Using the Advanced GUI

This chapter contains the following sections:

Toggling Between Basic and Advanced GUI Modes

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

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

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

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

1 In the GUI, click the **welcome, <login_name>** drop-down list and choose Toggle GUI Mode.

2 In the **Warning** dialog box, click Yes for

3 Wait for the application to complete loading and display the GUI in the changed mode.

About Getting Started with APIC Examples

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

Switch Discovery with the APIC

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

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

Switch Registration with the APIC Cluster

Note Before you begin registering a switch, make sure that all switches in the fabric are physically connected and booted in the desired configuration. For information about the installation of the chassis, see http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-installation-guides-list.html.

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

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

Registering the Unregistered Switches Using the GUI

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

Before You Begin

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

Procedure

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

Step 2 In the **Navigation** pane, click **Fabric Membership**.
In the **Work** pane, in the **Fabric Membership** table, a single leaf switch is displayed with an ID of 0. It is the leaf switch that is connected to apic1.

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

a) In the **ID** field, add the appropriate ID (leaf1 is ID 101, and leaf 2 is ID 102).
The ID must be a number that is greater than 100 because the first 100 IDs are for APIC appliance nodes.

b) In the **Switch Name** field, add the name of the switch, and click **Update**.

Note After an ID is assigned, it cannot be updated. The switch name can be updated by double-clicking the name and updating the **Switch Name** field.

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

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

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

a) In the **ID** field, add the appropriate ID (spine1 is ID 203 and spine 2 is ID 204).

Note It is recommended that leaf nodes and spine nodes be numbered differently. For example, number spines in the 200 range and number leafs in the 100 range.

b) In the **Switch Name** field, add the name of the switch, and click **Update**.

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

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

a) Double-click the switch, enter an **ID** and a **Name**, and click **Update**.

b) Repeat for the next switch in the list.

Switch Discovery Validation and Switch Management from the APIC

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

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

Validating the Registered Switches Using the GUI

Procedure

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

Step 2 In the **Navigation** pane, expand **Fabric Membership**.
The switches in the fabric are displayed with their node IDs. In the **Work** pane, all the registered switches are displayed with the IP addresses that are assigned to them.

Validating the Fabric Topology

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

Validating the Fabric Topology Using the GUI

Procedure

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

Step 2 In the **Navigation** pane, choose the pod that you want to view.

Step 3 In the **Work** pane, click the **TOPOLOGY** tab.
The displayed diagram shows all attached switches, APIC instances, and links.

Step 4 (Optional) To view the port-level connectivity of a leaf switch or spine switch, double-click its icon in the topology diagram.
To return to the topology diagram, in the upper left corner of the **Work** pane, click the **Previous View** icon.

Step 5 (Optional) To refresh the topology diagram, in the upper left corner of the **Work** pane, click the **Refresh** icon.

Unmanaged Switch Connectivity in VM Management

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

Note The ACI simulator only supports LLDP. Cisco Discovery Protocol (CDP) is not supported.

Figure 1: Unmanaged Layer 2 Switches in the APIC Fabric Inventory

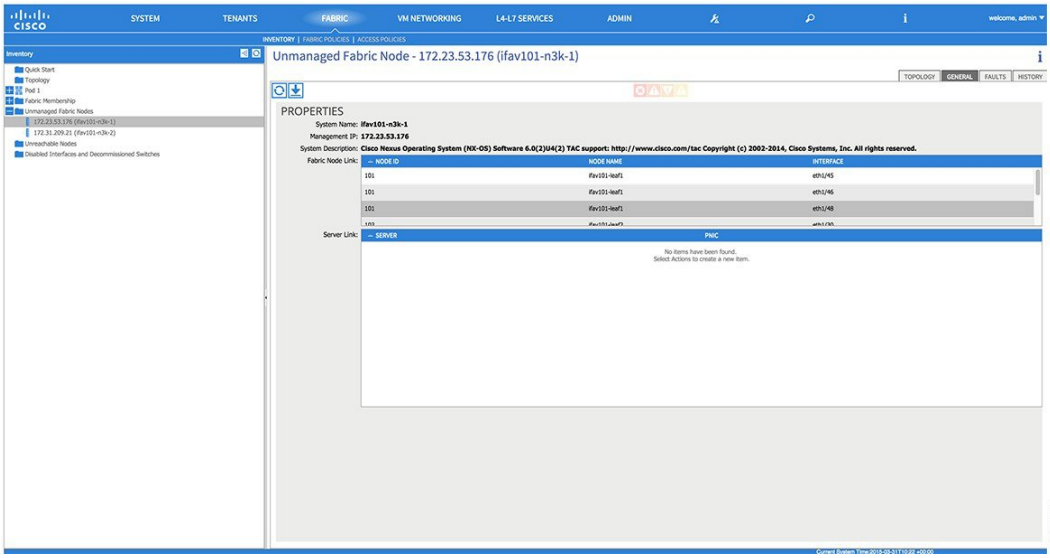

Configuring Network Time Protocol

Time Synchronization and NTP

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

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

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

In-Band Management NTP

Note

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
- See the Adding Management Access section in this guide for information about in-band management access.

- In-Band Management NTP—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication..

NTP over IPv6

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

Configuring NTP Using the Advanced GUI

Procedure

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

Step 2 In the **Navigation** pane, choose **Pod Policies** > **Policies**.

Step 3 In the **Work** pane, choose **Actions** > **Create Date and Time Policy**.

Step 4 In the **Create Date and Time Policy** dialog box, perform the following actions:

a) Enter a name for the policy to distinguish between the different NTP configurations in your environment. Click **Next**.

b) Click the + sign to specify the NTP server information (provider) to be used.

c) In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.

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

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

Step 5 In the **Navigation** pane, choose **Pod Policies** > **Policy Groups**.

Step 6 In the **Work** pane, choose **Actions** > **Create Pod Policy Group**.

Step 7 In the **Create Pod Policy Group** dialog box, perform the following actions:

a) Enter a name for the policy group.

b) In the **Date Time Policy** field, from the drop down list, choose the NTP policy that you created earlier. Click **Submit**.
The pod policy group is created. Alternatively, you can use the default pod policy group.

Step 8 In the **Navigation** pane, choose **Pod Policies** > **Profiles**.

Step 9 In the **Work** pane, double-click the desired pod selector name.

Step 10 In the Properties area, from the **Fabric Policy Group** drop down list, choose the pod policy group you created. Click **Submit**.

Verifying NTP Operation Using the GUI

Procedure

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

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

Step 3 In the **Work** pane, verify the details of the server.

Verifying NTP Policy Deployed to Each Node Using the CLI

Procedure

Step 1 SSH to an APIC in the fabric.

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

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

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

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

Step 4 View the NTP peer status.

Example:
```
leaf-1# show ntp peer-status
```
A reachable NTP server has its IP address prefixed by an asterisk (*), and the delay is a non-zero value.

