop-down list, choose **Consumer** or **Provider**.

Step 12 Click **Submit**.

Adding a Domain to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a domain and click **View Details**.
Step 6 Click the row with the template to which you want to add a domain and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add a domain and click **View Details**.

The EPGs available in the application profile is displayed.

Step 9 Click the row with the EPG to which you want to add a domain.
Step 10 From the **More Actions** drop-down list, choose **Add Domain to EPG**.
Step 11 On the **Add Domain to EPG** screen, complete the fields, including the following:

- Click **Select** to choose the site.
- From the **Domain Association Type** drop-down list, choose the type in which you want to associate the domain.
- Click **Select** to choose the domain profile.
- From the **Deployment Immediacy** drop-down list, choose one of the following options:
  - **Immediate**—Specifies that the policy is programmed in the hardware policy CAM as soon as it is downloaded in the leaf software.
  - **On Demand**—Specifies that the policy is programmed in the hardware policy CAM only when the first packet is received through the data path. This process helps to optimize the hardware space.
- From the **Resolution Immediacy** drop-down list, one of the following options:
  - **Immediate**—Specifies that EPG policies (including contracts and filters) are downloaded to the associated leaf switch software upon hypervisor attachment to VDS. LLDP or OpFlex permissions are used to resolve the hypervisor to leaf node attachments.
  - **On Demand**—Specifies that a policy (for example, VLAN, VXLAN bindings, contracts, or filters) is pushed to the leaf node only when a hypervisor is attached to VDS and a VM is placed in the port group (EPG).
  - **Pre-provision**—Specifies that a policy (for example, VLAN, VXLAN binding, contracts, or filters) is downloaded to a leaf switch even before a hypervisor is attached to the VDS. Therefore, this option pre-provisions the configuration on the switch.

- **Vlan Mode** drop-down list—This field appears only when VMM is chosen in the **Domain Association Type** drop-down list. Choose **Dynamic** to assign VLAN identifiers to EPG dynamically by the APIC or **Static** to assign VLAN identifiers to EPG manually by an administrator.
- **Allow Micro-segmentation** check box—This field appears only when VMM is chosen in the **Domain Association Type** drop-down list. Check this box to use Cisco APIC configure Micro-segmentation with Cisco ACI to put VMs that belong to different base EPGs or the same EPG into a new attribute-based EPG.

Step 12 Click **Submit**.
The domain is associated with the EPG. You can see the domains that are associated with the EPG under the **EPG Domain Association** tab (**ACI Multisite Account** > **Schema** > **Sites** > **Application Profile**).

Adding a Static Port to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a static port and click **View Details**.
Step 6 Click the row with the template to which you want to add a static port and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add a static port and click **View Details**.
The EPGs available in the application profile is displayed.
Step 9 Click the row with the EPG to which you want to add a static port.
Step 10 From the **More Actions** drop-down list, choose **Add Static Port to EPG**.
Step 11 On the **Create STATIC PORT** screen, complete the fields, including the following:

- Click **Select** to choose the site.
- From the **Path Type** drop-down list, choose port, virtual port channel or direct port channel as the static port path.
- The Leaf field appears only when you choose port as the path type. Click **Select** and choose the leaf for the static port.
- Click **Select** and choose the static port path.
- In the **Port Encap VLAN** field, enter the port encapsulation VLAN ID.
- In the **Primary Micro Seg VLAN** field, enter the primary micro segment VLAN ID.
- From the **Deployment Immediacy** drop-down list, choose on-demand or immediate as the deployment type.
- From the **Mode** drop-down list, choose the mode in which the static port has to be created.

Step 12 Click **Submit**.

You can view the static port added to EPG under the **EPG** tab (**ACI Multisite Account** > **Schema** > **Sites** > **Application Profile**). You also have options to update and delete the static port under the **EPG** tab (**ACI Multisite Account** > **Schema** > **Sites** > **Application Profile**).

Adding a Static Leaf to the EPG

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with ACI Multi-Site Controller account to which you want to add a static leaf and click **View Details**.

Step 4 Click the row with the schema to which you want to add a static leaf and click **View Details**.

Step 5 Click the row with the template to which you want to add a static leaf and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the application profile to which you want to add a static leaf and click **View Details**.

Step 8 Click **EPG**.

Step 9 Click the row with the EPG to which you want to add a static leaf.

Step 10 From the **More Action** drop-down list, choose **Add Static Leaf to EPG**

Step 11 On the **Add Static Leaf to EPG in Multi-Site Schema** screen, complete the following fields.

- Click **Select** and choose a site.
- Click **Select** to view the list of available leaves. Choose a leaf that you want to add to the EPG in the ACI Multi-site schema and click **Select**.
- Enter the ID of VLAN on leaf.

Step 12 Click **Submit**.

Creating an ACI Multi-Site Bridge Domain

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with ACI Multi-Site Controller account to which you want to add a bridge domain and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema to which you want to add a bridge domain and click **View Details**.

Step 6 Click the row with the template to which you want to add a bridge domain and click **View Details**.

Step 7 Click **Bridge Domain**.

Step 8 Click **Add**.

Step 9 On the **Create ACI Multi-Site Bridge Domain** screen, complete the fields, including the following:

- Enter a unique name for the bridge domain.

- Click **Select** and check the virtual routing and forwarding that you want to use for the bridge domain.
- Check the **L2 Stretch** check box to apply stretched layer-2 or VLAN extension on the bridge domain.
- Check the **Inter Site Bum Traffic Allow** check box to allow inter site bum traffic. This field appears only when the **L2 Stretch** check box is checked.
- Check the **Optimize WAN Bandwidth** check box to optimize the WAN bandwidth in bridge domain. This field appears only when the **Inter Site Bum Traffic Allow** check box is checked.
- From the **L2 Unknown Unicast** drop-down list, choose proxy or flood.

Step 10 Click **Submit**.

Adding a Layer 3 Out to the Site Bridge Domain

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with ACI Multi-Site Controller account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template and click **View Details**.
Step 8 Click **Bridge Domain**.
Step 9 Click the row with the bridge domain to which you want to add a layer 3 Out.
Step 10 On the **Add L3 OUT to Site Bridge Domain** screen, complete the fields, including the following:

- Click **Select** and check the site that you want to use.
- Click **Select** and check the check box of the layer 3 Out that you want to add to the site bridge domain.

Step 11 Click **Submit**.

Adding a Subnet to an ACI Multi-Site Bridge Domain

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with ACI Multi-Site Controller account to which you want to add a bridge domain subnet and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a bridge domain subnet and click **View Details**.
Step 6 Click the row with the template to which you want to add a bridge domain subnet and click **View Details**.

Step 7 Click **Bridge Domain**.

Step 8 Click the row with the bridge domain to which you want to add a subnet and click **View Details**.

Step 9 Click **Subnets**.

Step 10 Click **Add Subnet to Site Bridge Domain**.

Step 11 On the **Add Subnet to ACI Multi-Site Bridge Domain** screen, complete the fields, including the following:

- Click **Select** to view the list of available sites. Choose the site to which you want to add the bridge domain subnet and click **Select**.
- Enter the gateway IP address and a description for the subnet you intend to add.
- In the **Scope** field, select either **Private to VRF** or **Advertised Externally**.
- Click the check box for **Shared Between VRFs** if appropriate.
- Click the check box for **No Default SVI Gateway** if appropriate. If checked, the pervasive SVI will not be configured for this subnet and it is used to leak more specific prefix routes to other VRFs.

Step 12 Click **Submit**.

Adding an EPG to the Template

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema to which you want to add a contract and click **View Details**.

Step 6 Click **Application Profile**.

Step 7 Click the row with the application profile to which you want to add a contract and click **View Details**.

The EPGs available in the application profile is displayed.

Step 8 Click **Add**.

Step 9 On the **Create ACI Multi-Site EPG** screen, complete the fields, including the following:

- Enter a unique name for the EPG.
- Check the **USEG EPG** check box to consider the EPG as uSeg EPG. If unchecked, the EPG is considered as base EPG.
- From the **Intra EPG Isolation** drop-down list, choose unenforced or enforced.

  If an EPG is configured with intra-EPG endpoint isolation enforced, the following restrictions apply:

  - In ACI Multi-Site, intra-EPG isolation is not supported in AVS-VLAN mode. Setting Intra-EPG isolation to be enforced may cause the ports to go into a blocked state in these domains.
  - Intra-EPG isolation is not supported if the Bridge Domain is configured as "legacy BD mode".

- Preserving QoS CoS priority settings is not supported when traffic is flowing from an EPG with isolation-enforced to an EPG without isolation enforced.

- **Forwarding Control Proxy-ARP** check box—This field appears only when enforced is chosen in the **Intra EPG Isolation** field. Check the box to enable proxy ARP. The proxy ARP in Cisco ACI enables endpoints within a network or subnet to communicate with other endpoints without knowing the real MAC address of the endpoints.
- Click **Select** and choose a bridge domain from the list.

Step 10 Click **Submit**.

Adding a Filter to the Template

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a filter and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template to which you want to add a filter and click **View Details**.
Step 8 Click **Filters**.
Step 9 Click **Add**.
Step 10 On the **Create ACI Multi-Site Filter** screen, enter the display name of the filter.
Step 11 Click **Submit**.

Adding an Entry to an ACI Multi-Site Filter

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a filter entry and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template to which you want to add a filter entry and click **View Details**.
Step 8 Click **Filters**.
Step 9 Click the row with the filter to which you want to add a filter entry and click **View Details**.
Step 10 Click **Add**.
Step 11 On the **Add Entry to Filter** screen, complete the fields, including the following:

- Enter a unique name and description for the filter entry.
- Choose the type, ARP flag and IP protocol from the respective drop-down lists.
- Check the **Match Only Fragments** check box to match only fragments during filtering.
- Check the **Stateful** check box to enable stateful connection.
- Enter the starting and ending range of the source port number in the **Source Port range From** field and **Source Port Range To** field.
- Enter the starting and ending range of the destination port number in the **Destination Port range From** field and **Destination Port Range To** field.
- Choose one of the following as the TCP session rules:
  - Acknowledgment
  - Established
  - Finish
  - Synchronize
  - Reset
  - Unspecified

Step 12 Click **Submit**.

Adding an uSeg Attribute to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add an uSeg attribute and click **View Details**.
Step 6 Click the row with the template to which you want to add an uSeg attribute and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add an uSeg attribute and click **View Details**.

The EPGs available in the application profile is displayed.

Step 9 Click the row with the EPG to which you want to add an uSeg attribute and click **View Details**.

Note You can add the uSeg attribute to the EPG for which the USEG EPG is enabled.

