Provisioning Core ACI Fabric Services

This chapter contains the following sections:

- Link-level policies, on page 1
- Link flap policies, on page 3
- Configure a link flap policy by using the GUI, on page 4
- Time synchronization and NTP, on page 4
- DHCP relay policy, on page 11
- DNS service policy, on page 17
- Configuring Custom Certificates, on page 21
- Provisioning Fabric Wide System Settings, on page 25
- Provisioning Global Fabric Access Policies, on page 50
- Per Port Policies, on page 54
- Mis-cabling Protocol Interface policy, on page 56

Link-level policies

A link-level policy is an access policy that configures the following physical-layer (Layer 1) interface settings:

- Autonegotiation
- Port speed
- Link debounce

Electromagnetic interference retrain

Electromagnetic interference (EMI) retrain detects and filters EMI-generated noise to prevent unnecessary link flaps.

- Supported beginning with Cisco APIC Release 5.2(4).
- Configured through a link-level policy.
- Supported only on Cisco N9K-C93108TC-EX and N9K-C93108TC-FX leaf switches that use copper cables.

EMI retrain configuration requirements

To enable EMI retrain, choose **enable** for the **EMI Retrain** property in the link-level policy.

Configure a link-level policy by using the GUI

A link-level policy defines the physical-layer settings that are applied to switch interfaces.

Procedure

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

Step 2 In the **Navigation** pane, choose **Policies** > **Interface** > **Link Level** .

Step 3 Right-click **Link Level** and choose **Create Link Level Policy** .

Step 4 In the **Create Link Level Policy** dialog box, configure the fields for the interface.

For **Speed** , use the default value, **Inherit** . With this value, the Cisco APIC determines the interface speed based on the transceiver installed in the switch port.

See the field tooltips for additional information.

Note
The **Auto** speed option is not supported and is planned for deprecation. Use **Inherit** instead.

Step 5 Click **Submit** .

Port bring-up delay

Port bring-up delay is a link-level policy parameter that specifies, in milliseconds, how long the system delays decision feedback equalizer tuning during port initialization.

- Configurable beginning with Cisco APIC Release 4.2(5).
- Helps prevent cyclic redundancy check (CRC) errors during link bring-up.
- Use only when required for compatibility with specific third-party adapters.

Note The **Port bring-up delay (milliseconds)** parameter does not apply to fabric extender (FEX) ports.

Link-level policy limitations

Link-level policy limitations prevent unsupported transceiver combinations from establishing links between Cisco Nexus 9408 switches and the Cisco APIC fabric. The limitations include:

- Incompatibility between 50G-SR-S transceivers on N9K-X9400-22L line cards and 25G-SR-S transceivers connected to the Cisco APIC fabric

- Automatic removal of speed policies when a link connects to a Cisco APIC
- Unsupported use of 10G-GLCT transceivers for Cisco APIC fabric bring-up

Operational constraints for Cisco APIC fabric links

The following limitations apply to link-level policy configurations:

- Link bring-up is not supported between a 50G-SR-S transceiver on an N9K-X9400-22L line card in a Cisco Nexus 9408 ACI-Mode switch and a 25G-SR-S transceiver connected to the Cisco APIC fabric. This transceiver combination operates only between N9K switches that use encrypted symmetric keys.
- If a link connects to a Cisco APIC, the default policy is applied to the leaf port and the speed policy is removed.
- The 10G-GLCT transceiver is not supported for Cisco APIC fabric bring-up.

Link debounce interval

The link debounce interval specifies how long a fabric interface waits after a link flap before rechecking the link status.

- Applies to leaf switch uplink ports and spine switch ports.
- Delays link-status confirmation for the configured interval after a flap.
- Keeps the interface up if the link is operational when the interval expires.

Configuration methods and default values

Configure the link debounce interval by using one of the following methods:

- GUI: Use the **Link debounce interval (msec)** field in the Cisco Application Policy Infrastructure Controller (APIC).
- CLI: Use the **[no] link debounce time <time>** command.

Note The default debounce interval is 0 ms. The recommended value is 100 ms, but choose a value that is appropriate for your fabric.

Link flap policies

A link flap policy defines when the system disables a switch port that exceeds a specified number of link-state transitions within a configured time period.

- Identifies ports that flap because of faulty, unsupported, or nonstandard cables or small form-factor pluggable (SFP) transceivers.
- Defines the maximum number of permitted link flaps within a specified time period.

- Places the port in an error-disabled state when the configured threshold is exceeded.