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

Creating User Accounts

Configuring a Local User

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

Configuring a Remote User

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

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

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

Configuring a Local User Using the GUI

Before You Begin

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

Procedure

Step 1 On the menu bar, choose **ADMIN** > **AAA**.

Step 2 In the **Navigation** pane, click **AAA Authentication**.

Step 3 In the **Work** pane, verify that in the default **Authentication** field, the **Realm** field displays as Local.

Step 4 In the **Navigation** pane, expand **Security Management** > **Local Users**.
The admin user is present by default.

Step 5 In the **Navigation** pane, right-click **Create Local User**.

Step 6 In the **Security** dialog box, choose the desired security domain for the user, and click **Next**.

Step 7 In the **Roles** dialog box, click the radio buttons to choose the roles for your user, and click **Next**.
You can provide read-only or read/write privileges.

Step 8 In the **User Identity** dialog box, perform the following actions:

a) In the **Login ID** field, add an ID.

b) In the **Password** field, enter the password.
At the time a user sets their password, the APIC validates it against the following criteria:

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

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

d) Click **Finish**.

Step 9 In the **Navigation** pane, click the name of the user that you created. In the **Work** pane, expand the + sign next to your user in the **Security Domains** area.
The access privileges for your user are displayed.

AV Pair on the External Authentication Server

You can add a Cisco attribute/value (AV) pair to the existing user record to propagate the user privileges to the APIC controller. The Cisco AV pair is a single string that you use to specify the Role-Based Access Control (RBAC) roles and privileges for an APIC user. An example configuration for an open RADIUS server (/etc/raddb/users) is as follows:

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

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

Procedure

Step 1 On the menu bar, click **ADMIN** > **AAA**.

Step 2 In the **Navigation** pane, click **AAA Authentication**.

Step 3 In the **Work** pane, in the **Properties** area, from the **Remote user login policy** drop-down list, choose **Assign Default Role**.
The default value is **No Login**. The **Assign Default Role** option assigns the minimal read-only privileges to users that have missing or bad Cisco AV Pairs. Bad AV Pairs are those AV Pairs that fail the parsing rules.

Best Practice for Assigning AV Pairs

As best practice, Cisco recommends that you assign unique UNIX user ids in the range 16000-23999 for the AV Pairs that are assigned to users when in bash shell (using SSH, Telnet or Serial/KVM consoles). If a situation arises when the Cisco AV Pair does not provide a UNIX user id, the user is assigned a user id of 23999 or similar number from the range that also enables the user's home directories, files, and processes accessible to remote users with a UNIX ID of 23999.

Configuring an AV Pair on the External Authentication Server

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

Procedure

Configure an AV pair on the external authentication server.
The Cisco AV pair definition is as follows (Cisco supports AV pairs with and without UNIX user IDs specified):

Example:

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

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

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

The following is an example:

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

Configuring a Remote User Using the GUI

Before You Begin

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

Procedure

Step 1 On the menu bar, choose **ADMIN** > **AAA**. In the **Navigation** pane, expand **RADIUS Management**.

Step 2 Right-click **RADIUS Providers**, and click **Create RADIUS Provider**.

Step 3 In the **Create RADIUS Provider** dialog box, and perform the following actions:

a) In the **Host Name (or IP Address)** field, add the hostname.
b) In the **Authorization Port** field, add the port number required for authorization.
This number depends on the RADIUS server configured.
c) Click the required **Authorization Protocol** radio button.
d) In the **Key** and **Confirm Key** fields, enter the preshared key.
This key is the same information that is shared with the server key configured on the RADIUS server.

Step 4 In the **Navigation** pane, under **RADIUS Providers**, click the RADIUS provider that you created.
Details about the configurations for the RADIUS provider are displayed in the **Work** pane.

Step 5 In the **Navigation** pane, right-click **RADIUS Provider Groups**, and click **Create RADIUS Provider Group**.

Step 6 In the **Create RADIUS Provider Group** dialog box, perform the following actions:

a) In the **Name** field, enter a name.
b) Expand the **Providers** field, and from the **Name** field drop-down list, choose the provider created earlier.
c) In the **Priority** field, assign a priority. Click **Update**, and click **Submit**.
The radius provider group is created.

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

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

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

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

Adding Management Access

In-band management access—You can configure in-band management connectivity to the APIC and the ACI fabric. You first configure the VLANs that will be used by APIC when the APIC is communicating with the leaf switches, and then you configure the VLANs that the VMM servers will use to communicate with the leaf switches.

Note Do not configure the APIC selector (the set of leaf ports to which the APIC is connected) when configuring the simulator with in-band management access.

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

IPv4/IPv6 Addresses and In-Band Policies

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

IPv4/IPv6 Addresses in Out-of-Band Policies

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

Configuring Management Access

Configuring In-Band Management Access Using the GUI

When using the APIC simulator, the IP addresses are automatically assigned. If you configure any IP addresses in the following steps, the IP addresses you configure will not be effective. This is because the IP addresses are pre-configured and a customized configuration is not supported.

Procedure

Step 1 On the menu bar, choose **FABRIC** > **Access Policies**. In the **Navigation** pane, expand **Interface Policies**.

Step 2 In the **Navigation** pane, right-click **Switch Policies** and choose **Configure Interface, PC and VPC**.

Step 3 In the **Configure Interface, PC, and VPC** dialog box, to configure the ports connected to VMM servers, perform the following actions:

a) Click the large + icon next to the switch diagram to create a new profile and configure VLANs for the APIC.
b) In the **Switches** field, from drop-down list, check the check boxes for the switches to which the VMM servers are connected. (leaf1).
c) In the **Switch Profile Name** field, enter a name for the profile (vmmConnectedLeaves).
d) Click the + icon to configure the ports.
e) Verify that in the **Interface Type** area, the **Individual** radio button is selected.
f) In the **Interfaces** field, enter the ports to which VMM servers are connected (1/40).
g) In the **Interface Selector Name** field, enter the name of the port profile (vmmConnectedPorts).
h) In the **Interface Policy Group** field, click the **Create One** radio button.
i) In the **Attached Device Type** field, choose the appropriate device type to configure the domain (Bare Metal).
j) In the **Domain** field, from the drop-down list click the **Choose One** radio button
k) From the **Physical Domain** drop-down list, choose the domain created earlier.
l) In the **Domain Name** field, enter the domain name. (**inband**)
m) Click **Save**, and click **Save** again.

Step 4 In the **Configure Interface, PC, and VPC** dialog box, click **Submit**.

Step 5 On the menu bar, click **TENANTS** > **mgmt**. In the **Navigation** pane, expand **Tenant mgmt** > **Networking** > **Bridge Domains** to configure the bridge domain on the in-band connection.

Step 6 Right-click the in-band bridge domain (inb). Right-click **Subnets**. Click **Create Subnet** and perform the following actions to configure the in-band gateway:

a) In the **Create Subnet** dialog box, in the **Gateway IP** field, enter the in-band management gateway IP address and mask.
b) Click **Submit**.