Step 10 Click **USEG Attributes**.
Step 11 Click **Add**.
Step 12 On the Add uSeg Attribute to EPG in ACI Multi-Site Schema screen, complete the fields, including the following:

- Enter a unique name and description for the uSeg attribute.
- From the **Attribute Type** drop-down list, choose an attribute type. According to the selected attribute type, additional fields will be displayed. Enter the attribute values.

Step 13 Click **Submit**.

Adding a Subnet to the EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a subnet and click **View Details**.
Step 6 Click the row with the template to which you want to add a subnet and click **View Details**.
Step 7 Click **Application Profile**.
Step 8 Click the row with the application profile to which you want to add a subnet and click **View Details**.
Step 9 Click the row with the EPG to which you want to add a subnet and click **View Details**.
Step 10 Click **Subnets**.
Step 11 Click Add.
Step 12 On the **Add Subnet to ACI Multi-Site EPG** screen, complete the fields, including the following:

- Enter the gateway IP address and short description for the subnet.

  Note The gateway IP address must be entered in the format: <Valid IP address>/<Valid Prefix Length>

- From the **Scope** drop-down list, choose **Private to VRF** or **Advertised Externally**.
- Check the **Shared Between VRFs** and **No Default SVI Gateway** check boxes as required.

Step 13 Click **Submit**.

Adding a Subnet to the Site EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add a subnet and click **View Details**.
Step 6 Click the row with the template to which you want to add a subnet and click **View Details**.
Step 7 Click **Application Profile**.

Step 8 Click the row with the application profile to which you want to add a subnet and click **View Details**.

Step 9 Click the row with the EPG to which you want to add a subnet and click **View Details**.

Step 10 Click **Subnets**.

Step 11 Click **Add Subnet to Site EPG**.

Step 12 On the **Add Subnet to Site EPG** screen, complete the fields, including the following:

- Click **Select** to view a list of available sites. Check the check box of the site to which you want to add the subnet and click **Select**.
- Enter the gateway IP address and short description for the subnet.

  Note The gateway IP address must be entered in the format: <Valid IP address>/<Valid Prefix Length>

- From the **Scope** drop-down list, choose **Private to VRF** or **Advertised Externally**.
- Check the **Shared Between VRFs** and **No Default SVI Gateway** check boxes as required.

Step 13 Click **Submit**.

Adding an External EPG to the Template

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **Schema**.

Step 5 Click the row with the schema to which you want to add an external EPG and click **View Details**.

Step 6 Click **Template**.

Step 7 Click the row with the template to which you want to add an external EPG and click **View Details**.

Step 8 Click **External EPG**.

Step 9 Click **Add**.

Step 10 On the **Create External EPG** screen, complete the following fields:

Note When the template is assigned to the site and the user tries to add an External EPG, the L3 outside must be assigned to the External EPG for the site. If the template is not assigned to site, then specify only the EPG display name and the L3 outside, site, and VRF details are not needed.

a) Enter a unique name for the external EPG.

b) From the **Associated Sites** drop-down list, choose **Single** to associate only one site to EPG or **Multiple** to associate more than one sites to EPG.

When **Single** is chosen, the following fields appear:

1. Click **Select** and check the check box of site that you want to use for the external EPG.
2. Click **Select** and check the check box of Virtual Routing and Forwarding (VRF) that you want to use for the external EPG.
3. Click **Select** and check the check box of the L3Outs on the site to be used for the external EPG.

When **Multiple** is chosen, the following field appears:

1. Click **Select** and check the check box of L3Outs on the sites to be used for the external EPG. Ensure that you choose only one L3Out per site and VRF associated to the L3Outs must be same.

Step 11 Click **Submit**.

Adding a Contract to the External EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add the external EPG contract and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template to which you want to add the external EPG contract and click **View Details**.
Step 8 Click **External EPG**.
Step 9 Click the row with the external EPG to which you want to add a contract and click **View Details**.
Step 10 Click **Contracts**.
Step 11 Click **Add**.
Step 12 On the **Add Contract to ACI Multi-Site External EPG** screen, complete the fields, including the following:

- Click **Select** and choose the contract that you want to add to the external EPG.
- From the **Type** drop-down list, choose consumer or provider as the contract type.

Step 13 Click **Submit**.

Adding a Subnet to the External EPG

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the ACI Multi-Site account and click **View Details**.
Step 4 Click **Schema**.
Step 5 Click the row with the schema to which you want to add the external EPG contract and click **View Details**.
Step 6 Click **Template**.
Step 7 Click the row with the template to which you want to add the external EPG contract and click **View Details**.
Step 8 Click **External EPG**.
Step 9 Click the row with the external EPG to which you want to add a contract and click **View Details**.
Step 10 Click **Subnet**.

Step 11 Click **Add**.

Step 12 On the **Add Subnet to ACI Multi-Site External EPG** screen, enter a unique name for the subnet.

Step 13 Click **Submit**.

Deploying a Template to the Site

This section captures the list of actions that have to be executed to deploy a template to site.

Step 1 Add a site to an ACI Multi-Site controller account.

See Adding a Site to an ACI Multi-Site Controller Account, on page 113.

Step 2 Create a tenant and assign it to the site.

See Creating a Tenant, on page 114.

Step 3 Create a schema and map it to the tenant.

See Adding a Schema, on page 115.

Step 4 Optional. Add a template to the schema.

See Adding a Template to a Schema, on page 116.

Step 5 Add an application profile to the template.

See Adding an Application Profile to a Schema Template, on page 118.

Step 6 Add a VRF to the template.

See Adding a VRF to a Schema Template, on page 119.

Step 7 Add a bridge domain to the template.

See Creating an ACI Multi-Site Bridge Domain, on page 123.

Step 8 Add a subnet to the bridge domain.

See Adding a Subnet to an ACI Multi-Site Bridge Domain, on page 124.

Step 9 Add an EPG to the template.

See Adding an EPG to the Template, on page 125.

Step 10 Add a filter to the template.

See Adding a Filter to the Template, on page 126.

Step 11 Add an entry to filter.

See Adding an Entry to an ACI Multi-Site Filter, on page 126.

Step 12 Add a contract to the template.

See Adding a Contract to the Template, on page 119.

Step 13 Add a contract or multiple contracts to the EPG.

See Adding a Contract to the EPG, on page 120 and Adding Multiple Contracts to the EPG, on page 120.

Step 14 Associate a template to the site.

See Associating a Template to the Site, on page 114.

Step 15 Add a domain to the EPG.

See Adding a Domain to the EPG, on page 121.

Step 16 Add a static port to the EPG.

See Adding a Static Port to the EPG, on page 122.

Step 17 Add an uSeg attribute to the EPG.

See Adding an uSeg Attribute to the EPG, on page 127.

Step 18 Add the external EPG to the template.

See Adding an External EPG to the Template, on page 129.

Step 19 Add a contract to the external EPG.

See Adding a Contract to the External EPG, on page 130.

Step 20 Add a subnet to the external EPG.

See Adding a Subnet to the External EPG, on page 130.

Step 21 Deploy a template to the site.

See Deploying a Schema Template to the Site, on page 116.

Viewing ACI Multi-Site Controller Resources

After creating an ACI Multi-Site controller account in Cisco UCS Director, you can view related resources of the ACI Multi-Site controller account.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click one of the following tabs to view the details of a specific component in the ACI Multi-Site controller:

- **Summary** tab—Displays the system overview and control plane BGP of the ACI Multi-Site controller.
- **Site** tab—Displays a list of sites that are configured in the ACI Multi-Site controller. To associate a template to a site, select the site and then click **Associate Template**. In the **Associate Template to MSC Site** screen, click **Select** and check the template that you want to associate with the site. To disassociate a template from the site, select the site and then click **Disassociate Template**. In the **Disassociate Template from MSC Site** screen, click **Select** and check the template that you want to disassociate from the site.

- **Tenant** tab—Displays a list of tenants that are available in the ACI Multi-Site controller. To create a tenant in a site, click **Create Tenant (+)**. On the **Create Tenant** screen, do the following:
  - Enter a unique account name and description for the tenant.
  - Click **Select** and choose associated sites, security domains, and associated users for the tenant.
  - Click **Submit**.
- **User** tab—Displays a list of ACI Multi-Site account users. To create a user for a site, click **Create User**. For more details, see Creating a User, on page 113.
- **Schema** tab—Displays a list of schemas that are defined for the ACI Multi-Site controller. To add a schema to an ACI Multi-site, click **Add (+)**. On the **Add Schema to ACI Multi-Site** screen, enter the schema and template name, and then select a tenant.

  To view more details about schema, choose a schema and click **View Details**. The following tabs appear:
  - **Template**—Displays the templates associated with the schema. To add a template, click **Add**. To deploy a template to a site, click **Deploy Template**. To import APIC policy to ACI Multi-site schema, click **Import**. On the **Import APIC Policy to ACI Multi-site Schema** screen, click **Select** and choose the site ID, application profile name, EPG name, contract name, filter name, VRF name, and bridge domain name.

    To view more details about template, choose a template and click **View Details**. The following tabs appear:
    - **Contract**—Displays the contracts that are added to an ACI Multi-site schema. To add a contract, click **Add (+)**, enter the contract display name, and choose **application profile**, **vrf**, **tenant** or **global** as the scope of the contract. To view and add filters to contract in the schema, choose the row with the contract and click **View Details**. Click **Add**. On the **Add Filter to Contract in ACI Multi-Site Schema** screen, click **Select** and choose the filter that you want to add to the contract and choose **none** or **log** as **Directive**.
    - **Application Profile**—Displays the application profiles of the ACI Multi-site. To add an application profile, click **Add** and enter the display name of the application profile. To view the EPGs of the site, choose the row with an application profile and click **View Details**. Choose an EPG and click **View Details** to view the EPG contract association, USEG attributes, and Subnets.

      In the EPG tab, you can perform the following tasks:
      - To add a contract to the EPG, choose **Add Contract to EPG** from the **More Actions** drop-down list. For more details, see Adding a Contract to the EPG, on page 120.
      - To add a domain to the EPG, choose **Add Domain to EPG** from the **More Actions** drop-down list. For more details, see Adding a Domain to the EPG, on page 121.
      - To add a static port to the EPG, choose **Add Static Port to EPG** from the **More Actions** drop-down list. For more details, see Adding a Static Port to the EPG, on page 122.
      - To add a static port to the EPG, choose **Add Static Leaf to EPG** from the **More Actions** drop-down list. For more details, see Adding a Static Leaf to the EPG, on page 123.
    - **VRFs**—Displays the Virtual Routing and Forwarding (VRF) instances of the site. To add a VRF, click **Add** and enter the display name of the VRF.
    - **Bridge Domain**—Displays the bridge domains of the site. To add a bridge domain, click **Add**. For more details, see Creating an ACI Multi-Site Bridge Domain, on page 123. To add a layer 3 Out to the site bridge domain, choose a row with the bridge domain and click **Add L3 OUT to Site Bridge Domain**. For more details, see Adding a Layer 3 Out to the Site Bridge Domain, on page 124. To view the subnets of the

bridge domain, choose a row with the bridge domain and click **View Details**. For more details, see Adding a Subnet to an ACI Multi-Site Bridge Domain, on page 124.