Link flap policy requirements and limitations

The port remains in the error-disabled state until you use the Cisco Application Policy Infrastructure Controller (APIC) to disable and then reenable the port.

Note A link flap policy does not apply to fabric extender (FEX) host interface (HIF) ports or to leaf switch models whose product IDs do not have an -EX, -FX, -FX2, -GX, or later-generation suffix.

Configure a link flap policy by using the GUI

After you configure the link flap policy, associate it with a leaf or spine node interface policy to deploy it on the node access ports.

Procedure

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

Step 2 In the **Navigation** pane, choose **Policies** > **Interface** > **Link Flap** .

Step 3 Right-click **Link Flap** and choose **Create Link Flap Policy** .

Step 4 In the **Create Link Level Policy** dialog box, configure the fields for the link flap policy.

See the field tooltips and online help for additional information.

Step 5 Click **Submit** .

Time synchronization and NTP

Time synchronization maintains consistent clock settings across fabric nodes, enabling accurate traffic-flow analysis and event correlation.

- Supports accurate analysis of traffic flows, debug information, and timestamps.
- Enables full use of atomic counters when calculating application health scores.
- Uses Network Time Protocol (NTP) as the primary method for synchronizing device clocks.

NTP configuration considerations

Incorrect time synchronization might not immediately generate a fault, but it can make operational issues more difficult to diagnose. Consider the following factors when you deploy time synchronization:

- Management scheme: The configuration depends on whether the Cisco ACI nodes and APICs use in-band or out-of-band management.

- Time-source reliability: Evaluate the reliability of the source when choosing between a private internal clock and an external public clock.
- Authentication schemes: Supported schemes include MD5, SHA-1, and AES128-CMAC. Beginning with Cisco APIC Release 6.1(1), AES128-CMAC is supported and is the recommended scheme because MD5 is considered weak and insecure.

Note When Federal Information Processing Standards (FIPS) mode is enabled, only AES128-CMAC and SHA-1 authentication schemes are supported.

In-band management NTP

Network Time Protocol (NTP) synchronization enables fabric nodes to use a consistent clock source through in-band or out-of-band management connectivity.

- Out-of-band management provides connectivity to fabric nodes from outside the Cisco ACI fabric.
- In-band management uses the internal Cisco ACI management network for IP connectivity.
- Date and Time policies define the NTP server configuration for fabric nodes.

Management NTP configuration

The following requirements apply to each management type:

- In-band management NTP: In a fabric that uses in-band management, ensure that the NTP server is reachable from the Cisco ACI in-band management network. In-band IP addresses in the Cisco ACI fabric are not reachable from outside the fabric. To use an external NTP server with in-band management, configure a policy that enables communication between the fabric and the server.

Note See the Adding Management Access section in this guide for information about in-band management access .

Note If you apply a monitoring or NTP policy to in-band management, faults are generated for a standby APIC that does not have an in-band configuration. Configure in-band management on the standby APIC to clear the faults.

NTP over IPv6

Network Time Protocol (NTP) over IPv6 enables time synchronization with servers and peers that use IPv6 addresses.

- Supports hostnames that resolve to IPv6 addresses and configurations that specify IPv6 peer addresses.

- Uses `gai.conf` address-selection settings to give IPv6 precedence over IPv4.
- Supports hostname-resolution preferences for different deployment requirements.

Configure NTP by using the GUI

Note A hostname-based NTP server might fail name resolution if the DNS server is reachable through in-band or out-of-band management. If you specify an NTP server by hostname, configure a DNS service policy that provides connectivity to the DNS providers. Also configure the appropriate DNS label for the in-band or out-of-band virtual routing and forwarding (VRF) instance associated with the management endpoint group (EPG) that you selected in the DNS profile policy.

Procedure

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

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

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

Step 4 In the **Create Date and Time Policy** dialog box, configure the policy:

a) In the **Name** field, enter a name that distinguishes the policy from other NTP configurations.

b) For **Authentication State** , choose **enabled** , expand the **NTP Client Authentication Keys** table, and enter the authentication-key information.

c) Click **+** to create an authentication key. Enter the **ID** and **Key** , check the check box in the **Trusted** column, and choose the algorithm from the **Authentication Type** drop-down list.

The available authentication algorithms are SHA-1, MD5, and AES128-CMAC.

d) Click **Update** , and then click **Next** .

e) Click **+** to add an NTP server as a provider.

f) In the **Create Providers** dialog box, configure the provider, including the following fields: **Name** , **Description** , **Minimum Polling Intervals** , and **Maximum Polling Intervals** .

