



Managing Tenants

- [Tenants, on page 2](#)
- [Virtual Routing and Forwarding \(VRF\), on page 6](#)
- [BGP Timers, on page 11](#)
- [Bridge Domains, on page 12](#)
- [Application Profiles, on page 17](#)
- [Endpoint Groups, on page 19](#)
- [Contracts, on page 26](#)
- [Adding Contracts to EPGs, on page 29](#)
- [Contract Labels, on page 32](#)
- [Fabric Extender \(FEX\), on page 34](#)
- [Fabric Ext Connection Policies, on page 35](#)
- [Filter Chain, on page 37](#)
- [Data Plane Policing, on page 39](#)
- [FHS Trust Policy, on page 40](#)
- [Creating a BGP Address Family Context Policy, on page 41](#)
- [NetFlow Monitor Policy, on page 42](#)
- [Bidirectional Forwarding Detection, on page 42](#)
- [Hot Standby Router Protocol, on page 44](#)
- [Routed Outside, on page 47](#)
- [Dynamic Host Configuration Protocol, on page 56](#)
- [IGMP Interface Policy, on page 59](#)
- [Route Tag Policy, on page 60](#)
- [EIGRP Address Family Context Policy, on page 60](#)
- [OSPF Timers, on page 61](#)
- [IGMP Snoop Policy, on page 62](#)
- [MLD Snoop Policy, on page 63](#)
- [Monitoring Policy, on page 64](#)
- [NetFlow Monitor Policy, on page 65](#)
- [Flow Record, on page 66](#)
- [Route Maps, on page 66](#)
- [Creating a Set Rules for Route Map, on page 70](#)
- [Adding a Context to a Route Map or Profile, on page 73](#)

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 4](#).

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

See [Creating a VRF, on page 7](#).

Step 3 Add Bridge Domain to the VRF.

See [Adding a Bridge Domain to VRF, on page 13](#).

Step 4 Create Application Profiles.

See [Creating an Application Profile for the Tenant, on page 18](#).

Step 5 Create EPGs.

See [Adding an EPG, on page 19](#).

Step 6 Add domain to EPGs.

See [Adding a Domain to an EPG, on page 20](#).

Step 7 Add Static path to EPGs.

See [Adding a Static Path to EPG, on page 22](#).

Step 8 Create Contracts.

See [Creating Contracts, on page 27](#).

Step 9 Add contracts to EPGs.

See [Adding a Consumed Contract to an EPG, on page 30](#).

See [Adding a Provided Contract to an EPG, on page 29](#).

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:
- Enter the globally unique identifier (GUID) for an SCVMM provider. GUID must be in the format :
xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx.
 - 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:
- Enter a unique name for the private network and alias.
 - Click **Tags** and check the tag that you want to use.
 - 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**.
 - 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.
 - Check the **BD Enforcement Status** check box to enable BD enforcement status.
 - 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**.
 - Enter the description for the private network.
 - 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.
 - 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.
 - 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.
 - Enter the comma separated DNS labels for private network.
 - 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:
- From the **Type** drop-down list, choose **Import** or **Export**.
 - 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:
- From the **BGP Address Family Type** drop-down list, choose IPv4 or IPv6.
 - 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:
- From the **OSPF Address Family Type** drop-down list, choose IPv4 or IPv6.
 - 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 14](#).

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 15](#).



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:
- Add a unique name and description for the Bridge Domain.
 - Click **Select** and check the tag you want to use for the account.
 - From the **Type** drop-down list, choose regular or FC based on your requirement.
 - Check **Advertise Host Routes** to enable the host route.
 - 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.

- Click **Select** and check the VRF you want to use for the account.
- 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:
- From the **IP Type** drop-down list, choose **IPv4** or **IPv6**.
 - In the **Gateway IP (Address)** field, enter the IP address of the default gateway.
 - In the **Gateway IP (Prefix)** field, enter a prefix in the range of 1-32 that starts with "/"x" for IPv4 and 1-128 that starts with "/"x" for IPv6.
 - 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.
 - Check the **Treat as virtual IP address** check box to treat the given IP as virtual address.
 - Check the **Make this IP address as primary** check box to make this IP address as primary for the subnet.
 - Check the **Public Subnet** check box to export it to a routed connection.
 - Check the **Private Subnet** check box to apply the subnet only within its tenant.
 - Check the **Scope (Shared between VRFs)** check box to enable subnet control for a shared subnet.
 - Check the **Scope (Advertised Externally)** check box to enable subnet control for a public subnet.
 - Check the **Scope (Private to VRF)** check box to enable subnet control for a private subnet.
 - Enter a description for the subnet.
 - Check the **Subnet Control (No Default SVI Gateway)** check box to enable the No Default SVI Gateway.
 - Check the **Subnet Control (Querier IP)** check box to apply specific protocol to the subnet. Querier IP enables IGMP Snooping on the subnet.
 - 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.
 - Click **Select** and check the route profile that you want to use for this bridge domain.