- **Filters**—Displays the filters of the site. To add a filter, click **Add** and enter the display name of the filter. To view the filter entries, choose a row with the filter and click **View Details**. Click **Add** to add filter entries. For more details, see Adding an Entry to an ACI Multi-Site Filter, on page 126.
- **External EPG**—Displays the external EPGs of the site. To add an external EPG, click **Add** and enter the external EPG Name. Choose a row with the external EPG and click **View Details** to view the contracts and subnets of the external EPG. You can also add contract and subnet to the external EPG, under the respective tab.

- **Sites**—Displays the sites that are associated with the schema template. To undeploy a template from the site, choose the row with the site from which you want to undeploy the template from and then choose **Undeploy Template**. The template will be undeployed from the site after confirmation.

  To view more details about sites, choose a site and click **View Details**. The following tabs appear:

  - **Application Profile**—Displays the application profiles of the site. Choose a row with an application profile and click **View Details**, to view the EPG associated with the application profile.

    To view more details about EPGs, choose an EPG and click **View Details**. The following tabs appear:

    - **EPG Domain Association**—You can view and update the domains that are associated with the EPG under the **EPG** tab (**Schema > Template > Application Profile**).
    - **Static Port**—You can view and update the static ports of the EPG.
    - **Subnets**—You can view and update the subnets of the EPG. To add a subnet to the site EPG, click **Add**. In the **Add Subnet to Site EPG** screen, enter the gateway IP address and short description for the subnet. From the **Scope** drop-down list, choose **Private to VRF** or **Advertised Externally**. Check the **Shared Between VRFs** and **No Default SVI Gateway** check boxes as required.
    - **Static Leaf**—You can view static leaves of the EPG.

  - **Bridge Domain**—You can view the bridge domains of the site. Choose a row with a bridge domain and click **View Details**, to view the subnets of the bridge domain.
  - **Deployed Status**—You can view the status of the template deployment of the site.
  - **External EPG**—You can view and update the external EPG of the site.

- **OSPF Policies** tab—Displays the OSPF policies of the site. To add a OSPF policy, click **Add**. For more details, see Creating an OSPF Policy, on page 135. To configure control plane BGP, click **Configure Control Plane BGP**. For more details, see Configuring Control Plane BGP, on page 136.
- **Site Settings** tab—Displays the settings of all sites in the ACI Multi-Site account. To view more details about site settings, choose a site setting and click **View Details**. The pods associated with the site is displayed. Choose a pod and click **View Details** to view the spines. Choose a spine and click **View Details** to view the ports of the spine. To add a port to the spine, click **Add**. On the **Add Port to Spine** screen, do the following:
  - In the **Port ID** field, enter the port ID in the slot number/port number format For example, 1/10.
  - In the **IP address** field, enter the IP address in the valid IP address/valid prefix length format.
  - In the **MTU** field, enter the MTU. The range is 576 to 9000 or inherit.

- In the **OSPF Policy** field, click **Select** and choose the OSPF policy.
- Click **Submit**.

Note You can update the site settings, pod, and spine using the **Update** option available under respective tabs.

Creating an OSPF Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account and click **View Details**.

Step 4 Click **OSPF Policies**.

Step 5 Click **Create**.

Step 6 On the **Create OSPF Policy** screen, complete the fields, including the following:

- Enter a unique name for the OSPF policy.
- From the **Network Type** drop-down list, choose broadcast, point-to-point, or unspecified as the interface type.
- Enter the routing device's priority in the range of 0 to 255, for becoming the designated router. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. The default value is 1.
- Enter the cost of an OSPF interface in the range of 0 to 65535. The cost is a routing metric that is used in the link-state calculation. The default value is 0.
- Choose one of the following as the interface control:
  - **Advertise-subnet**—To advertise subnet.
  - **BFD**—To enable Bidirectional Forwarding Detection (BFD) at the interface.
  - **MTU-ignore**—To ignore any IP MTU mismatch with neighbors.
  - **Passive-participation**—To suppress routing updates on the interface.
- In the **Hello Interval (SECONDS)** field, enter how often the routing device sends hello packets out the interface. The hello interval must be in the range of 1 to 65535. The default value is 10.
- In the **Dead Interval (SECONDS)** field, enter how long OSPF waits before declaring that a neighboring routing device is unavailable. The dead interval must be in the range of 1 to 65535. The default value is 40.
- In the **Retransmit Interval (SECONDS)** field, enter how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors. The retransmit interval must be in the range of 1 to 65535. The default value is 5.
- In the **Transmit Delay (SECONDS)** field, enter the estimated time required to transmit a link-state update on the interface. The transmit delay must be in the range of 1 to 450. The default value is 1.

Step 7 Click **Submit**.

Configuring Control Plane BGP

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the ACI Multi-Site account for which you want to configure the control plane BGP.

Step 4 Click **Configure Control Plane BGP**.

Step 5 On the **Configure Control Plane BGP** screen, complete the fields, including the following:

- From the **Bgp Peering Type** drop-down list, choose full-mesh or route-reflector. The default value is full-mesh.
- In the **Keep Alive Interval (SECONDS)** field, enter the keep alive value to retain the route information learned from BGP in the routing table. The keep alive interval must be in the range of 0 to 3600. The default value is 60.
- In the **Hold Interval (SECONDS)** field, enter the hold-time value to use when negotiating a connection with the peer. The hold interval must be in the range of 0 to 3600. The default value is 180.
- In the **Stale Interval (SECONDS)** field, enter the period of time for which stale routes must be preserved by using the long-lived graceful restart capability for BGP sessions on the restarting router. The stale interval must be in the range of 1 to 3600. The default value is 300.
- Check **Graceful Helper** to enable or turn on the helper mode to assist a neighboring router attempting a graceful restart.
- In the **Maximum AS Limit** field, enter the maximum allowed number of autonomous system (AS) in the range of 0 to 2000. The default value is 0.
- In the **Bgp TTL Between Peers** field, enter the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets in the range of 1 to 255. The default value is 10.

Step 6 Click **Submit**.
You can view the control plane BGP configured for the site under the **Summary** tab (**ACI Multi-Site Account** > **Summary**).

Generating the ACI Multi-Site Troubleshooting Report

For troubleshooting issues that you may face in ACI Multi-Site, you can generate the troubleshooting report and the infrastructure log file for all the schemas, sites, tenants, and users that are managed by ACI Multi-Site.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with ACI Multi-Site Controller account for which you want to generate troubleshooting report.

Step 4 Click **Download Troubleshooting Report**.

Step 5 On the **ACI Multi-Site Troubleshooting Report Download** screen, check the check box of the objects for which you want to generate and view the report:

- Sites—Site definitions in the JSON format.
- Tenants—Tenant definitions in the JSON format.
- Schemas—All schemas available in the Multi-Site in the JSON format.
- Users—User definitions in the JSON format
- Infra Logs—Logs of the containers in the infra_logs.txt file.

Step 6 Click **Submit**.
Cisco UCS Director will fetch the report for selected objects from the ACI Multi-Site.

Step 7 Click the download link to download the troubleshooting report. If you want to change the objects that you have chosen for report generation, do the necessary changes and click **Generate Download Link**.

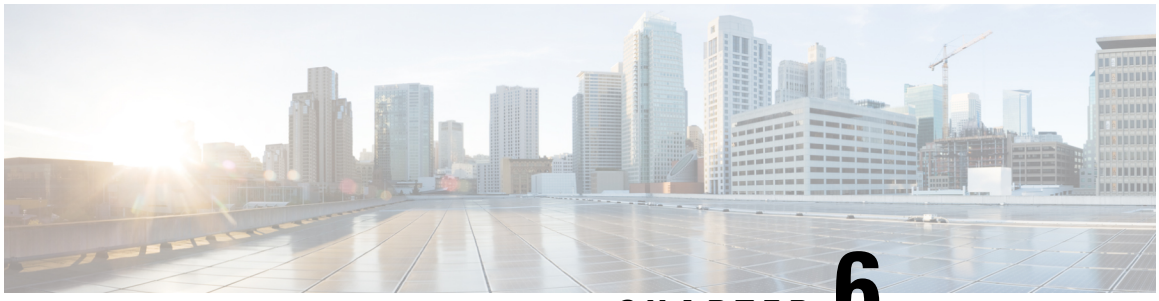

CHAPTER 6

Configuring L4-L7 Services

Unmanaged Mode

In unmanaged mode, Cisco APIC allocates only the network resources for the service graph and programs only the fabric side during graph instantiation. You must configure the unmanaged device in an external application or tool.

When you add an unmanaged network device, Cisco APIC does not require the device package for that device.

For more information about unmanaged mode, see the Cisco APIC Layer 4 to Layer 7 Services Deployment Guide.

Setting Up an Unmanaged Device

For unmanaged devices, the Application Policy Infrastructure Controller (APIC) allocates only the network resources for the service graph and programs only the fabric side during graph instantiation. You cannot configure an unmanaged device in the APIC.

Step 1 Add an unmanaged device cluster.

See Adding an Unmanaged L4-L7 Devices, on page 141.

Step 2 Add at least one concrete device to the unmanaged device cluster.

See Adding a Concrete Device to an Unmanaged L4-L7 Device Cluster, on page 145.

Step 3 For a virtual concrete device, add at least one vNIC to the concrete interface on the device.

See Adding a vNIC to an Unmanaged Virtual Concrete Device, on page 146.

Step 4 For a physical concrete device, add at least one path to the concrete interface on the device.

See Adding a Path Interface to an Unmanaged Physical Concrete Device, on page 147.

Step 5 Add at least one logical interface to the unmanaged device cluster.

See Adding a Cluster Interface to an Unmanaged L4-L7 Device, on page 143.

Step 6 Create an L4-L7 service graph template with the configuration parameters you want to use for the device cluster.

See Creating a Service Graph Template, on page 157.

Step 7 Apply the L4-L7 service graph template to configure the device cluster.

See Applying a Service Graph Template, on page 158.