- If you configure multiple providers, check the **Preferred** check box for the most reliable NTP source.
- From the **Management EPG** drop-down list, choose **Out-of-Band** if all fabric nodes can reach the NTP server through out-of-band management. If you use in-band management, see In-Band Management NTP. Click **OK** .

Repeat the provider configuration for each NTP server that you want to add.

Note
If you add a fully qualified domain name (FQDN) as an NTP provider and then add the corresponding IP address as another provider, a *Duplicate IP* fault is generated. Delete the most recently added provider configuration. For example, if you added the IP address after the FQDN, delete the IP-address provider configuration.

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

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

Step 7 In the **Create Pod Policy Group** dialog box, configure the policy group:

a) In the **Name** field, enter a name for the policy group.

b) From the **Date Time Policy** drop-down list, choose the NTP policy that you created, and click **Submit** .

The pod policy group is created. Alternatively, you can use the default pod policy group.

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

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

Step 10 In the **Properties** area, choose the pod policy group from the **Fabric Policy Group** drop-down list, and click **Submit** .

Configure NTP by using the REST API

Note A hostname-based NTP server might fail name resolution if the DNS server is reachable through in-band or out-of-band management. If you specify an NTP server by hostname, configure a DNS service policy that provides connectivity to the DNS providers. Also configure the appropriate DNS label for the in-band or out-of-band virtual routing and forwarding (VRF) instance associated with the management endpoint group (EPG) that you selected in the DNS profile policy.

Procedure

Step 1 Send the following POST request to create the NTP Date and Time policy and provider.

Example:

```
POST URL: https://APIC-IP/api/node/mo/uni/fabric/time-CiscoNTPPol.xml
                         POST payload:
                         <imdata totalCount="1">
                         <datetimePol adminSt="enabled" authSt="disabled" descr=""
dn="uni/fabric/time-CiscoNTPPol" name="CiscoNTPPol" ownerKey="" ownerTag="">
                         <datetimeNtpProv descr="" keyId="0" maxPoll="6" minPoll="4"
name="10.10.10.11" preferred="yes">
                       <datetimeRsNtpProvToEpg tDn="uni/tn-mgmt/mgmtp-default/inb-default"/>

                         </datetimeNtpProv>
                         </datetimePol>
                         </imdata>
```

Step 2 Send the following POST request to associate the Date and Time policy with the pod policy group.

Example:

```
POST URL: https://APIC-IP/api/node/mo/uni/fabric/funcprof/podpgrp-calo1/rsTimePol.xml
                         POST payload:
                         <imdata totalCount="1">
                         <fabricRsTimePol tnDatetimePolName="CiscoNTPPol">
                         </fabricRsTimePol>
                         </imdata>
```

Step 3 Send the following POST request to associate the pod policy group with the default pod profile.

Example:

```
POST URL:
https://APIC-IP/api/node/mo/uni/fabric/podprof-default/pods-default-typ-ALL/rspodPGrp.xml
                        POST payload:
                        <imdata totalCount="1">
                        <fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-calo1"
status="created">
                        </fabricRsPodPGrp>
                        </imdata>
```

Verify NTP operation by using the GUI

Procedure

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

Step 2 In the **Navigation** pane, choose **Pod Policies** > **Policies** > **Date and Time** > ***ntp_policy*** > ***server_name*** .

The *ntp_policy* variable identifies the previously created NTP policy. The **Host Name/IP Address** field supports IPv6 addresses. If you specify a hostname that resolves to both IPv4 and IPv6 addresses, configure address resolution to prioritize IPv6 over IPv4.

Step 3 In the **Work** pane, verify that the displayed NTP server details are correct.

NTP server

The NTP server feature enables a switch to provide time synchronization to downstream clients in accordance with NTP version 3 specifications.

- Supports IPv4 and IPv6 addressing in all tenant and management VRFs.
- Does not maintain client state.
- Responds to requests through the same VRF in which it received the requests.

NTP server operation

The NTP server feature has the following operational characteristics:

- A switch can synchronize as an IPv4 client and operate as an IPv6 server, or synchronize as an IPv6 client and operate as an IPv4 server.
- A switch can synchronize through either the out-of-band or in-band management VRF.
- The feature does not require additional contracts or IP routing-table configurations.
- When the switch is synchronized with an upstream server, it sends time information with a stratum number that is one greater than that of its system peer.