Step 7 In the **Navigation** pane, expand **Tenant mgmt** > **Node Management EPGs**. Click the default in-band EPG (In-Band EPG - default), and in the **Work** pane, perform the following actions to set the VLAN on the in-band EPG used to communicate with the APIC:

a) In the **Work** pane, in the **In-Band EPG default** area, verify that the default is displayed.
b) In the **Encap** field, enter the VLAN (vlan-10).
c) Expand **Provided Contracts**. In the **Name** field, from the drop-down list, choose the default contract to enable EPG to provide the default contract that will be consumed by the EPGs on which the VMM servers are located.
d) Click **Update**, and click **Submit**.

Step 8 In the **Navigation** pane, right-click **Node Management Addresses**. Click **Create Node Management Addresses**, and perform the following actions to configure the IP addresses for the leaf and spine switches in the fabric:

a) In the **Create Node Management Addresses** dialog box, in the **Policy Name** field, enter the policy name (switchInb).
b) In the **Nodes** field, **Select** column, check the check boxes next to the nodes that will be part of this fabric (leaf1, leaf2, spine1, spine2).
c) In the **Config** field, click the **In-Band Addresses** checkbox.
d) In the **Node Range** fields, enter the range.
e) In the **In-Band IP Addresses** area, in the **In-Band Management EPG** field, from the drop-down list, choose default. The default in-band management EPG is now associated.

f) In the **In-Band IP Addresses** and **Gateway** fields, enter the IPv4 or IPv6 addresses as desired. This range of IP addresses must be different from the range of addresses assigned to the APICs.

g) Click **Submit**. In the **Confirm** dialog box, click **Yes**. The IP addresses for the leaf and spine switches are now configured.

Step 9 In the **Navigation** pane, under **Node Management Addresses**, click the switches policy name (switchInb). In the **Work** pane, the IP addresses that are assigned to switches and the gateway addresses they are using are displayed.

IPv6 Table Modifications to Mirror the Existing IP Tables Functionality

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

Existing IP Tables

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

Modifications to IP Tables

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

   ```
   -A INPUT -s <OOB Address Ipv4/Ipv6> -j apic-default
   ```

5. When a new rule is added, it presents in the IP tables-new file and not in the IP tables file, and it signifies that there is some error in the IP tables-new file. Only if the restoration is successful, the file is saved and new rules are seen in the IP tables file.

Note

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

Management Connectivity Modes

Establish connection to external entities using the out-of-band or in-band network depending upon whether you have configured out-of-band and/or in-band management connectivity. The following two modes are available to establish connectivity to external entities such as the vCenter server:

- Layer 2 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 2.
- Layer 3 management connectivity—Use this mode when the external entities are attached to the leaf node using Layer 3 through a router. The leaf is connected to a router through which external entities can be reached.

Note

- The inband IP address range must be separate and distinct from the IP address range used on the Layer 3 connection from the leaf node to outside the fabric.
- The Layer 3 inband management design does not provide inband management access to the spine fabric nodes in the topology.

The following diagram displays the two modes available to establish connectivity.

Figure 2: Layer 2 and Layer 3 Management Connectivity Examples

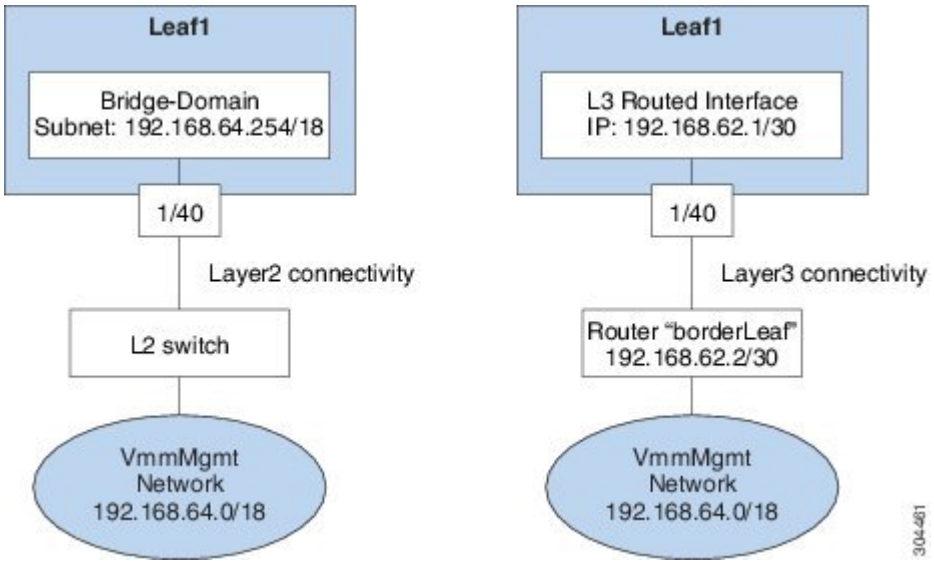

Configuring Layer 2 Management Connectivity Using the Advanced GUI

Note

Before You Begin

Before you create a vCenter domain profile, you must establish connectivity to establish an external network using in-band management network.

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

Step 1 On the menu bar, choose **Tenants** > **mgmt**.

Step 2 In the **Navigation** pane, expand **Tenant mgmt** > **Networking**, right-click **Bridge Domains**, and click **Create Bridge Domain**.

Step 3 In the **Create Bridge Domain** dialog box, perform the following actions:

a) In the **Name** field, enter a bridge domain name.

b) In the **VRF** field, from the drop-down list, choose the network (mgmt/inb). Click **Next**.

c) Click the **L3 Configuration** tab, and in the **Subnets** field, click the + icon to add a subnet. Add the Gateway IP address as required.

d) In the **Create Bridge Domain** dialog box, click **Next** and then click **Submit**.

The bridge domain created.

Step 4 In the **Navigation** pane, expand **Tenant mgmt** > **Application Profiles**.

Step 5 Right-click **Application Profiles** and click **Create Application Profile**.

Step 6 In the **Create Application Profile** dialog box, perform the following actions:

a) In the **EPGs** field, click the + icon to add an EPG, and in the **Name** field, enter a name.
b) From the **BD** drop-down list, choose the appropriate BD.
c) From the **Domain** field drop-down list, choose the appropriate domain.
d) In the **Static Path** field, (enter the appropriate values similar to the following example, 101/1/40).
e) In the **Static Path VLAN** field, enter the appropriate VLAN (enter the appropriate value similar to the following example vlan-11).
f) In the **Consumed Contract** field, from the drop-down list, choose the appropriate value. Click **Update** and **Submit**.

In the **Navigation** pane, under **Networking** a bridge domain is created, and under **Application Profiles**, an application profile and an application EPG are created. The layer 2 management connectivity is now configured.

Configuring Layer 3 Management Connectivity Using the Advanced GUI

Note

- The name vmm is used as an example string in this task.