Managed Mode

By default, when a device is registered with Cisco APIC, the device is set to be in managed mode. When a device is configured as managed, Cisco APIC manages the device and programs the device during graph instantiation.

Setting Up a Managed Device

Step 1 Add a managed device cluster.

See Adding a Managed L4-L7 Device, on page 142.

Step 2 Add at least one concrete device to the managed device cluster.

See Adding a Concrete Device to a Managed L4-L7 Device Cluster, on page 145.

Step 3 For a virtual concrete device, add at least one vNIC to the concrete interface on the device.

See Adding a vNIC to a Managed Virtual Concrete Device, on page 146.

Step 4 For a physical concrete device, add at least one path to the concrete interface on the device.

See Adding a Path Interface to a Managed Physical Concrete Device, on page 147.

Step 5 Add at least one logical interface to the managed device cluster.

See Adding a Cluster Interface to a Managed Device Cluster, on page 144.

Step 6 Create an L4-L7 service graph template with the configuration parameters you want to use for the device cluster.

See Creating a Service Graph Template, on page 157.

Step 7 Apply the L4-L7 service graph template to configure the device cluster.

See Applying a Service Graph Template, on page 158.

Layer 4 to Layer 7 Device Clusters

A L4-L7 device cluster, also known as a logical device, contains one or more concrete devices that act as a single device. A L4-L7 device cluster has cluster interfaces, also known as logical interfaces, which describe the interface information for the device cluster.

L4-L7 device clusters can be managed or unmanaged.

When the Application Policy Infrastructure Controller (APIC) renders and instantiates service graph templates, it does the following:

- Associates function node connectors with the cluster interfaces.
- Allocates network resources for a function node connector, such as VLAN or Virtual Extensible Local Area Network (VXLAN) resources
- Programs those network resources onto the cluster logical interfaces

The service graph template uses a specific device that is based on a device selection policy, known as a logical device context.

Each device cluster can have a maximum of two concrete devices in active/standby mode.

Adding an Unmanaged L4-L7 Devices

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click **Add**.

Step 8 On the **Create L4-L7 Devices** screen, complete the following fields:

a) Ensure that **Managed** is not checked.

b) In the **Device Cluster** field, enter a unique name for the cluster.

c) Choose **True** from the **Promiscuous Mode** drop-down list, to enable promiscuous mode. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or to the Services Processing Unit (SPU) regardless of the destination MAC address of the packet.

d) From the **Context Aware** drop-down list, choose one of the following:

- **Single**—The device cluster must be given to a specific tenant and cannot be shared across multiple tenants. This is the default value.
- **Multiple**—The device cluster can be shared across multiple tenants of a given type that you are hosting on a provider network.

e) From the **Function Type** drop-down list, choose one of the following:

- **Go To**—A Go To device has a specific destination. This is the default value.

- **Go Through**—A Go Through device is a transparent device. A packet goes through the device without being addressed to it, and the endpoints are not aware of the device.

f) From the **Service Type** drop-down list, choose one of the following:

- **ADC**—One-arm and two-arm deployment modes.
- **Firewall**—Routed and transparent deployment modes.
- **IDS/IPS**—IDS and IPS deployment modes.
- **Other**—Any other mode.

g) From the **Device Type** drop-down list, choose one of the following:

- **Physical**
- **Virtual**

h) Click **Select** and check the domain that you want to use.

The domain must be a VMM domain if the device type is Virtual, and a physical domain if the device type is Physical.

Step 9 Click **Submit**.

Adding a Managed L4-L7 Device

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Devices**.
Step 7 Click **Add**.
Step 8 On the **Create L4-L7 Devices** screen, complete the following fields:

a) Ensure that **Managed** is checked.
b) In the **Device Cluster** field, enter a unique name for the cluster.
c) Click **Select** and check the device package that you want to use.
d) Click **Model** and check the device package model that you want to use.
e) Choose **True** from the **Promiscuous Mode** drop-down list, to enable promiscuous mode. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or to the Services Processing Unit (SPU) regardless of the destination MAC address of the packet.
f) From the **Context Aware** drop-down list, choose one of the following:

- **Single**—The device cluster must be given to a specific tenant and cannot be shared across multiple tenants. This is the default value.
- **Multiple**—The device cluster can be shared across multiple tenants of a given type that you are hosting on a provider network.

g) From the **Function Type** drop-down list, choose one of the following:

- **Go To**—A Go To device has a specific destination. This is the default value.
- **Go Through**—A Go Through device is a transparent device. A packet goes through the device without being addressed to it, and the endpoints are not aware of the device.

h) From the **Service Type** drop-down list, choose one of the following:

- **ADC**—One-arm and two-arm deployment modes.
- **Firewall**—Routed and transparent deployment modes.
- **IDS/IPS**—IDS and IPS deployment modes.
- **Other**—Any other mode.

i) From the **Device Type** drop-down list, choose one of the following:

- **Physical**
- **Virtual**

j) Click **Select** and check the domain that you want to use.

The domain must be a VMM domain if the device type is Virtual, and a physical domain if the device type is Physical.

k) From the **APIC to Device Management Connectivity** drop-down list, choose the type of connectivity. Choose **Out-of-Band** when you are connecting to a device that is outside of the fabric or **In-Band** when you are connecting to a device through the fabric.

l) Click **Select** and check the EPG that you want to use.

m) Enter the virtual IP address, port, user name, and password for the cluster management interface.

Step 9 Click **Submit**.

Cluster Interface

Adding a Cluster Interface to an Unmanaged L4-L7 Device

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click the row with the unmanaged L4-L7 device that you want to update and click **View Details**.

Check the **Managed** column to determine if the device cluster is managed or unmanaged.

Step 8 Click **Cluster Interface**.

Step 9 Click **Add**.

Step 10 On the **Add Cluster Interface to L4-L7 Device** screen, complete the following fields:

a) Enter a unique name for the logical interface.

b) In the **Encapsulation** field, enter the traffic encapsulation identifiers for the logical interface.

The valid VLAN range for encapsulation is between 1 and 4094.

Step 11 Click **Submit**.

Adding a Cluster Interface to a Managed Device Cluster

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click the row with the managed L4-L7 device cluster that you want to update and click **View Details**.

Check the **Managed** column to determine if the device cluster is managed or unmanaged.

Step 8 Click **Cluster Interface**.

Step 9 Click **Add**.

Step 10 On the **Add Cluster Interface to L4-L7** screen, complete the following fields:

a) Enter a unique name for the logical interface.

b) Click **Select** and check the logical interface type that you want to use for managed L4-L7 device cluster.

Step 11 Click **Submit**.

Concrete Devices

A concrete device has concrete interfaces. When a concrete device is added to a logical device cluster, concrete interfaces are mapped to the logical interfaces. During service graph template instantiation, VLANs and VXLANs are programmed on concrete interfaces that are based on their association with logical interfaces.

You can create multiple cluster interfaces on a concrete device and then specify which cluster interface will be used for the connector in the device selection policy. This cluster interface can be shared by using multiple service graph instantiations.

Adding a Concrete Device to an Unmanaged L4-L7 Device Cluster

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click the row with the unmanaged physical L4-L7 device cluster that you want to update and click **View Details**.

Check the **Device Type** column to determine if the device cluster is physical or virtual.

Step 8 Click **Concrete Device**.

Step 9 Click **Add**.

Step 10 On the **Add Concrete Device to L4-L7** screen, complete the following fields:

a) In the **Device Name** field, enter a unique name for the concrete device.

b) In the **Device Context Label** field, enter the label for the device cluster context.

c) For a virtual device cluster, in the **VM Name** field, enter the name of the VM where the device is hosted.

d) For a virtual device cluster, in the **vCenter Name** field, enter the name of the VMware vCenter where the VM is located.

Step 11 Click **Submit**.

Adding a Concrete Device to a Managed L4-L7 Device Cluster

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4 - L7 Devices**.

Step 7 Click the row with the managed physical L4-L7 device cluster where you want to create the concrete device and click **View Details**.

Check the **Device Type** column to determine if the device cluster is physical or virtual.

Step 8 Click **Concrete Device**.

Step 9 Click **Add**.

Step 10 On the **Add Concrete Device to L4-L7** screen, complete the fields, including the following:

a) In the **Device Name** field, enter a unique name for the concrete device.

b) In the **Device Context Label** field, enter the label for the device cluster context.

c) For a virtual device cluster, in the **VM Name** field, enter the name of the VM where the device is hosted.

d) For a virtual device cluster, in the **vCenter Name** field, enter the name of the VMware vCenter where the VM is located.
e) Enter the virtual IP address, port, user name, and password for the cluster management interface.

Step 11 Click **Submit**.

Adding a vNIC to an Unmanaged Virtual Concrete Device

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Devices**.
Step 7 Click the row with the L4-L7 device cluster that you want to update and click **View Details**.
Step 8 Click **Concrete Device**.
Step 9 Click the row with the concrete device that you want to update and click **View Details**.
Step 10 Click **vNIC to Concrete Interface**.
Step 11 Click **Add**.
Step 12 On the **Add Concrete Interface to Device** screen, complete the fields, including the following:
a) In the **Concrete Interface** field, enter a unique name for the interface.
b) Click **Select** and check the path that you want to add to the interface.
c) In the **vNIC** field, enter the vNIC assigned to this interface.
d) Click **Select** and check the logical interface where you want to add the path.
Step 13 Click **Submit**.

Adding a vNIC to a Managed Virtual Concrete Device

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Devices**.
Step 7 Click the row with the managed L4-L7 device cluster that you want to update and click **View Details**.
Check the **Managed** column to determine if the device cluster is managed or unmanaged.
Step 8 Click **Concrete Device**.

Step 9 Click the row with the concrete device that you want to update and click **View Details**.

Step 10 Click **vNIC to Concrete Interface**.

Step 11 Click **Add**.

Step 12 On the **Add Concrete Interface to Device** screen, complete the following fields:

a) In the **Concrete Interface** field, enter a unique name for the interface.
b) Click **Select** and check the path that you want to add to the interface.
c) In the **vNIC** field, enter the vNIC assigned to this interface.
d) Click **Select** and check the logical interface where you want to add the path.

Adding a Path Interface to an Unmanaged Physical Concrete Device

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click the row with the unmanaged L4-L7 device cluster that you want to update and click **View Details**.

Check the **Managed** column to determine if the device cluster is managed or unmanaged.

Step 8 Click **Concrete Device**.

Step 9 Click the row with the concrete device that you want to update and click **View Details**.

Step 10 Click **Path to Concrete Interface**.

Step 11 Click **Add**.