- If the switch clock is not synchronized, the switch sends time information with stratum 16, which prevents clients from synchronizing with the switch.

Note Clients can use the in-band or out-of-band IP address of the leaf switch, the bridge domain switched virtual interface (SVI) of the endpoint group (EPG), or any Layer 3 Outside (L3Out) IP address as the NTP server address. Fabric switches must not synchronize with other switches in the same fabric. Configure fabric switches to synchronize with external NTP servers.

NTP server configuration

The NTP server feature is disabled by default. Enable the feature explicitly in the configuration policy.

Enable the NTP server by using the GUI

Use this procedure to enable the NTP server while configuring NTP in the Cisco APIC GUI.

Procedure

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

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

The **Date and Time** node appears in the **Navigation** pane.

Step 3 In the **Navigation** pane, right-click **Date and Time** and choose **Create Date and Time Policy** .

The **Create Date and Time Policy** dialog box appears in the **Work** pane.

Step 4 In the **Create Date and Time Policy** dialog box, configure the policy.

a) In the **Name** field, enter a name that distinguishes this policy from the other NTP configurations in your environment.

b) For **Server State** , click **enabled** .

Enabling **Server State** allows switches to operate as NTP servers and provide time information to downstream clients.

Note
Cisco recommends that you configure an upstream NTP peer for the server. An upstream peer provides a consistent time source that the server can use for its downstream clients.

When **Server State** is enabled, the following behavior applies:

- When the switch is synchronized with an upstream server, the NTP server advertises a stratum value that is one greater than the stratum value of its system peer.
- When the switch clock is not synchronized with an upstream server, the NTP server advertises stratum 16. Clients cannot synchronize with a server that advertises stratum 16.

c) For **Master Mode** , click **enabled** .

Enabling **Master Mode** allows the designated NTP server to provide its unsynchronized local clock time to downstream clients by using the configured stratum value. For example, a leaf switch that operates as the NTP server can provide its local clock time to leaf switches that operate as clients.

Note

- **Master Mode** applies only when the server clock is not synchronized with an upstream time source.
- The default **Stratum Value** for master mode is 8.

d) In the **Stratum Value** field, enter the stratum level that the NTP server advertises to its clients. The valid range is from 1 through 14.

e) Click **Next** .

f) Click **+** to add an NTP provider.

g) In the **Create Providers** dialog box, complete the applicable fields, including **Name** , **Description** , **Minimum Polling Intervals** , and **Maximum Polling Intervals** .

- If you create multiple providers, check the **Preferred** check box for the most reliable NTP source.
- From the **Management EPG** drop-down list, choose **Out-of-Band** if all fabric nodes can reach the NTP server through out-of-band management. If you use in-band management, follow the requirements described in *In-band and out-of-band management NTP* .

h) Click **OK** .

Repeat the provider substeps for each NTP provider that you want to create.

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

The **Policy Groups** work pane appears.

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

The **Create Pod Policy Group** dialog box appears.

Step 7 In the **Create Pod Policy Group** dialog box, configure the pod policy group.

a) Enter a name for the policy group.

b) From the **Date Time Policy** drop-down list, choose the NTP policy that you created.

c) Click **Submit** .
The pod policy group is created. Alternatively, you can use the default pod policy group.

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

Step 9 In the **Work** pane, double-click the pod selector that you want to configure.

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

Step 11 Click **Submit** .

Configure the date and time format by using the GUI

Use this procedure to configure how Cisco APIC displays the system date and time.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Date and Time** .

Step 3 In the **Work** pane, configure the following settings:

- **Display Format** —Click **local** to display the local date and time, or click **utc** to display the date and time in Coordinated Universal Time (UTC). The default is **local** .
- **Time Zone** —From the drop-down list, choose the time zone for your domain. The default is **Coordinated Universal Time** .
- **Offset State** —Click **enable** or **disable** . When enabled, Cisco APIC displays the difference between the local time and the reference time. The default is **enable** .

DHCP relay policy

A DHCP relay policy enables the fabric to operate as a relay agent when DHCP clients and servers reside on different subnets.

- Supports DHCP Option 82 for relay-agent information.
- Supports DHCPv6 Option 79 for client link-layer address identification.
- Provides the DHCP server preference settings that control VRF selection for relay packets.