Before You Begin

Before you create a VMM domain profile, you must establish connectivity to an external network using the inband-management network.

Make sure that the IP address range configured as part of management connectivity policy does not overlap with the infrastructure IP address range used by the ACI fabric.

Procedure

Step 1 On the menu bar, choose **TENANTS** > **mgmt**.

Step 2 In the **Navigation** pane, perform the following actions:

a) Expand **Tenant mgmt** > **Networking** > **External Routed Networks**.
b) Right-click **Create Routed Outside**.

Step 3 In the **Create Routed Outside** dialog box, perform the following actions:

a) In the **Name** field, enter the name of the Layer 3 routed outside policy (vmm).
This name can be up to 64 alphanumeric characters. You cannot change the name after the object is saved.
b) From the **VRF** drop-down list, choose the in-band default network (mgmt/inb).
Note You must choose the default in-band network.

Step 4 Expand the **Nodes and Interfaces Protocol Profiles** area. In the **Create Node Profile** dialog box, perform the following actions:

a) In the **Name** field, enter a name. (borderLeaf)

b) Expand **Nodes** to display the **Select Node** dialog box. In the **Node ID** field, choose a leaf switch from the drop-down list (leaf1).
c) In the **Router ID** field, enter the router ID.
d) Expand **Static Routes**.
e) In the **Create Static Route** dialog box, in the **Prefix** field, enter the subnet prefix for the static route of the external management system (for example, the VMware vCenter, the syslog server, or the AAA server) with which you are trying to communicate.
f) Expand **Next Hop Addresses** and in the **Next Hop IP** field, enter the IP address of the router that is connected to the leaf switch. In the **Preference** field, choose a preference. Click **Update**.
g) Click **OK**. In the **Select Node** dialog box, click **OK**.

Step 5 Expand **Interface Profiles**. In the **Create Interface Profile** dialog box, perform the following actions:
a) In the **Name** field, enter a name. (portProfile1)
b) Expand **Routed Interfaces**. In the **Select Routed Interface** area, in the **Path** field, from the drop-down list, choose the path that associates with leaf1.
c) In the **IPv4 Primary/IPv6 Preferred Address** field, enter the IP address and subnet mask for the routed interface on the leaf. Click **OK**
d) In the **Create Interface Profile** dialog box, click **OK**. In the **Create Node Profile** dialog box, click **OK**.

Step 6 In the **Create Routed Outside** dialog box, click **Next**, and expand **External EPG Networks**.

Step 7 In the **Create External Network** dialog box, perform the following actions:
a) In the **Name** field, enter a name (vmmMgmt).
b) Expand the + icon for **Subnet**.
c) In the **Create Subnet** dialog box, in the **IP address** field, enter the subnet address.
d) Click **OK** two times, and click **Finish**.

The L3 management connectivity is configured.

Validating Management Connectivity

This validation process applies to both Layer 2 and Layer 3 modes and can be used to verify connectivity that is established by using the APIC GUI, REST API, or CLI.

After completing the steps to establish management connectivity, log in to the APIC console. Ping to the IP address of the vCenter server that is reachable (for example, 192.168.81.2) and verify that the ping works. This action indicates that the policies have been successfully applied.

Configuring a VMM Domain

Configuring Virtual Machine Networking Policies

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

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

About the VM Manager

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

The following are details of some VM manager terms:

- A VM controller is an external virtual machine management entity such as VMware vCenter, and the VMware vShield. The APIC communicates with the controller to publish network policies that are applied to virtual workloads. A VM controller administrator provides an APIC administrator with a VM controller authentication credential; multiple controllers of the same type can use the same credential.
- Credentials represent the authentication credentials to communicate with VM controllers. Multiple controllers can use the same credentials.
- A virtual machine mobility domain (vCenter mobility domain) is a grouping of VM controllers with similar networking policy requirements. This mandatory container holds one or more VM controllers with policies such as for a VLAN pool, server to network MTU policy, or server to network access LACP policy. When an endpoint group gets associated with a vCenter domain, network policies get pushed to all the VM controllers in the vCenter domain.
- A pool represents a range of traffic encapsulation identifiers (for example, VLAN IDs, VNIDs, and multicast addresses). A pool is a shared resource and can be consumed by multiple domains such as VMM and Layer 4 to Layer 7 services. A leaf switch does not support overlapping VLAN pools. You must not associate different overlapping VLAN pools with the VMM domain. The two types of VLAN-based pools are as follows:
  - Dynamic pools—Managed internally by the APIC to allocate VLANs for endpoint groups (EPGs). A vCenter Domain can associate only to a dynamic pool.
  - Static pools—The EPG has a relation to the domain, and the domain has a relation to the pool. The pool contains a range of encapsulated VLANs and VXLANs. For static EPG deployment, the user defines the interface and the encapsulation. The encapsulation must be within the range of a pool that is associated with a domain with which the EPG is associated.
- For a VMware vCenter to be deployed, it must operate in VLAN mode or VXLAN mode. A VMM domain must be associated with a VLAN pool and a vShield must be associated with the vCenter.

About Attachable Entity Profile

Attach Entity Profiles

The ACI fabric provides multiple **attachment points** that connect through leaf ports to various **external entities** such as baremetal servers, hypervisors, Layer 2 switches (for example, the Cisco UCS fabric interconnect), and Layer 3 routers (for example Cisco Nexus 7000 Series switches). These attachment points can be physical ports, port channels, or a virtual port channel (vPC) on the leaf switches.

An **attachable entity profile** (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies, for example, Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), maximum transmission unit (MTU), and Link Aggregation Control Protocol (LACP).

A VM manager (VMM) domain automatically derives the physical interfaces policies from the interface policy groups that are associated with an AEP.

- An override policy at AEP can be used to specify a different physical interface policy for a VMM domain. This policy is useful in scenarios where a hypervisor is connected to the leaf switch through an intermediate Layer 2 node, and a different policy is desired at the leaf switch and hypervisor physical ports. For example, you can configure LACP between a leaf switch and a Layer 2 node. At the same time, you can disable LACP between the hypervisor and the Layer 2 switch by disabling LACP under the AEP override policy.

An AEP is required to deploy any VLAN pools on the leaf switches. It is possible to reuse the encapsulation pools (for example, VLAN) across different leaf switches. An AEP implicitly provides the scope of the VLAN pool (associated to the domain) to the physical infrastructure.

Note

- An AEP provisions the VLAN pool (and associated VLANs) on the leaf. The VLANs are not actually enabled on the port. No traffic flows unless an EPG is deployed on the port.
- Without VLAN pool deployment using an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
  - A particular VLAN is provisioned or enabled on the leaf port based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter.
  - If you wish to set the VMM encapsulation statically in the EPG, you must use a static pool. If you have a mix of static and dynamic allocations, create a dynamic pool and add a block within that pool with static mode.
- A leaf switch does not support overlapping VLAN pools. Different overlapping VLAN pools must not be associated with the same AEP that is associated through a domain.