Step 12 On the **Add Concrete Interface to Device** screen, complete the fields, including the following:

a) In the **Concrete Interface** field, enter a unique name for the interface.
b) Click **Select** and check the path that you want to add to the interface.
c) For virtual device, enter vNIC.
d) Click **Select** and check the logical interface where you want to add the path.

Step 13 Click **Submit**.

Adding a Path Interface to a Managed Physical Concrete Device

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Devices**.

Step 7 Click the row with the managed L4-L7 device cluster that you want to update and click **View Details**.

Check the **Managed** column to determine if the device cluster is managed or unmanaged.

Step 8 Click **Concrete Device**.

Step 9 Click the row with the concrete device that you want to update and click **View Details**.

Step 10 Click **Path to Concrete Interface**.

Step 11 Click **Add**.

Step 12 On the **Add Concrete Interface to Device** screen, complete the following fields:

a) In the **Concrete Interface** field, enter a unique name for the interface.
b) Click **Select** and check the path that you want to add to the interface.
c) For virtual device, enter vNIC.
d) Click **Select** and check the logical interface where you want to add the path.

Step 13 Click **Submit**.

APIC Function Profiles

An APIC function profile provides default values for the parameters of a particular function associated with a device package that is managed by Cisco APIC. You can then include one or more APIC function profiles in an L4-L7 service graph template. For example, you can create a function profile that provides default values for the Cisco ASA firewall function.

In Cisco UCS Director, you create APIC function profiles within function profile groups.

Function profile groups organize function profiles to make it easier to identify the profiles that you want to include in a specific service graph template.

Note Cisco UCS Director supports only the configuration of Cisco ASA devices through APIC function profiles and service graphs. You can create service graph template and function profile with load balancer service. But the parameters that are added for the function profile through individual workflow task, user interface action, and REST API in Cisco UCS Director, are supported for firewall service alone.

Creating an APIC Function Profile Group

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Function Profile Group**.

Step 7 Click **Add**.

Step 8 On the **Add Function Profile Group** screen, enter a name and description for the group and click **Submit**.

What to do next

Add one or more APIC function profiles to the function profile group.

Creating an APIC Function Profile

Note Cisco UCS Director supports only the configuration of Cisco ASA devices through APIC function profiles and service graphs.

Before you begin

- Create an APIC function profile group.
- Create an APIC firewall policy.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Function Profile Group**.

Step 7 Click the row with the group where you want to add a function profile and click **View Details**.

Step 8 Click **Function Profile**.

Step 9 Click **Add**.

Step 10 On the **Create Function Profile** screen, complete the following fields:

a) Add a unique name and description for the function profile.

b) Click **Select** and check the row with the APIC account, device, and function that you want to use.

For example, to configure a firewall for a Cisco ASA 1.2, check a row that has a Device Package Name of CISCO-ASA-1.2 and a Function of Firewall. After you validate your selection, the function displays next to **Function Name**.

c) If you chose a load balancing function, in the **Load Balancer Parameters** area, complete the following fields:

- **External ID**
- **External Netmask**
- **Internal ID**
- **Internal Netmask**
- **Services**—Use a comma-separated list to include multiple services.

- **LB IPv4 IP**

d) Optional. If you chose a firewall function, click **Select** and check the APIC firewall policy that you want to assign to this function profile.

If the list does not include the firewall policy you need, click **Add** to create a new policy.

Alternately, navigate to **Policy** > **Resource Groups** > **APIC Firewall Policy**, and then click **Add** to create a firewall policy.

e) Click **Submit**.

What to do next

Click **View Details** and add one or more parameters to the function profile from **Function Profile Parameters**. After you add the parameters, they are displayed on either **L4L7 Function Profile Parameters** or **Function Profile Function Parameters**.

Adding ACL Parameters to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Function Profile Group**.

Step 7 Click the row with the group where you want to update the function profile and click **View Details**.

Step 8 Click **Function Profile**

Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.

Step 10 Click **Function Profile Parameter**.

Step 11 Choose **Add ACL to Function Profile**.

Step 12 On the **Add ACL to Function Profile** screen, complete the following fields:

a) In the **ACL List Name** field, enter the name of the Access Control List.

b) In the **ACE Name** field, enter the name of the Access Control Entry in the ACL to specify the permit or deny rule for packets.

c) From the **Protocol** drop-down list, choose one of the following protocols:

- **ip**
- **tcp**
- **udp**
- **icmp**

d) Check **Source Any** if you want the ACL to apply to any source IP address.

If you do not check this box, enter a single IP address, an IP address range, or a network address or subnet address in the **Source Address** field.

e) Check **Destination Any** if you want the ACL to apply to any destination IP address.

If you do not check this box, you can enter a single IP address, an IP address range, or a network address or subnet address in the **Destination Address** field.

f) From the **Action** drop-down list, choose one of the following:

- **deny** if you want this ACL to drop the packet.
- **allow** if you want this ACL to forward the packet. The ACL denies all packets that you do not specifically allow.

g) In the **Order** field, enter the order of this entry in the ACL.

Step 13 Click **Submit**.

Adding an Interface to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Function Profile Group**.

Step 7 Click the row with the group where you want to update the function profile and click **View Details**.

Step 8 Click **Function Profile**

Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.

Step 10 Click **Function Profile Parameters**.

Step 11 Choose **Add Interface to Function Profile**.

Step 12 On the **Add Interface to Function Profile** screen, complete the following fields:

a) Enter a unique name for the interface.

b) From the **Type** drop-down list, choose one of the following:

- **External**
- **Internal**

c) In the **IPv4 Address** field, enter the IPv4 address for the interface.

d) In the **Security Level** field, enter the security level for the interface.

The security level can be from 0 (lowest) to 100 (highest). The Cisco ASA uses the security level to determine the type of traffic allowed to and from the interface. For example, you can assign a higher security level to an interface that handles internal traffic and a lower security level to an interface that handles external traffic.

e) Click **Select** and check the bridge group that you want to use for this interface.

f) Click **Select** and check the ACL that you want to use for inbound traffic.
g) Click **Select** and check the ACL that you want to use for outbound traffic.

Step 13 Click **Submit**.

Adding a Bridge Group Interface to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Function Profile Group**.
Step 7 On the **Function Profile Group** page, choose the row with the group where you want to update the function profile and click **View Details**.
Step 8 Click **Function Profile**
Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.
Step 10 Click **Function Profile Parameters**.
Step 11 Choose **Add Bridge Group Interface to Function Profile**.
Step 12 On the **Add Bridge Group Interface to Function Profile** screen, complete the following fields:

- **Bridge Group ID**—Enter an integer between 1 and 100.
- **IPv4 Address Value**—Enter the IPv4 address for the bridge group interface.

Step 13 Click **Submit**.

Adding a Static Route to an Interface on an APIC Function Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Function Profile Group**.
Step 7 Click the row with the group where you want to update the function profile and click **View Details**.
Step 8 Click **Function Profile**
Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.
Step 10 Click **Function Profile Parameters**.
Step 11 Choose **Add Static Route to Interface on APIC Function Profile**.

Step 12 On the **Add Static Route to Interface on APIC Function Profile** screen, complete the following fields:

a) Click **Select** and check the interface you want to update.
b) From the **Type** drop-down list, choose either **IPv4** or **IPv6**.
c) If you chose **IPv4**, complete the following fields:

- **Gateway Address**
- **Network Mask**
- **Network**
- **Metric**

d) If you chose **IPv6**, complete the following fields:

- **Gateway Address**
- **Hop Count**
- **Prefix**
- **Tunneled**

Step 13 Click **Submit**.

Adding a Network Object to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Function Profile Group**.
Step 7 On the **Function Profile Group** page, choose the row with the group where you want to update the function profile and click **View Details**.
Step 8 Click **Function Profile**
Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.
Step 10 Click **Function Profile Parameter**.
Step 11 Choose **Add Network Object to Function Profile**.
Step 12 On the **Add Network Object to Function Profile** screen, complete the following fields:

a) In the **Network Object Name** field, enter a unique name for the network object.
b) From the **Network Object Type** drop-down list, choose one of the following types:

- **FQDN**
- **Host IP Address**
- **IP Address Range**

- **Network IP Address**

c) If you chose **FQDN**, enter the fully qualified domain name for this network object.
d) If you chose **Host IP Address**, enter the IP address that you want to use for this network object.
e) If you chose **IP Address Range**, enter the range of IP addresses that you want to use for this network object.
f) If you chose **Network IP Address**, enter the IP address that you want to use for this network object.

Step 13 Click **Submit**.

Adding a Service Object to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Function Profile Group**.
Step 7 Click the row with the group where you want to update the function profile and click **View Details**.
Step 8 Click **Function Profile**
Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.
Step 10 Click **Function Profile Parameters**.
Step 11 Choose **Add Service Object to Function Profile**.
Step 12 On the **Add Service Object to Function Profile** screen, complete the following fields:
a) In the **Service Object Name** field, enter a unique name for the service object.
b) Enter a description of the service object.
c) In the **Protocol Type** field, enter the IP protocol name or number for the service object.
d) From the **Service Object Type** drop-down list, choose one of the following types:

- **icmp**
- **icmp6**
- **tcp**
- **udp**

After you choose the type, you are prompted to enter additional parameters for that type.

e) If you chose **icmp**, enter the **Code** and **Type** for the service object.
f) If you chose **icmp6**, enter the **Code** and **Type** for the service object.
g) If you chose **tcp**, enter the **High Port**, **Low Port**, and **Operator** for the following fields:

- **TCP Destination**
- **TCP Source**

h) If you chose **udp**, enter the **High Port**, **Low Port**, and **Operator** for the following fields:

- **UDP Destination**—
- **UDP Source**—Enter the **High Port**, **Low Port**, and **Operator**.

Step 13 Click **Submit**.

Creating a NAT Rule for an APIC Function Profile

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Function Profile Group**.
Step 7 Click the row with the group where you want to update the function profile and click **View Details**.
Step 8 Click **Function Profile**.
Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.
Step 10 Click **Function Profile Parameters**.
Step 11 Choose **Create NAT Rule**.
Step 12 On the **Create NAT Rule** screen, complete the following fields:

a) Enter a unique name for the NAT rule.
b) Click **Select** and check the source real object that you want to use.
c) Click **Select** and check the source mapped object that you want to use.
d) From the **Type** drop-down list, choose one of the following:
   - **Static**
   - **Dynamic**
e) Click **Select** and check the destination real object that you want to use.
f) Click **Select** and check the destination mapped object that you want to use.
g) Click **Select** and check the service real object that you want to use.
h) Click **Select** and check the service mapped object that you want to use.
i) In the **DNS** field, enter the IP address or the fully qualified domain name (FQDN) of the DNS server that you want to use.
j) In the **Order** field, enter the order of the rule in an access list.