DHCP relay policy components and options

The following terms apply to DHCP relay policy configurations:

- *Client VRF* —The VRF instance that contains the host that initiates the DHCP request.
- *Server VRF* —The VRF instance that contains the DHCP server or provides the path to the DHCP server.
- *Client EPG* —The endpoint group (EPG) that contains the host that initiates the DHCP request.
- *Server EPG* —The EPG to which the DHCP server is connected. If the DHCP server is outside the fabric, the server EPG is the external EPG.

The **DHCP Server Preference** field determines which VRF the leaf switch uses to route DHCP relay packets:

- **None** —Routes the DHCP relay packet from the client VRF. This option is the default. Inter-VRF DHCP relay requires a shared-services contract that leaks the server VRF network into the client VRF.
- **Use Server VRF** —Routes the DHCP relay packet from the server VRF. This option does not require an inter-VRF contract or route leaking.

Figure 1: DHCP relay asymmetric forwarding scenario

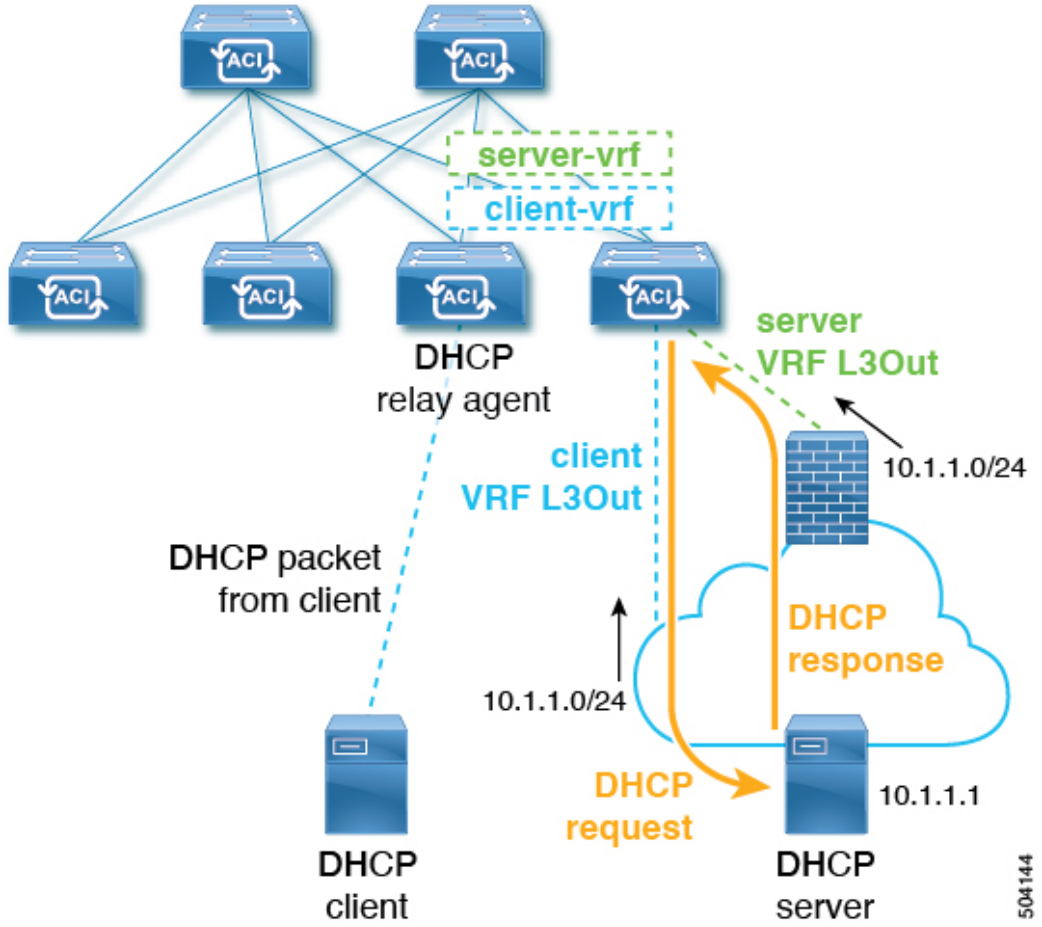

DHCPv6 Option 79 configuration

Beginning with Cisco APIC Release 5.2(4), bridge domains configured as DHCP relay agents can include DHCPv6 Option 79. When this option is enabled, the leaf switch includes the client link-layer address in the relay packet. In dual-stack deployments, this address helps associate DHCPv4 and DHCPv6 messages with the same client interface.

Guidelines and limitations for a DHCP relay policy

Use these guidelines when you configure and manage DHCP relay policies.