Prerequisites for Creating a VMM Domain Profile

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

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

Custom User Account with Minimum VMware vCenter Privileges

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

- Alarms
- Distributed Switch
- dvPort Group
- Folder
- Host
    - Advanced Setting
    - Local operations.Reconfigured Virtual Machine
    - Network Configuration
- Network
- Virtual Machine
    - Virtual machine.Configuration.Modify device settings
    - Virtual machine.Configuration.Settings

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

Creating a VMM Domain Profile

In this section, examples of a VMM domain are vCenter domain or vCenter and vShield domains.

Creating a vCenter Domain Profile Using the Advanced GUI

An overview of the tasks performed in the creation of a vCenter Domain are as follows (details are in the steps that follow):

- Create/select a switch profile
- Create/select an interface profile
- Create/select an interface policy group
- Create/select VLAN pool
- Create vCenter domain
- Create vCenter credentials

Procedure

Step 1 On the menu bar, click **FABRIC** > **Access Policies**.

Step 2 In the **Navigation** pane, click **Switch Policies**.

Step 3 Right-click **Switch Policies**, and click **Configured Interfaces, PC, and VPC**.

Step 4 In the **Work** pane, in the **Configured Switch Interfaces** area, expand **Switch Profile** and perform the following actions:

Figure 3: Representative Screenshot for Configure Interface, PC, and VPC Dialog Box

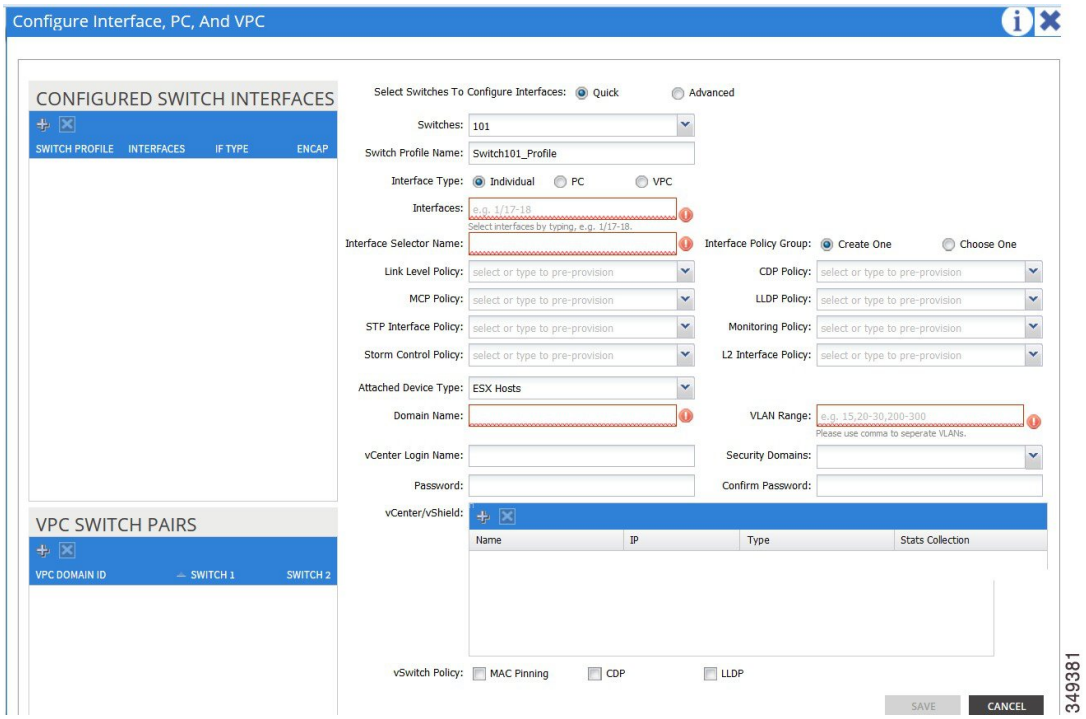

a) In the **Select Switches to Configure Interfaces** field, the **Quick** radio button is automatically checked.
b) From the **Switches** field drop-down list, choose the appropriate leaf ID.
c) In the **Switch Profile Name** field, the switch profile name automatically populates.
d) Click the + icon to configure the switch interfaces.
e) In the **Interface Type** field, check the appropriate radio button.
f) In the **Interfaces** field, enter the desired interface range.
g) In the **Interface Selector Name** field, enter the selector name where the ESX ports will be connected.
h) From the **Link Level Policy** drop-down list, choose the desired link level policy.
i) From the **CDP Policy** drop-down list, choose the desired CDP policy.
Note Similarly choose the desired interface policies from the available policy fields.
j) In the **Attached Device Type** field, choose **ESX Hosts**.
k) In the **Domain Name** field, enter the domain name.
l) In the **VLAN Range** field, enter the VLAN range as appropriate.

Note We recommend a range of at least 200 VLAN numbers. Do not define a range that includes VLAN 4, because that VLAN is for internal use.

m) In the **vCenter Login Name** field, enter the login name.
n) (Optional) From the **Security Domains** drop-down list, choose the appropriate security domain.
o) In the **Password** field, enter a password.
p) In the **Confirm Password** field, reenter the password.
q) Expand **vCenter/vShield**.

Step 5 In the **Create vCenter/vShield Controller** dialog box, enter the appropriate information, and click **Save**.

Step 6 In the **Configure Interface, PC, And VPC** dialog box, in the **vSwitch Policy** field, check the desired check box to enable CDP or LLDP. Click **Save**, and click **Submit**.

Step 7 Verify the new domain and profiles, by performing the following actions:

a) On the menu bar, choose **VM Networking** > **Inventory**.
b) In the **Navigation** pane, expand **VMware** > ***Domain_name*** > **vCenter_name**.

In the **Work** pane, under **Properties**, view the VMM domain name to verify that the controller is online. In the **Work** pane, the vCenter properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the vCenter server is established, and the inventory is available.

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

An overview of the tasks performed in the creation of a vCenter and vShield domains are as follows (details are in the steps that follow):

- Create/select a switch profile
- Create/select an interface profile
- Create/select an interface policy group
- Create/select VLAN pool
- Create vCenter and vShield domains
- Create vCenter and vShield credentials

Before You Begin

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

Procedure

Step 1 On the menu bar, click **FABRIC** > **Access Policies**.

Step 2 In the **Navigation** pane, click **Switch Policies**.

Step 3 Right-click **Switch Policies**, and click **Configured Interfaces, PC, and VPC**.