The route profile specifies policies for external networks.
 - 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:
- 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**.
 - Click **Select** and check the DHCP relay policy that you want to use for the tenant bridge domain.
 - Click **Select** and check the DHCP option policy that you want to use.
 - 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:
- 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.
 - Check the **Auto Configuration** check box to enable the auto-configuration and use the configured prefix in stateless address allocation.
 - 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.
 - 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 22](#).

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 24](#) and [Adding a Domain to an EPG, on page 20](#).

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:
- 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
 - 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.
 - 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`.
 - 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.

- 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.

- In the **Primary VLAN for Micro-Seg** field, enter the encapsulation for a primary VLAN.
- 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:
 - Click **Select** and check the node that you want to add to the EPG.
 - 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.
 - 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:
- Click **Select** and check the application profile that you want to use in the EPG contract.
 - 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:
- Click **Select** and check the contract that you want to add to the EPG.
 - 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
 - Enter the label for the provided contract.
 - Enter the subject label for the provided contract.
 - 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:
- Click **Select** and check the contract that you want to add to the EPG.
 - 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
 - Enter the label for the consumed contract.
 - Enter the subject label for the consumed contract.
 - 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:
- Click **Select** and check the contract interface that you want to add to the EPG.
 - 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
 - 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:
- Enter a unique name and description for the fabric external routing profile.
 - 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:
- Click **Select** and choose a pod to which you want to add the connection profile.
 - (Optional) The **Password** field appears only when you choose a virtual pod. Enter the BGP password for the connection profile.
 - Re-enter the password for confirmation.
 - 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.
 - 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:
- Enter a unique name and description for the FHS trust policy.
 - 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.
 - 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.
 - 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.
 - Check the **Trust ARP** check box to allow the port to trust ARP packets.
 - Check the **Trust ND** check box to allow the port to trust neighbor discovery (ND) protocol packets.
 - 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:
- Enter a unique name and short description for the BGP Address Family Context policy.
 - In the **eBGP 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.
 - 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.
 - 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.
 - 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.
 - 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.
 - 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:
- Enter the name and description for the BFD interface policy.
 - From the **Admin State** drop-down list, choose **Enabled** to enable the admin state for the BFD interface policy.
 - Check the **Control State** check box to enable the control state for the BFD interface policy.
 - 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.
 - 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.
 - 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.
 - 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.
 - 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:
- 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**.
 - 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:
- Enter a unique name, description, and alias for the external routed network.
 - Click **Select** and check the tag that you want to use for the external routed network.
 - Click **Select** and check the network that you want to use for the external routed network.
 - 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**).
- Click **Select** and check the external routed domain that you want to use for the external routed network.
 - Check the **BGP** check box to configure BGP as the routing protocol.
 - 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.
 - 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.
 - Check the **PIM** check box to configure protocol independent multicast (PIM) protocol as the routing protocol.
 - 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.
 - Click **Select** and check the route profile for interleaf.
 - Enter the label for the provider.
 - Enter the label for the consumer.
 - 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:
- Enter a unique name and description for the logical node profile.
 - 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:
- 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.
 - Click **Select** and choose a logical node that needs to be added to the logical node profile.
 - Enter the IP address of the local routing device.
 - 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:
- Enter an IP address for the static route with mask. For example, 10.10.10.0/24.
 - 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.
 - Check the **Route Control BFD** check box to enable BFD on static route to provide fast forwarding path failure detection.
 - 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:
- 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.
 - 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:
- 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.
 - 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:
- **I3-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:
- **I3-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:

- In the **Peer Address** field, enter the peer IP address in the IPv4, IPv6, IPv4/prefix, or IPv6/prefix format.
- Enter a short description for the BGP peer connectivity profile.
- Check the **Allow Self AS** check box to allow the readvertisement of all prefixes that contain duplicate autonomous system numbers (ASNs).
- 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.
- Check the **Disable Peer AS Check** check box to disable checking of the peer AS number when advertising.
- Check the **Next-hop Self** check box to always set the next hop attribute to the local peering address.
- Check the **Send Community** check box to send the community attribute to the neighbor.
- Check the **Send Extended Community** check box to send the extended community attribute to the neighbor.
- In the **Password** field, enter the password for BGP MD5 authentication.
- In the **Confirm Password** field, re-enter the password for BGP MD5 authentication.
- The **Allow Self AS Count** field displays the allowed AS number count.
- Check the **Bidirectional Forwarding Detection** check box to enable BFD in the BGP peer connectivity profile.
- Check the **Disable Connected Check** check box to skip the connected check and attempt the connection to the peering address.
- 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.
- 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.
- 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.
- Check **Remove private AS** check box to remove private AS numbers from outbound route updates to an eBGP peer.
- 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.
- In the **BGP Peer Prefix Policy** field, click **Select** and choose a BGP peer prefix policy to be associated with the BGP peer connectivity profile.
- In the **Remote Autonomous System Number** field, enter a value in the range of 1 to 429467925 as the remote ASN.
- From the **Local-AS Number Config** drop-down list, choose any one of the option as the method for configuring local ASN.
- 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:
- From the **Scope** drop-down list, choose **infra** or **tenant**.
 - 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:
- 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 Account > 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:
- Enter the name and description for the IGMP interface policy.
 - Check the **Allow v3 ASM** check box to accept IGMPv3 reports for non SSM groups.
 - Check the **Fast Leave** check box to configure fast leave for the VLAN/BD.
 - Enter the group timeout in seconds to configure group membership timeout in all VLAN/BDs. The default value is 260 seconds.
 - Enter the query interval in seconds to configure interval between group-specific query transmissions. The default value is 125 seconds.
 - Enter the query response interval in seconds. The default value is 10 seconds.
 - Enter the last member count. The default value is 2.
 - Enter the response time of last member in seconds. The default value is 1 second.
 - Enter the start-up query count to configure the number of queries to be sent at start-up. The default value is 2.
 - Enter the start-up query interval in seconds to configure the query interval at start-up. The default value is 31 seconds.
 - Enter the querier timeout for IGMP in seconds. The default value is 255 seconds.
 - 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.
 - 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**.
 - Click **Select** and choose a report policy route map for the IGMP.
 - Click **Select** and choose a static report route map for the IGMP.
 - Enter the maximum number of allowed multicast entries.
 - Enter the reserved multicast entries.
 - 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:
- Enter a unique name and short description for the EIGRP Address Family Context policy.
 - 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.
 - 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.
 - 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.
 - 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.
 - 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:
- Enter a unique name and short description for the OSPF timer policy.
 - In the **Bandwidth Reference (MBPS)** field, enter a value in the range of 0 to 4000000. The default value is 40,000.
 - 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.
 - 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:
- Enter the name and description for the IGMP snoop policy.
 - From the **Admin State** drop-down, choose state for the IGMP snoop policy. The default is **enabled**.
 - Check the **Fast Leave** check box to enable IGMP V2 immediate dropping of queries through this policy.
 - Check the **Enable querier** check box to enable the IGMP querier activity through this policy.
 - 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.
 - Enter the query response interval in seconds. When a host receives the query packet, it starts counting to a random value, less than the maximum response time. When this timer expires, host replies with a report. The default value is 10 seconds.
 - 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.
 - Enter the start-up query count to configure the number of queries to be sent at start-up. The default value is 2.
 - 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:
- Enter the name and description for the MLD Snoop policy.
 - From the **Admin State** drop-down list, choose **Enabled** or **Disabled** to enable or disable this entire policy. The default option is **Enabled**.
 - Check **Fast leave** to enable MLD v1 immediate dropping of queries through this policy.
 - 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:
- Enter the name and description for the NetFlow monitor policy.
 - 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:
- Enter the IPv4 or IPv6 address as the specific prefix for match rule.
 - 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.
 - 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:
- Enter a unique name for the match community factor.
 - Choose **Transitive** or **Non transitive** as the scope of the community.
 - 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:
- Enter a value in the range of 0 to 65535 as the preferred order for the route map entry.
 - Enter a valid group IP address with prefix to match an IP multicast packet based on group. For example: 10.2.3.1/22.
 - Enter a valid source IP address with prefix to match an IP multicast packet based on source. For example: 10.1.2.3/22.
 - 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.
 - 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:
- Enter a unique name and description for the action rule profile.
 - 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.
 - Click to set the route tag for the action rule profile. Enter the tag value in the range of 0 to 4294967295.
 - 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:
- 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.
 - Enter the AS number in the range of 1 to 4294967295.
 - 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:
- 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.
 - 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:
- Enter a unique name for the context.
 - Enter the order in which the context has to be inserted in to the route map or profile.
 - Choose **Permit** to permit routes that match criteria defined in match rule.
 - Enter a short description for the context.
 - 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**.
-