- A DHCP relay policy created under the `infra` or `common` tenant is not available to other tenants when you configure DHCP relay for a bridge domain. For inter-tenant DHCP relay communication, create a global DHCP relay policy as described in Create a global DHCP relay policy by using the GUI, on page 52 .
- The DHCP relay IP address is always the primary switched virtual interface (SVI) IP address.

- You cannot use a DHCP relay policy when both the clients and servers connect through Layer 3 Outside (L3Out) endpoint groups (EPGs). The clients, servers, or both must connect through a regular, non-L3Out EPG.
- When the DHCP relay provider EPG is an L3Out external EPG in a different VRF from the client bridge domain, the relay process sends a DHCP relay packet by using each node router ID and each L3Out interface IP address as a source address.
- When an L3Out SVI or floating SVI is configured as a DHCP relay agent by using a DHCP relay label, do not connect the DHCP server or provider to the leaf switches on which the L3Out SVI is deployed. For floating SVIs, this restriction applies to all anchor and non-anchor switches.
- The DHCP relay option policy supports only Option 79, and Option 79 is supported only for DHCPv6.
- An incoming unicast DHCPv4 packet uses source port `68` (`0x44`) and destination port `67` (`0x43`), with specific source and destination IP addresses.

  Because the packet uses source port 68 and destination port 67, it is subject to supervisor (SUP) copying. One copy is sent directly to the server, and another copy is punted to the DHCP process in Cisco ACI.

  When the DHCP process receives a unicast DHCPv4 request or response that has a nonzero gateway IP address that is not a local address, the process forwards the packet to its destination even if DHCP relay is not configured. As a result, the server receives an additional copy of the packet.

Configure a DHCP server policy for the APIC infrastructure by using the GUI

Before you begin

Configure Layer 2 or Layer 3 management connectivity.

Use this procedure to deploy a DHCP relay policy for an endpoint group (EPG).

Follow these guidelines and restrictions:

- The port and encapsulation that the application EPG uses must belong to a physical or virtual machine manager (VMM) domain. If the EPG is not associated with a domain, Cisco APIC deploys the EPG but raises a fault.
- Cisco APIC supports DHCP relay for IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6 addresses. DHCPv6 relay occurs only when IPv6 is enabled on the fabric interface and at least one DHCPv6 relay server is configured.
- Cisco APIC supports DHCP relay only for the primary IP address pool.
- The following guidelines and restrictions apply to the **DHCP Server Preference** field, which was introduced in Cisco APIC Release 5.2(4):
  - For an L3Out DHCP relay configuration in which the DHCP server is reachable through an L3Out and **DHCP Server Preference** is set to **Use Server VRF** , deploy an EPG, bridge domain, and bridge domain subnet on each leaf switch on which a client bridge domain is deployed if the server VRF does not already have an interface on those leaf switches.
  - When a DHCP server is behind an EPG and **DHCP Server Preference** is set to **Use Server VRF** , the IPv4 and IPv6 routes and the server bridge domain switched virtual interface (SVI) are created on the client leaf switch.
  - The **Use Server VRF** option does not support intersite DHCP traffic.

- The following restrictions apply to Option 79:
  - Option 79 is supported only for DHCPv6.
  - Option 79 is not supported in the *infra* tenant.

Procedure

	Command or Action	Purpose
Step 1	On the menu bar, choose **Tenants** > ***tenant_name*** .	
Step 2	In the **Navigation** pane, under **Tenant** *tenant_name* , expand **Policies** > **Protocol** > **DHCP** .	
Step 3	Right-click **Relay Policies** and choose **Create DHCP Relay Policy** .	
Step 4	In the **Create DHCP Relay Policy** dialog box, configure the DHCP relay policy.	The DHCP relay policy is created.
Step 5	In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels** .	
Step 6	Right-click **DHCP Relay Labels** and choose **Create DHCP Relay Label** .	
Step 7	In the **Create DHCP Relay Label** dialog box, configure the DHCP relay label.	The DHCP server is associated with the bridge domain.
Step 8	In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels** to view the DHCP server.	

DHCPv6 Option 79 REST API configuration

DHCPv6 Option 79 identifies the client link-layer address in DHCPv6 relay messages.