Step 4 In the **Work** pane, in the **Configured Switch Interfaces** area, expand **Switch Profile** and perform the following actions:

Figure 4: Representative Screenshot for Configure Interface, PC, and VPC Dialog Box

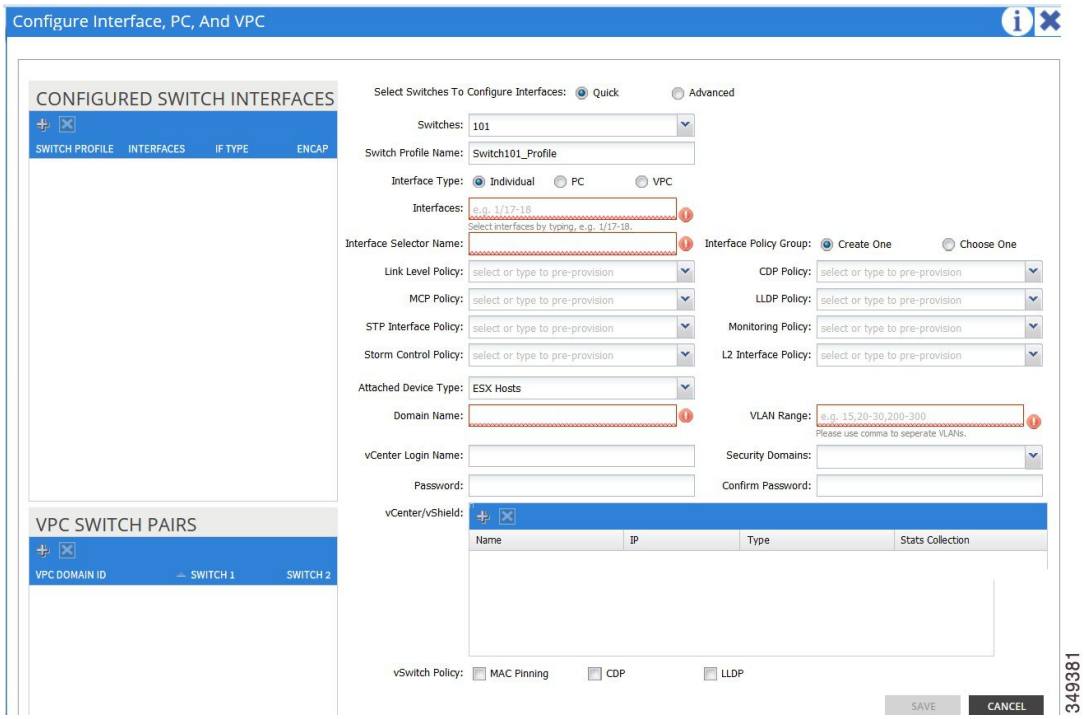

a) In the **Select Switches to Configure Interfaces** field, the **Quick** radio button is automatically checked.
b) From the **Switches** field drop-down list, choose the appropriate leaf IDs.
c) In the **Switch Profile Name** field, the switch profile name automatically populates.
d) Click the + icon to configure the switch interfaces.
e) In the **Interface Type** field, check the appropriate radio button.
f) In the **Interfaces** field, enter the desired interface range.
g) In the **Interface Selector Name** field, enter the selector name where the ESX ports will be connected.
h) From the **Link Level Policy** drop-down list, choose the desired link level policy.
i) From the **CDP Policy** drop-down list, choose the desired CDP policy.
Note Similarly choose the desired interface policies from the available policy fields.
j) From the **Attached Device Type** drop-down list, choose the appropriate device type.
k) In the **Domain Name** field, enter the domain name.
l) In the **VLAN Range** field, enter the VLAN range as appropriate.

Note We recommend a range of at least 200 VLAN numbers. Do not define a range that includes VLAN 4, because that VLAN is for internal use.

m) In the **vCenter Login Name** field, enter the login name.
n) In the **Password** field, enter a password.
o) In the **Confirm Password** field, reenter the password.
p) Expand **vCenter/vShield**.

Step 5 In the **Create vCenter/vShield Controller** dialog box, enter the appropriate information.

Step 6 In the **vSwitch Policy** field, check the check boxes for the desired vSwitch policies. Click **Save**.

Step 7 In the **Configure Interface, PC, and vPC** dialog box, in the **vSwitch Policy** field, check the desired check box to enable CDP or LLDP. Click **Save**, and click **Submit**.

Step 8 Verify the new domain and profiles, by performing the following actions:

a) On the menu bar, choose **VM Networking** > **Inventory**.
b) In the **Navigation** pane, expand , and click **VMware** > ***Domain_name*** > **vCenter_name**.

In the **Work** pane, under **Properties**, view the VMM domain name to verify that the controller is online. In the **Work** pane, the vCenter properties are displayed including the operational status. The displayed information confirms that connection from the APIC controller to the vCenter server is established, and the inventory is available.

Creating Tenants, VRF, and Bridge Domains

Tenants Overview

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

Tenant Creation

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

VRF and Bridge Domains

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

For details about enabling IPv6 Neighbor Discovery see the related KB article, *KB: Creating a Tenant, VRF, and Bridge Domain with IPv6 Neighbor Discovery* .

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

- If you have a public subnet when you configure the routed outside, you must associate the bridge domain with the outside configuration.

Procedure

Step 1 On the menu bar, click **TENANT** > **Add Tenant**.

Step 2 In the **Create Tenant** dialog box, perform the following tasks:

a) In the **Name** field, enter a name.
b) Click the **Security Domains** + icon to open the **Create Security Domain** dialog box.
c) In the **Name** field, enter a name for the security domain. Click **Submit**.
d) In the **Create Tenant** dialog box, check the check box for the security domain that you created, and click **Submit**.

Step 3 In the **Navigation** pane, expand **Tenant-name** > **Networking**, and in the **Work** pane, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:

a) In the **Name** field, enter a name.
b) Click **Submit** to complete the VRF configuration.

Step 4 In the **Networking** pane, drag the **BD** icon to the canvas while connecting it to the **VRF** icon. In the **Create Bridge Domain** dialog box that displays, perform the following tasks:

a) In the **Name** field, enter a name.
b) Click the **L3 Configurations** tab.
c) Expand **Subnets** to open the **Create Subnet** dialog box, enter the subnet mask in the **Gateway IP** field and click **OK**.
d) Click **Submit** to complete bridge domain configuration.

Step 5 In the **Networks** pane, drag the **L3** icon down to the canvas while connecting it to the **VRF** icon. In the **Create Routed Outside** dialog box that displays, perform the following tasks:

a) In the **Name** field, enter a name.
b) Expand **Nodes And Interfaces Protocol Profiles** to open the **Create Node Profile** dialog box.
c) In the **Name** field, enter a name.
d) Expand **Nodes** to open the **Select Node** dialog box.
e) In the **Node ID** field, choose a node from the drop-down list.
f) In the **Router ID** field, enter the router ID.
g) Expand **Static Routes** to open the **Create Static Route** dialog box.
h) In the **Prefix** field, enter the IPv4 or IPv6 address.
i) Expand **Next Hop Addresses** and in the **Next Hop IP** field, enter the IPv4 or IPv6 address.
j) In the **Preference** field, enter a number, then click **UPDATE** and then **OK**.
k) In the **Select Node** dialog box, click **OK**.
l) In the **Create Node Profile** dialog box, click **OK**.

m) Check the **BGP**, **OSPF**, or **EIGRP** check boxes if desired, and click **NEXT**. Click **OK** to complete the Layer 3 configuration.