   The order of the rules in an access list determines how traffic is handled and which rule the Cisco ASA applies to the traffic. For an access list with multiple rules, the Cisco ASA goes through the rules in order and applies the first rule that matches the traffic.

k) In the **Uni-Direction** field, enter unidirectional so that the destination addresses cannot initiate traffic to the source addresses.
l) Click **Select** and check the source interface that you want to use.
m) Click **Select** and check the destination interface that you want to use.

Step 13 Click **Submit**.

Adding a Network Object Group to an APIC Function Profile

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Function Profile Group**.

Step 7 On the **Function Profile Group** page, choose the row with the group where you want to update the function profile and click **View Details**.

Step 8 Click **Function Profile**.

Step 9 On the **Function Profiles** page, choose the profile where you want to add parameters and click **View Details**.

Step 10 Click **Function Profile Parameters**.

Step 11 From **More Actions** drop-down list, choose **Add Network Object Group to Function Profile**.

Step 12 On the **Add Network Object Group to Function Profile** screen, complete the following fields:

a) Enter a unique name and description for the network object group.

b) From the **Network Object Group Type** drop-down list, choose one of the following:

- **Host IP Address**
- **Network Address**
- **Network Object**

c) If you chose **Host IP Address**, enter an IPv4 or IPv6 address for the host.

d) If you chose **Network Address**, enter one of the following:

- An IPv4 address with netmask in the following format: 10.10.10.10/255.255.255.255
- An IPv6 address with prefix in the following format: X:X:X:X:X:X/<0-128>

e) If you chose **Network Object**, click **Select** and check the network objects that you want to include.

Step 13 Click **Submit**.

Service Graph Templates

A service graph template contains configuration parameters, which you can specify through one or more of the following:

- Device package
- EPG

- Application profile
- Tenant context

You can apply a service graph template to multiple devices and ensure that all of those devices have the same configuration.

A function node within a service graph template can require one or more configuration parameters. You can lock the parameter values to prevent any additional changes.

The values of the configuration parameters in a service graph are passed to the device script within the device package. The device script converts the parameter data to the configuration that is downloaded onto the device.

Creating a Service Graph Template

Before you begin

Create at least one function profile for the function and device.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Service Graph**.

Step 7 Click **Create L4 L7 Service Graph Template**.

Step 8 On the **Create L4 L7 Service Graph Template** screen, complete the following fields:

a) Enter a unique name and description for the service graph template.

b) From the **Type** drop-down list, choose the type of template you want to create.

The template type determines which configuration parameters you can include in the service graph template. The template type can be one of the following:

- **Single Node - Firewall in Transparent Mode**—A single node graph that inserts a firewall into the traffic path. The graph configures the firewall in transparent mode without routing. The bridge domain is used for traffic that is derived from the EPGs. You do not need to provide any additional bridge domains.
- **Single Node - Firewall in Routed Mode**—A single node graph that inserts a firewall into the traffic path. The graph configures the firewall in routed mode, which performs the routing. The bridge domain is used for traffic that is derived from the EPGs. You do not need to provide any additional bridge domains.
- **Single Node - ADC in One-Arm Mode**—A single node graph that inserts an ADC into the traffic path. The graph configures the ADC in 1-ARM mode. The bridge domain is used for traffic that is explicitly provided.
- **Single Node - ADC in Two-Arm Mode**—A single node graph that inserts an ADC into the traffic path. The graph configures the ADC in 2-ARM mode without routing. The bridge domain is used for traffic that is derived from the EPGs. You do not need to provide any additional bridge domains.
- **Two Nodes - Firewall in Transparent and ADC in One-Arm Mode**—A two node graph that inserts a firewall and an ADC into the traffic path. The graph configures the firewall in transparent mode without routing and the

ADC in 1-ARM mode. The bridge domain that is used for the traffic from the firewall to the ADC and the ADC to the provider EPG is explicitly provided.

- **Two Nodes - Firewall in Routed and ADC in One-Arm Mode**—A two node graph that inserts a firewall and an ADC into the traffic path. The graph configures the firewall in routed mode and the ADC in 1-ARM mode. The bridge domain that is used for the traffic in to and out of the ADC is explicitly provided.
- **Two Nodes - Firewall in Routed and ADC in Two-Arm Mode**—A two node graph that inserts a firewall and an ADC into the traffic path. The graph configures the firewall in routed mode and the ADC in 2-ARM mode. The bridge domain that is used for the traffic from the firewall to the ADC and the firewall to the consumer EPG is explicitly provided.
- **Two Nodes - Firewall in Transparent and ADC in Two-Arm Mode**—A two node graph that inserts a firewall and an ADC into the traffic path. The graph configures the firewall in transparent mode and the ADC in 2-ARM mode. The bridge domain that is used for the traffic for Firewall to ADC is explicitly provided.

c) Complete the following fields to add the configurations to the service graph template. If you have chosen Two Nodes template type, you have to specify the following details for two nodes.

If you chose a template type with a firewall, the firewall is always Node One, whether you choose a Single Node or Two Nodes template type. If you chose a template type with an ADC, the ADC is Node One for a Single Node template type and Node Two for a Two Nodes template type.

- **Managed**—Specifies whether the device is managed or unmanaged. For unmanaged device, you can enable policy-based route redirect by choosing **true** from **Route Redirect** drop-down list. For a managed device, complete the following fields.
  - **Function Name**—Specifies the virtual function for a managed device. Click **Select** and choose a function name that you want to use. This is a single virtual function on a service device such as a firewall, a load balancer, or an SSL offloading device.
  - **Function Profile**—Specifies the function profile for a managed device. Click **Select** and choose a function profile that you want to use. The profile includes the abstract device configuration, the abstract group configuration, and the abstract function configuration.
- **Route Redirect**—This field is applicable for ADC and Firewall Routed mode. Choose **true** from the drop-down list to enable policy-based route redirect on the ADC or Firewall Routed mode.

Step 9 Click **Submit**.

Applying a Service Graph Template

Before you begin

Depending upon the configuration parameters you plan to use, create the following:

- Consumer EPG or external network
- Provider EPG or external network
- Contract
- Device clusters

- Cluster interfaces
- Bridge domains, if you plan to use a general connector type
- Router configuration, if you plan to use route peering
- Redirect Policy

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Service Graph**.

Step 7 Choose the service graph template with managed or unmanaged node that you want to apply.

Check the **Managed** column of the **Nodes** tab of service graph template to determine if the node is managed or unmanaged.

Step 8 Click **Apply L4 L7 Service Graph Template**.

Step 9 On the **Apply L4 L7 Service Graph Template** screen, complete the following fields:

a) From the **Consumer EPG/External Network** drop-down list, choose either **EPG** or **External Network** and then do one of the following:

- Click **Select** and check the consumer EPG you want to use.
- Click **Select** and check the consumer external network you want to use.

b) From the **Provider EPG/External Network** drop-down list, choose either **EPG** or **External Network** and then do one of the following:

- Click **Select** and check the provider EPG you want to use.
- Click **Select** and check the provider external network you want to use.

c) From the **Create a New Contract/Choose an Existing Contract Subject** drop-down list, choose one of the following:

- **Create a New Contract** and then complete the contract name and filters fields for that contract.
- **Choose an Existing Contract Subject** and then click **Select** and check the contract subject that you want to use.

d) In the **Node One Consumer Connector** area, the status of the policy-based routing is displayed. When the policy-based routing status is true, the **Redirect Policy** drop-down list appears from which you can chose a redirect policy for the contract subject. You have to choose the device cluster, function profile, the consumer connector type and the consumer layer 3 destination virtual IP (VIP) address, and then complete the appropriate fields as per the chosen consumer connector type.

Note The function profile is enabled only when the function profile is not provided as input while creating the service graph template.

- **General**—Choose the **Consumer Bridge Domain** and **Consumer Cluster Interface**.
- **Route Peering**—Choose the **Router Configuration**, **Consumer Cluster Interface**, and **Consumer External Network**.

e) In the **Node One Provider Connector** area, the status of the policy-based routing is displayed. You can choose the provider connector type and the provider layer 3 destination VIP address, and then complete the appropriate fields as per the chosen provider connector type.

- **General**—Choose the **Provider Bridge Domain** and **Consumer Cluster Interface**.
- **Route Peering**—Choose the **Router Configuration**, **Provider Cluster Interface**, and **Provider External Network**.

f) If your service graph or service graph template is a Two Node type, complete the **Node Two Connector** fields for that node.

Step 10 Click **Submit**.

Service Graphs

Service graphs identify the set of network or service functions that are needed by an application. You can instantiate service graphs on the ACI fabric through Cisco UCS Director.

By using a service graph, you can install a service, such as an ASA firewall, once and deploy it multiple times in different logical topologies. Each time the graph is deployed, ACI takes care of changing the configuration on the firewall to enable the forwarding in the new logical topology.

A service graph represents the network using the following elements:

- Function node—A function node represents a function that is applied to network traffic, such as a transform (SSL termination, VPN gateway), filter (firewalls), or terminal (intrusion detection systems). A function within the service graph might require one or more parameters and have one or more connectors.
- Terminal node—A terminal node enables input and output from the service graph.
- Connector—A connector enables input and output from a node.
- Connection—A connection determines how traffic is forwarded through the network.

After you configure a service graph, the network services are automatically configured according to the service function requirements in the service graph. This does not require any change in the service device.

A service graph is represented as two or more tiers of an application with the appropriate service function inserted between them.

A service appliance (or device) performs a service function within the graph. One or more service appliances might be required to render the services required by a graph. One or more service functions can be performed by a single-service device.

Service graphs and service functions have the following characteristics:

- Traffic sent or received by an endpoint group can be filtered based on a policy, and a subset of the traffic can be redirected to different edges in the graph.

- Service graph edges are directional.
- Taps (hardware-based packet copy service) can be attached to different points in the service graph.
- Logical functions can be rendered on the appropriate (physical or virtual) device, based on the policy.
- The service graph supports splits and joins of edges, and it does not restrict the administrator to linear service chains.
- Traffic can be reclassified again in the network after a service appliance emits it.
- Logical service functions can be scaled up or down or can be deployed in a cluster mode or 1:1 active-standby high-availability mode, depending on the requirements.

For more information about the requirements of service graphs and their deployment, see the Cisco APIC Layer 4 to Layer 7 Services Deployment Guide.

Adding a Service Graph

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Service Graph**.
Step 7 Click **Add**.
Step 8 On the **Add Service Graph** screen, complete the fields, including the following:
a) Enter a unique name and description for the service graph.
b) Click **Select** and check the node that you want to use.