REST API request

Send the following POST request to add DHCPv6 Option 79 to a DHCP option policy:

```
POST URL: https://apic-ip-address/api/mo/uni.xml
                POST payload:
                <dhcpOptionPol dn="uni/tn-dhcp_client/dhcpoptpol-dhcp_option_policy"
name="dhcp_option_policy" status="">
                <dhcpOption data="" id="79" name="option_79"/>
                </dhcpOptionPol>
```

Configure a DHCP server policy for the APIC infrastructure by using the NX-OS-style CLI

Before you begin

Configure Layer 2 or Layer 3 connectivity to the DHCP server address.

Follow these guidelines when you configure the DHCP server policy:

- The port and encapsulation that the application endpoint group (EPG) uses must belong to a physical or virtual machine manager (VMM) domain. If the EPG is not associated with a domain, Cisco APIC deploys the EPG but raises a fault.
- Cisco APIC supports DHCP relay for IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6 addresses. DHCPv6 relay operates only when IPv6 is enabled on the fabric interface and at least one DHCPv6 relay server is configured.

Procedure

Enter the commands for the applicable DHCP relay deployment.

Example:

DHCP relay policy for an endpoint group

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

DHCP relay policy for a Layer 3 Outside

```
ifav28-ifc2(config)# tenant dhcpTn
      ifav28-ifc2(config-tenant)# template dhcp relay policy DhcpRelayPol
      ifav28-ifc2(config-tenant-template-dhcp-relay)# ip address 11.1.1.11 tenant dhcpTn application ap epg serverEpg
      ifav28-ifc2(config-tenant-template-dhcp-relay)# exit
      ifav28-ifc2(config-tenant)# exit
      ifav28-ifc2(config)# leaf 2001
      ifav28-ifc2(config-leaf)# interface ethernet 1/4
      ifav28-ifc2(config-leaf-if)# no switchport
      ifav28-ifc2(config-leaf-if)# vrf member tenant dhcpTn vrf v1
      ifav28-ifc2(config-leaf-if)# dhcp relay policy tenant DhcpRelayPol
      ifav28-ifc2(config-leaf-if)# exit
```

Configure a DHCP server policy for the APIC infrastructure by using the REST API

Follow these guidelines when you configure the DHCP server policy:

- Complete this task before you create a vShield domain profile.
- The port and encapsulation that the application endpoint group (EPG) uses must belong to a physical or virtual machine manager (VMM) domain. If the EPG is not associated with a domain, Cisco APIC deploys the EPG but raises a fault.

- Cisco APIC supports DHCP relay for IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6 addresses. DHCPv6 relay operates only when IPv6 is enabled on the fabric interface and at least one DHCPv6 relay server is configured.

Before you begin

Configure Layer 2 or Layer 3 management connectivity.

Procedure

Send the applicable REST API requests to configure the DHCP relay policy.

Note
The attachable access entity profile configuration deploys the relay policy on all leaf ports that connect to hypervisors. For information about configuring an attachable access entity profile, see the examples for creating VMM domain profiles.

Example:

DHCP relay policy for an EPG

```
POST URL: https://apic-ip-address/api/mo/uni.xml
            POST payload:
            <!-- api/policymgr/mo/.xml -->
            <polUni>
            <fvTenant name="infra">
            <dhcpRelayP name="DhcpRelayP" owner="tenant">
            <dhcpRsProv tDn="uni/tn-infra/ap-access/epg-default" addr="10.0.0.1"/>
            </dhcpRelayP>
            <fvBD name="default">
            <dhcpLbl name="DhcpRelayP" owner="tenant"/>
            </fvBD>
            </fvTenant>
            </polUni>
```

Example:

DHCP relay policy for a Layer 3 Outside

Note
Specify the DHCP relay label under `l3extLIfP` with the appropriate name and owner.

```
POST URL: https://apic-ip-address/api/mo/uni.xml
            POST payload:
            <polUni>
            <fvTenant name="dhcpTn">
            <l3extOut name="Out1">
            <l3extLNodeP name="NodeP">
            <l3extLIfP name="Intf1">
            <dhcpLbl name="DhcpRelayPol" owner="tenant"/>
            </l3extLIfP>
            </l3extLNodeP>
            </l3extOut>
            </fvTenant>
            </polUni>
```

Example:

Set **DHCP Server Preference** to **Use Server VRF**

```
POST URL: https://apic-ip-address/api/mo/uni.xml
            POST payload:
            <!-- api/policymgr/mo/.xml -->
            <polUni>
            <dhcpRelayP descr="" dn="uni/tn-dhcp_client/relayp-dhcp_relay_pol" status="">
            <dhcpRsProv addr="100.1.1.1/24" pref="use-server-vrf"
tDn="uni/tn-dhcp_server/ap-ap_server/epg-epg_server"/>
            </dhcpRelayP>
            </polUni>
```