To confirm L3 configuration, in the **Navigation** pane, expand **Networking** > **VRFs**.

Configuring External Connectivity for Tenants

Note The MP-BGP route reflector and the OSPF external routed network protocols do not work if you are using the simulator.

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

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

Configuring an MP-BGP Route Reflector Using the Advanced GUI

Procedure

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

Step 2 In the **Navigation** pane, expand **Pod Policies** > **Policies** > **BGP Route Reflector default**, right-click **BGP Route Reflector default**, and click **Create Route Reflector Node Policy EP**.

Step 3 In the **Create Route Reflector Node Policy EP** dialog box, from the **Spine Node** drop-down list, choose the appropriate spine node. Click **Submit**.

Note Repeat the above steps to add additional spine nodes as required.

The spine switch is marked as the route reflector node.

Step 4 In the **BGP Route Reflector default** properties area, in the **Autonomous System Number** field, choose the appropriate number. Click **Submit**.

Note The autonomous system number must match the leaf connected router configuration if Border Gateway Protocol (BGP) is configured on the router. If you are using routes learned using static or Open Shortest Path First (OSPF), the autonomous system number value can be any valid value.

Step 5 In the **Navigation** pane, expand and right-click **Policy Groups**, and click **Create POD Policy Group**.

Step 6 In the **Create POD Policy Group** dialog box, in the **Name** field, enter the name of a pod policy group.

Step 7 In the **BGP Route Reflector Policy** drop-down list, choose the appropriate policy (default). Click **Submit**. The BGP route reflector policy is associated with the route reflector pod policy group, and the BGP process is enabled on the leaf switches.

Step 8 In the **Navigation** pane, choose **Pod Policies** > **Profiles** > **default**. In the **Work** pane, from the **Fabric Policy Group** drop-down list, choose the pod policy that was created earlier. Click **Submit**.
The pod policy group is now applied to the fabric policy group.

Verifying the MP-BGP Route Reflector Configuration

Procedure

Step 1 Verify the configuration by performing the following actions:

a) Use secure shell (SSH) to log in as an administrator to each leaf switch as required.

b) Enter the **show processes | grep bgp** command to verify the state is S.
If the state is NR (not running), the configuration was not successful.

Step 2 Verify that the autonomous system number is configured in the spine switches by performing the following actions:

a) Use the SSH to log in as an administrator to each spine switch as required.

b) Execute the following commands from the shell window

Example:
cd /mit/sys/bgp/inst

Example:
grep asn summary

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

Creating an OSPF External Routed Network for Management Tenant Using the Advanced GUI

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

Procedure

Step 1 On the menu bar, choose **TENANTS** > **mgmt**.

Step 2 In the **Navigation** pane, expand **Networking** > **External Routed Networks**.

Step 3 Right-click **External Routed Networks**, and click **Create Routed Outside**.

Step 4 In the **Create Routed Outside** dialog box, perform the following actions:

a) In the **Name** field, enter a name (RtdOut).
b) Check the **OSPF** check box.
c) In the **OSPF Area ID** field, enter an area ID.
d) In the **OSPF Area Control** field, check the appropriate check box.
e) In the **OSPF Area Type** field, choose the appropriate area type.
f) In the **OSPF Area Cost** field, choose the appropriate value.
g) In the **VRF** field, from the drop-down list, choose the VRF (inb).
Note This step associates the routed outside with the in-band VRF.
h) From the **External Routed Domain** drop-down list, choose the appropriate domain.
i) Click the + icon for **Nodes and Interfaces Protocol Profiles** area.

Step 5 In the **Create Node Profile** dialog box, perform the following actions:

a) In the **Name** field, enter a name for the node profile. (borderLeaf).
b) In the **Nodes** field, click the + icon to display the **Select Node** dialog box.
c) In the **Node ID** field, from the drop-down list, choose the first node. (leaf1).
d) In the **Router ID** field, enter a unique router ID.
e) Uncheck the **Use Router ID as Loopback Address** field.
Note By default, the router ID is used as a loopback address. If you want them to be different, uncheck the **Use Router ID as Loopback Address** check box.
f) Expand **Loopback Addresses**, and enter the IP address in the **IP** field. Click **Update**, and click **OK**.
Enter the desired IPv4 or IPv6 IP address.
g) In the **Nodes** field, expand the + icon to display the **Select Node** dialog box.
Note You are adding a second node ID.
h) In the **Node ID** field, from the drop-down list, choose the next node. (leaf2).
i) In the **Router ID** field, enter a unique router ID.
j) Uncheck the **Use Router ID as Loopback Address** field.
Note By default, the router ID is used as a loopback address. If you want them to be different, uncheck the **Use Router ID as Loopback Address** check box.
k) Expand **Loopback Addresses**, and enter the IP address in the **IP** field. Click **Update**, and click **OK**. Click **OK**.
Enter the desired IPv4 or IPv6 IP address.

Step 6 In the **Create Node Profile** dialog box, in the **OSPF Interface Profiles** area, click the + icon.

Step 7 In the **Create Interface Profile** dialog box, perform the following tasks:

a) In the **Name** field, enter the name of the profile (portProf).
b) In the **Interfaces** area, click the **Routed Interfaces** tab, and click the + icon.
c) In the **Select Routed Interfaces** dialog box, in the **Path** field, from the drop-down list, choose the first port (leaf1, port 1/40).

d) In the **IP Address** field, enter an IP address and mask. Click **OK**.
e) In the **Interfaces** area, click the **Routed Interfaces** tab, and click the + icon.
f) In the **Select Routed Interfaces** dialog box, in the **Path** field, from the drop-down list, choose the second port (leaf2, port 1/40).
g) In the **IP Address** field, enter an IP address and mask. Click **OK**.
Note This IP address should be different from the IP address you entered for leaf1 earlier.
h) In the **Create Interface Profile** dialog box, click **OK**.

The interfaces are configured along with the OSPF interface.

Step 8 In the **Create Node Profile** dialog box, click **OK**.

Step 9 In the **Create Routed Outside** dialog box, click **Next**.
The **Step 2 External EPG Networks** area is displayed.

Step 10 In the **External EPG Networks** area, click the + icon.

Step 11 In the **Create External Network** dialog box, perform the following actions:
a) In the **Name** field, enter a name for the external network (extMgmt).
b) Expand **Subnet** and in the **Create Subnet** dialog box, in the **IP address** field, enter an IP address and mask for the subnet.
c) In the **Scope** field, check the desired check boxes. Click **OK**.
d) In the **Create External Network** dialog box, click **OK**.
e) In the **Create Routed Outside** dialog box, click **Finish**.
Note In the **Work** pane, in the **External Routed Networks** area, the external routed network icon (RtdOut) is now displayed.