If the node you want to use is not in the list, click **Add** to create the node.

Step 9 Click **Submit**.

Adding a Filter to a Service Graph Node

A filter policy is a group of resolvable filter entries. Each filter entry is a combination of network traffic classification properties.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Service Graph**.

Step 7 Click the service graph you want to update and click **View Details**.

Step 8 Click the node where you want to add a filter and click **View Details**.

Step 9 Click **Connectors**.

Step 10 Click **Add**.

Step 11 On the **Add Filter to Service Graph Node** screen, complete the following fields:

a) From the **Connector Mode** drop-down list, choose **internal** or **external**.

b) Click **Select** and check the filter that you want to use.

Step 12 Click **Submit**.

Custom Quality of Service

Achieving the required Quality of Service (QoS) by effectively managing the priority of applications on the fabric is important when deploying an end-to-end solution. Thus, QoS is the set of techniques to manage data center fabric resources.

When QoS is used in ACI to classify packets, packets are classified using layer 2 Dot1P policy, layer 3 differentiated services code point (DSCP) policy, or contracts. DSCP/Dot1p Policy is configured and applied at the EPG level through custom QoS policy.

Adding a Custom QOS Policy

Create a custom QoS policy and then associate the policy with a cluster interface context.

Before you begin

Create the tenant, application, and EPGs that will consume the custom QoS policy.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Custom QOS Policy**.

Step 7 Click **Add**.

Step 8 On the **Create Custom QOS Policy** screen, complete the following fields:

a) Enter a unique name for the custom QOS policy.

b) Enter a short description for the custom QOS policy.

Step 9 Click **Submit**.

Adding a DSCP to a Priority Map

DSCP policy within the custom QoS policy is a set of rules; each rule gives mapping of a range of DSCP values to a DSCP target.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Custom QOS Policy**.

Step 7 Click the row with the custom QoS policy to which you want to add DSCP policy and click **View Details**.

Step 8 Click **DSCP to Priority Map**.

Step 9 Click **Add**.

Step 10 On the **Add DSCP to Priority Map** screen, complete the following:

a) (Optional) Choose the priority level of the DSCP policy in QoS as **Unspecified**, **Level3**, **Level2**, or **Level1**. By default, the unspecified is set as priority.

b) From the **DSCP Range From** and **DSCP Range To** drop-down lists, choose the starting and ending value for the DSCP range. To set the DSCP range from 0 to 63, choose **Enter customized value** from the drop-down lists and enter the actual value in the **Enter DSCP Range From** and **Enter DSCP Range To** fields.

c) (Optional) Choose a DSCP target to which the DSCP range has to be mapped. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 63 as the DSCP target in the **Enter DSCP Target** field.

d) (Optional) Choose a target class of service (CoS) from the drop-down list. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 8 as the target CoS in the **Enter Target Cos** field.

Step 11 Click **Submit**.

Adding a Dot1P Classifier

Dot1P policy within the custom QoS policy is a set of rules; each rule gives mapping of a range of Dot1P values to a DSCP target.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **Custom QOS Policy**.

Step 7 Click the row with the custom QoS policy to which you want to add Dot1P policy and click **View Details**.

Step 8 Click **Dot1P Classifier**.

Step 9 Click **Add**.

Step 10 On the **Add Dot1P Classifier** screen, complete the following:

a) Choose the priority level of the Dot1P policy in QoS as **Unspecified**, **Level3**, **Level2**, or **Level1**. By default, the unspecified is set as priority.

b) From the **Dot1P Range From** and **Dot1P Range To** drop-down lists, choose the starting and ending value for the Dot1P range.

c) Choose a DSCP target to which the Dot1P range has to be mapped. Choose **Enter customized value** from the drop-down list, to set any value between 0 and 64 as the DSCP target in the **Enter DSCP Target** field.
d) Choose a target cost of service (CoS) from the drop-down list.

Step 11 Click **Submit**.

Adding a Logical Device Context

The service graph uses a specific device based on a device selection policy, known as a logical device context.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Logical Device Context**.
Step 7 Click **Add**.
Step 8 On the **Add Tenant Logical Device Context** screen, complete the following fields:
a) Click **Select** and check the device cluster that you want to use.
b) Click **Select** and check the contract name that you want to use.
c) Click **Select** and check the graph name that you want to use.
d) Click **Select** and check the node name that you want to use.
e) Enter a unique name for the device context. The name should not exceed 64 characters.
Step 9 Click **Submit**.

Adding a Subnet to a Logical Device Context

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Logical Device Context**.
Step 7 Click the row with the logical device context to which you want to add a subnet and click **View Details**.
Step 8 Click **Cluster Interface Context**.
Step 9 Click the row with the cluster interface context to which you want to add a subnet and click **View Details**.
Step 10 Click **Subnets**.
Step 11 Click **Add**.
Step 12 On the **Add Subnet to Cluster Interface Context** screen, complete the following fields:
a) Enter a gateway IP address in the format: <valid IP address>/<valid prefix length>. For example, 10.10.10.1/24.

If this gateway is for Anycast, the netmask must be **/32** and check the **Subnet Control (No Default SVI Gateway)** check box to not to set the subnet as the default SVI gateway.

b) From the **Scope** drop-down list, choose **Private to VRF** or **Advertised Externally**. By default, the **Private to VRF** is set as the scope. The **Private to VRF** implies that the subnet can only be used in the tenant. The Advertised Externally option is used to advertise tenant subnets externally on the L3Out.
c) Check the **Shared Between VRFs** check box to define subnets under an endpoint group, with the shared option configured, to route leak to other tenants within the fabric.
d) Enter short description for the subnet.
e) Check the **Subnet Control (ND RA Prefix)** check box to apply control specific to ND RA prefix protocols to the subnet. By default, the check box is checked.
f) Check the **Subnet Control (No Default SVI Gateway)** check box to not to configure Pervasive SVI for the subnet. This setting is used to leak more specific prefix routes to other VRFs. If the **Subnet Control (No Default SVI Gateway)** check box is checked, you can use /32 netmask for the subnet, especially for Anycast services. By default, the check box is left unchecked.
g) Check the **Subnet Control (Querier IP)** check box to enable IGMP snooping on the subnet. By default, the check box is left unchecked.
h) Check the **Preferred** check box to set the subnet as the preferred subnet for the device context. By default, the check box is left unchecked.
i) Check the **Type Behind Subnet** check box to enable AnyCast MAC address. By default, the check box is left unchecked. The **MAC address** field appears only when the **Type Behind Subnet** check box is checked. Enter a MAC address in the format: xx:xx:xx:xx:xx:xx. For example, aa:11:bb:11:cc:11.

Step 13 Click **Submit**.

Adding a Cluster Interface Context

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Logical Device Context**.
Step 7 Click the row with the logical device context that you want to update and click **View Details**.
Step 8 Click **Cluster Interface Context**.
Step 9 Click **Add**.
Step 10 On the **Create Cluster Interface Context** screen, complete the following fields:
a) Click **Select** and check the logical device context to which you want to add an interface.
b) Enter the connector name. By default, **any** is set as the connector name.
c) Click **Select** and check the logical interface name that you want to add to the logical device context.
d) Click **Select** and check the bridge domain name that you want to add to the logical device context.
e) Click **Select** and check the Layer 3 network that you want to add to the logical device context.
f) From the **L3 Destination (VIP)** drop-down list, choose an appropriate option.

- **Unspecified**—This is the default option. Choose this option to not to specify rule for Layer 3 destination.

- **True**—If the PBR policy is not configured on a specific service node, the node connector is treated as an L3 Destination and will continue to be in the new Cisco APIC version.
- **False**—Set false to enable user to choose the PBR policy for logical interface.

g) The **L4-L7 Policy Based Redirect** field appears only when **False** is selected as L3 Destination (VIP). Click **Select** and check the L4-L7 PBR that need to be associated to the interface context.
h) Click **Select** and check the custom QoS policy that need to be associated to the interface context.
i) Choose **True** from the **Permit Logging** drop-down list to enable permit logging for cluster interface context.

Step 11 Click **Submit**.

Adding a Virtual IP Address to a Cluster Interface Context

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **Logical Device Context**.
Step 7 Click the row with the logical device context that you want to update and click **View Details**.
Step 8 Click **Cluster Interface Context**.
Step 9 Click the row with the cluster interface context that you want to update and click **View Details**.
Step 10 Click **Virtual IP Address**.
Step 11 Click **Add**.
Step 12 On the **Add Virtual IP Address to Logical Interface** screen, enter the IPv4 address for the logical interface.
Step 13 Click **Submit**.

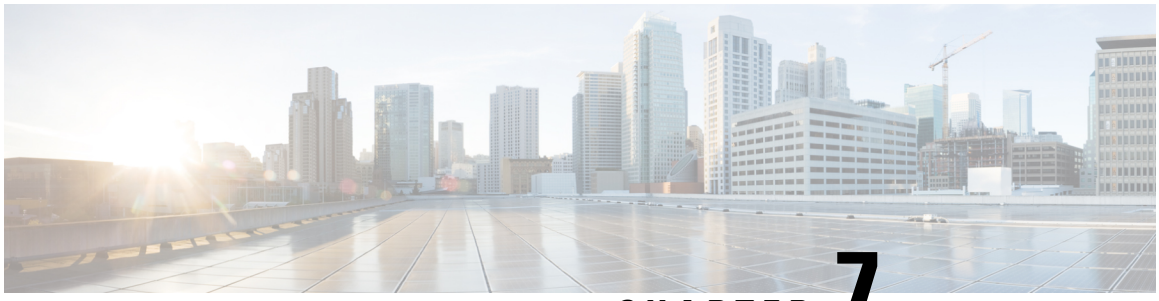

CHAPTER 7

Configuring Policy Based Redirect

Policy-Based Redirect

Cisco Application Centric Infrastructure (ACI) policy-based redirect (PBR) enables provisioning service appliances, such as firewalls or load balancers, as managed or unmanaged nodes without requiring a Layer 4 to Layer 7 package. Typical use cases include provisioning service appliances that can be pooled, tailored to application profiles, scaled easily, and have reduced exposure to service outages. PBR simplifies the deployment of service appliances by enabling the provisioning consumer and provider endpoint groups to be all in the same virtual redirect and forwarding (VRF) instance.

PBR deployment consists of configuring a route redirect policy and a cluster redirect policy, and creating a service graph template that uses the route and cluster redirect policies. After the service graph template is deployed, use the service appliance by enabling endpoint groups to consume the service graph provider endpoint group. This can be further simplified and automated by using vzAny. While performance requirements may dictate provisioning dedicated service appliances, virtual service appliances can also be deployed easily using PBR.

Creating Layer 4-Layer 7 Policy Based Redirect

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **L4-L7 Policy Based Redirect**.

Step 7 Click **Add**.

Step 8 On the **Create Policy Based Redirect** screen, complete the following fields:

- Enter a unique name and description for the Policy Based Redirect.

- Check the **Enable Pod ID Aware Redirection** check box to enable pod ID aware redirection and associate the pod IDs with the preferred PBR nodes to program redirect destinations in the leaf switches located in the specific pods.
- Choose one of the following hashing algorithms:
  - dip—Destination IP address
  - sip—Source IP address
  - sip-dip-prototype—Source IP address, Destination IP address and Protocol Type (also called Symmetric) based algorithm
- Check the **Resilient Hashing Enabled** check box to enable resilient hashing for mapping traffic flows to physical nodes and for avoiding the rehashing of any traffic other than the flows from the failed node.
- Check the **Anycast Endpoint** check box to enable anycast endpoint.
- Click **Select** and check the IP SLA monitoring policy that you want to use for PBR tracking.
- The **Threshold Enable** check box appears when you choose an IP SLA monitoring policy. Check this check box to enable threshold when you want to disable the redirect destination group completely and prevent any redirection. When there is no redirection, the traffic is directly sent between the consumer and the provider. The following threshold settings are available:
  - **Min Threshold Percent (Percentage)** field—Enter the minimum threshold percentage. If the traffic goes below the minimum percentage, the packet is permitted instead of being redirected. The default value is 0. The allowed threshold range is from 0 to 100.
  - **Max Threshold Percent (Percentage)** field—Enter the maximum threshold percentage. When the minimum threshold is reached, to revert to the operational state, the maximum threshold percentage must be reached first. The default value is 0. The allowed threshold range is from 0 to 100.
  - **Threshold Down Action** drop-down list—Choose **permit action** or **deny action** from the drop-down list to apply the threshold settings on traffic.

Step 9 Click **Submit**.

Creating Layer 4 - Layer 7 Redirect Health Group

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4 L7 Redirect Health Group**.
Step 7 Click **Add**.
Step 8 On the **Create L4-L7 Redirect Health Group** screen, enter a unique name and description for L4-L7 Redirect Health Group.

Step 9 Click **Submit**.

When a redirect health group is no longer consumed by the PBR, you can delete the redirect health group. To delete the redirect health group, click the row with the redirect health group on the **L4 L7 Redirect Health Group** screen and click **Delete**.

Creating a Destination of Redirect Traffic

Before you begin

The redirect health group that needs to be associated with the redirect traffic is created.

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **L4-L7 Policy Based Redirect**.
Step 7 Click the row with the L4-L7 policy-based redirect record that you want to update and click **View Details**.
Step 8 Click **Destination of Redirect Traffic**.
Step 9 Click **Add**.
Step 10 On the **Add Destination of Redirected Traffic** screen, complete the following fields:

- Enter the IP address for the Layer 4 to Layer 7 device. The IP address must be in the same subnet as the IP address that you have given to the bridge domain.
- Enter a short description for the destination of redirected traffic.
- Enter the MAC address for the Layer 4 to Layer 7 device. You should use a MAC address that is valid upon failover of the Layer 4 to Layer 7 device.
- Enter the secondary IP address for the Layer 4 to Layer 7 device.
- Enter the pod identification value. By default, 1 is set as the pod ID. The valid pod ID range is from 1 to 255.
- Click **Select** and check the check box for the redirect health group that you want to associate to an existing health group.

Step 11 Click **Submit**.

Creating an IP SLA Monitoring Policy

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **IP SLA Monitoring Policy**.

Step 7 Click **Add**.

Step 8 On the **Create IP SLA Monitoring Policy** screen, complete the following fields:

a) Enter a unique name and description for the IP SLA Monitoring Policy.

b) In the **SLA Frequency** field, enter the interval probe time to track a packet. The allowed SLA frequency range is 1 to 65535 seconds. The default value is 60 seconds.

c) Choose **icmp** or **tcp** as the SLA type. If you choose **tcp**, then enter the SLA port number in the **SLA Port** field.

Step 9 Click **Submit**.

vzAny

The vzAny managed object provides a convenient way of associating all endpoint groups (EPGs) in a Virtual Routing and Forwarding (VRF) instance to one or more contracts, instead of creating a separate contract relation for each EPG.

To view vzAny, choose **Physical** > **Network** > **Multi-Domain Managers** > **APIC Accounts** > **Tenant(s)** > **VRF** > **vzAny**.

Creating a vzAny Provided Contract

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with the VRF to which you want to add vzAny Provided Contract, and click **View Details**.

Step 8 Click **vzAny Provided Contract**.

Step 9 Click **Add**.

Step 10 On the **Add vzAny Provided Contract to VRF** screen, complete the following fields:

a) Click **Select** and choose the contract that you want to use for the vzAny.

b) Choose one of the following as the priority level of the quality of service (QoS):

- **Unspecified**—Default value.
- **Level3**—Class 3 Differentiated Services Code Point (DSCP) value.
- **Level2**—Class 2 DSCP value.
- **Level1**—Class 1 DSCP value.

c) Choose one of the following as the match criteria for the provided contract:

- **All**—Only matches when both endpoint groups have all labels, excluding blank labels.
- **AtleastOne**—At least 1 label matches on Provider and Consumer endpoint groups. Blank labels are considered a match.
- **AtmostOne**—Matches only when all labels on the endpoint groups are exactly the same. Blank labels are considered a match.
- **None**—None of the subject labels match.

Step 11 Click **Submit**.

Creating a vzAny Consumed Contract

Step 1 Choose **Physical** > **Network**.
Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.
Step 3 Click the row with the APIC account and click **View Details**.
Step 4 Click **Tenant(s)**.
Step 5 Click the row with the tenant that you want to update and click **View Details**.
Step 6 Click **VRF**.
Step 7 Click the row with the VRF to which you want to add vzAny Provided Contract, and click **View Details**.
Step 8 Click **vzAny Consumed Contract**.
Step 9 Click **Add**.
Step 10 On the **Add vzAny Consumed Contract to VRF** screen, complete the following fields:

a) Click **Select** and choose the contract that you want to use for the vzAny.
b) Choose one of the following as the priority level of the quality of service (QoS):

- **Unspecified**—Default value.
- **Level3**—Class 3 Differentiated Services Code Point (DSCP) value.
- **Level2**—Class 2 DSCP value.
- **Level1**—Class 1 DSCP value.

Step 11 Click **Submit**.

Creating a vzAny Contract Interface

A contract interface is used to associate an EPG from the destination tenant with the imported contract.

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click **Tenant(s)**.

Step 4 Click the row with the tenant that you want to update and click **View Details**.

Step 5 Click **VRF**.

Step 6 Click the row with the VRF to which you want to add vzAny contract interface and click **View Details**.

Step 7 Click **vzAny Contract Interface**.

Step 8 Click **Add**.

Step 9 On the **Add Contract Interface** screen, complete the following fields:

a) Click **Select** and check the contract interface that you want to use.

b) Choose one of the following as the priority level of the service contract:

- **Unspecified**—Default value.
- **Level3**—Class 3 Differentiated Services Code Point (DSCP) value.
- **Level2**—Class 2 DSCP value.
- **Level1**—Class 1 DSCP value.

Step 10 Click **Submit**.

Labels

Labels are managed objects with only one property: a name. Labels enable classifying which objects can and cannot communicate with one another. Label matching is done first. If the labels do not match, no other contract or filter information is processed. The label match attribute can be one of these values: at least one (the default), all, none, or exactly one.

Labels determine which EPG consumers and EPG providers can communicate with one another. Label matching determines which subjects of a contract are used with a given EPG provider or EPG consumer of that contract.

The two types of labels are as follows:

- Subject labels are applied to EPGs. Subject label matching enables EPGs to choose a subset of the subjects in a contract.
- Provider or consumer labels are applied to EPGs. Provider or consumer label matching enables consumer EPGs to choose their provider EPGs and vice versa.

Creating a vzAny EPG Consumed Any Labels

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with the VRF to which you want to add vzAny EPG consumed any label and click **View Details**.

Step 8 Click **vzAny EPG Consumed Any Labels**.

Step 9 Click **Add**.

Step 10 On the **Add Consumed Any EPG Label to vzAny VRF** screen, complete the following fields:

a) Enter a unique name for the consumed any label.

b) From the **Label Tag** drop-down list, choose a color for the label.

Step 11 Click **Submit**.

Creating a vzAny EPG Provided Any Labels

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with the VRF to which you want to add vzAny EPG provided any label and click **View Details**.

Step 8 Click **vzAny EPG Provided Any Labels**.

Step 9 Click **Add**.

Step 10 On the **Add Provided Any EPG Label to vzAny VRF** screen, complete the following fields:

a) Enter a unique name for the provided any label.

b) From the **Label Tag** drop-down list, choose a color for the label.

c) Check the **Complement** check box to enable the complement for the provided any label. By default, the complement is disabled.

Step 11 Click **Submit**.

Creating APIC vzAny Provided Subject Label to VRF

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with the VRF to which you want to add vzAny provided subject label and click **View Details**.

Step 8 Click **vzAny Subject Label Provided**.

Step 9 Click **Add**.

Step 10 On the **Add APIC vzAny Provided Subject Label to VRF** screen, complete the following fields:

a) Enter a unique name for the provided subject label.
b) From the **Label Tag** drop-down list, choose a color for the label.
c) Check the **Complement** check box to enable the complement for the provided subject label. By default, the complement is disabled.

Step 11 Click **Submit**.

Creating APIC vzAny Consumed Subject Label to VRF

Step 1 Choose **Physical** > **Network**.

Step 2 On the **Network** page, choose the account under **Multi-Domain Managers**.

Step 3 Click the row with the APIC account and click **View Details**.

Step 4 Click **Tenant(s)**.

Step 5 Click the row with the tenant that you want to update and click **View Details**.

Step 6 Click **VRF**.

Step 7 Click the row with the VRF to which you want to add vzAny consumed subject label and click **View Details**.

Step 8 Click **vzAny Subject Label Consumed**.

Step 9 Click **Add**.

Step 10 On the **Add APIC vzAny Consumed Subject Label to VRF** screen, complete the following fields:

a) Enter a unique name for the consumed subject label.
b) From the **Label Tag** drop-down list, choose a color for the label.
c) Check the **Complement** check box to enable the complement for the consumed subject label. By default, the complement is disabled.

Step 11 Click **Submit**.