Deploying an Application Policy

Three-Tier Application Deployment

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

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

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

To deploy an application policy, you must create the required application profiles, filters, and contracts. Typically, the APIC fabric hosts a three-tier application within a tenant network. In this example, the application

is implemented by using three servers (a web server, an application server, and a database server). See the following figure for an example of a three-tier application.

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

Figure 5: Three-Tier Application Diagram

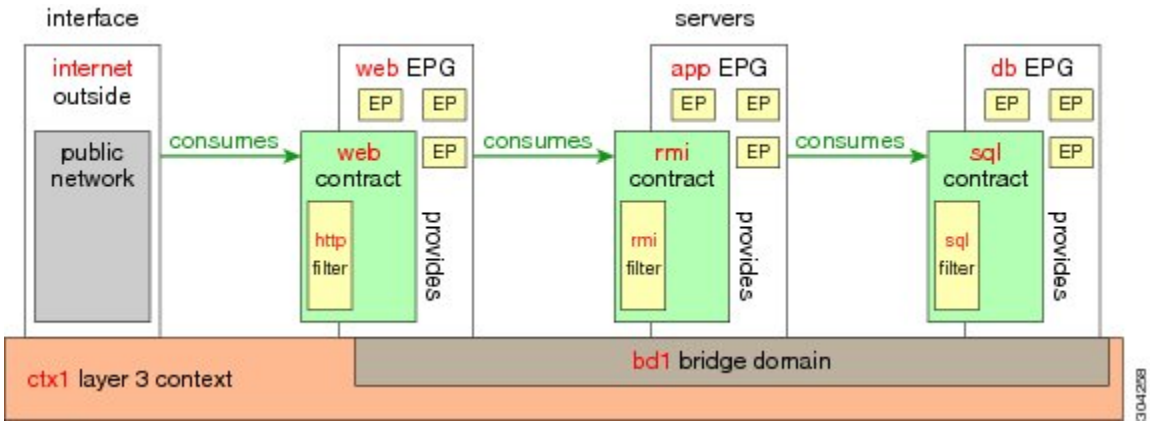

Parameters to Create a Filter for http

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

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

Parameters to Create Filters for rmi and sql

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

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

Example Application Profile Database

The application profile database in this example is as follows:

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

Deploying an Application Policy Using the GUI

Creating a Filter Using the GUI

Create three separate filters. In this example they are HTTP, RMI, SQL. This task shows how to create the HTTP filter. The task is identical for creating the other filters.

Before You Begin

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

Procedure

Step 1 On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the **tenant** > **Security Policies**, right-click **Filters**, and click **Create Filter**.

Note In the **Navigation** pane, you expand the tenant where you want to add filters.

Step 2 In the **Create Filter** dialog box, perform the following actions:

a) In the **Name** field, enter the filter name (http).
b) Expand **Entries**, and in the **Name** field, enter the name (Dport-80).
c) From the **EtherType** drop-down list, choose the EtherType (IP).
d) From the **IP Protocol** drop-down list, choose the protocol (tcp).
e) From the **Destination Port/Range** drop-down lists, choose **http** in the **From** and **To** fields. (http)
f) Click **Update**, and click **Submit**.
The newly added filter appears in the **Navigation** pane and in the **Work** pane.

Step 3 Expand **Entries** in the **Name** field. Follow the same process to add another entry with HTTPS as the **Destination** port, and click **Update**.
This new filter rule is added.

Step 4 Follow the same process in the earlier steps to create two more filters (rmi and sql) and use the parameters provided in Parameters to Create Filters for rmi and sql.

Creating a Contract Using the GUI

Procedure

Step 1 On the menu bar, choose **TENANTS** and the tenant name on which you want to operate. In the **Navigation** pane, expand the **tenant** > **Security Policies**.

Step 2 Right-click **Contracts** > **Create Contract**.

Step 3 In the **Create Contract** dialog box, perform the following tasks:

a) In the **Name** field, enter the contract name (web).
b) Click the + sign next to **Subjects** to add a new subject.
c) In the **Create Contract Subject** dialog box, enter a subject name in the **Name** field. (web)
d) **Note** This step associates the filters created that were earlier with the contract subject.

In the **Filter Chain** area, click the + sign next to **Filters**.
e) In the dialog box, from the drop-down menu, choose the filter name (http), and click **Update**.

Step 4 In the **Create Contract Subject** dialog box, click **OK**.

Step 5 Create two more contracts for rmi and for sql following the same steps in this procedure. For the rmi contract, choose the rmi subject and for sql, choose the sql subject.

Creating an Application Profile Using the GUI

Procedure

Step 1 On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the tenant, right-click **Application Profiles**, and click **Create Application Profile**.

Step 2 In the **Create Application Profile** dialog box, in the **Name** field, add the application profile name (OnlineStore).

Creating EPGs Using the GUI

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

Procedure

Step 1 Expand **EPGs**. In the **Create Application EPG** dialog box, perform the following actions:

a) In the **Name** field, add the EPG name (db).
b) In the **Bridge Domain** field, choose the bridge domain from the drop-down list (bd1).
c) Check the **Associate to VM Domain Profiles** check box. Click **Next**.
d) In the **Step 2 for Specify the VM Domains** area, expand **Associate VM Domain Profiles** and from the drop-down list, choose the desired VMM domain. Click **Update** and click **OK**.

Step 2 In the **Create Application Profile** dialog box, create two more EPGs. The three EPGs should be db, app, and web in the same bridge domain and data center.

Consuming and Providing Contracts Using the GUI

You can associate contracts that were created earlier to create policy relationships between the EPGs.

When you name the provided and consumed contracts, verify that you give the same name for both provided and consumed contracts.

Procedure

Step 1 Click and drag across the APIC GUI window from the db EPG to the app EPG.
The **Add Consumed Contract** dialog box is displayed.

Note The db, app, and web EPGs are displayed as icons.

Step 2 In the **Name** field, from the drop-down list, choose **sql** contract. Click **OK**.
This step enables the db EPG to provide the sql contract and the app EPG to consume the sql contract.

Step 3 Click and drag across the APIC GUI screen from the app ePG to the web EPG.

The **Add Consumed Contract** dialog box is displayed.

Step 4 In the **Name** field, from the drop-down list, choose **rmi** contract. Click **OK**.
This step enables the app EPG to provide the rmi contract and the web EPG to consume the rmi contract.

Step 5 Click the web EPG icon, and click the + sign in the **Provided Contracts** area.
The **Add Provided Contract** dialog box is displayed.

Step 6 In the **Name** field, from the drop-down list, choose **web** contract. Click **OK**. Click **Submit**.
You have created a three-tier application profile called OnlineStore.

Step 7 To verify, in the **Navigation** pane, navigate to and click **OnlineStore** under **Application Profiles**.
In the **Work** pane, you can see the three EPGs app, db, and web are displayed.

Step 8 In the **Work** pane, choose **Operational** > **Contracts**.
You can see the EPGs and contracts displayed in the order that they are consumed and provided.