Example:

Set **DHCP Server Preference** to **None**

```
POST URL: https://apic-ip-address/api/mo/uni.xml
            POST payload:
            <!-- api/policymgr/mo/.xml -->
            <polUni>
            <dhcpRelayP descr="" dn="uni/tn-dhcp_client/relayp-dhcp_relay_pol" status="">
            <dhcpRsProv addr="100.1.1.1/24" pref=""
tDn="uni/tn-dhcp_server/ap-ap_server/epg-epg_server"/>
            </dhcpRelayP>
            </polUni>
```

DNS service policy

A DNS service policy is a shared configuration object that defines the DNS providers and domains that the fabric uses for name resolution.

- Requires a DNS profile label under each applicable tenant.
- Supports DNS profile configurations for individual tenants and VRFs.
- Enables connectivity to external services by hostname.

DNS policy configuration requirements

To configure a DNS policy for the ACI fabric, complete the following tasks:

1. Configure the management endpoint group (EPG) for the DNS policy. If the management EPG is not configured, the switches do not apply the policy.
2. Create the `default` DNS profile and specify the DNS providers and domains.
3. Associate the `default` DNS profile, or another DNS profile, with a DNS label under the applicable tenant.

Note The management EPG supports only the `default` DNS policy.

You can create additional DNS profiles and apply them to specific VRFs in specific tenants by using the appropriate DNS labels. For example, if you create a DNS profile named `acme` , add a DNS label named `acme` to the applicable **Networking** > **VRF** policy in the tenant configuration.

External destinations with an in-band DNS service policy

An external destination is a network service or server that communicates with fabric components through in-band or out-of-band management connectivity.

- In-band management connectivity
- Out-of-band management connectivity
- External server placement

External destination connectivity

The connectivity requirements for an external destination depend on the source fabric component.

Table 1: External destination connectivity requirements

Source	In-band management	Out-of-band management	External server location
APIC	IP address or fully qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN **Note** The DNS policy must specify the out-of-band management endpoint group (EPG) to provide connectivity to the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN **Note** The DNS policy must specify the out-of-band management EPG to provide connectivity to the DNS server.	Directly connected to a leaf switch

The following external servers and destinations are supported:

- Call Home SMTP server
- Syslog server
- SNMP trap destination
- Statistics export destination
- Configuration export destination
- Tech support export destination

- Core export destination

Follow these guidelines when you configure connectivity to external servers:

- Connect external servers to leaf access ports.
- Use in-band connectivity for leaf switches to avoid additional management-port cabling.
- Use out-of-band management connectivity for spine switches. Connect the spine out-of-band management network to a leaf port in the in-band management virtual routing and forwarding (VRF) instance so that the spine and leaf switches can reach the same external servers.
- Specify external servers by IP address.

Dual-stack IPv4 and IPv6 DNS servers

DNS servers use resource records to associate domain names with IP addresses in IPv4 and IPv6 networks.

- A records map domain names to IPv4 addresses.
- AAAA records map domain names to IPv6 addresses.
- Dual-stack environments require specific name-server ordering to prevent DNS resolution delays.

Dual-stack DNS configuration requirements

To prevent DNS resolution delays in a dual-stack environment, configure `/etc/resolv.conf` as follows:

1. List the IPv6 DNS servers before the IPv4 DNS servers.
2. Add the `options single-request-reopen` parameter to send the IPv4 and IPv6 queries sequentially and reopen the socket for the second query.

Example resolv.conf configuration

The following `/etc/resolv.conf` example lists the IPv6 DNS servers before the IPv4 DNS servers and includes the `single-request-reopen` option:

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

Dual-stack IPv4 and IPv6 environment

A dual-stack environment supports both IPv4 and IPv6. By default, the GNU C Library (glibc) `getaddrinfo()` function returns IPv6 addresses before IPv4 addresses.

- The Linux IPv6 stack supports IPv4-mapped IPv6 addresses that use the `::ffff:0:0/96` prefix.
- An IPv6-capable application can use a single socket to accept or establish IPv4 and IPv6 connections.
- The glibc address-selection preferences for `getaddrinfo()` are configured in `/etc/gai.conf` .

Configuration and application behavior

Consider the following requirements when you configure address resolution and connectivity:

- Add `multi on` to `/etc/host.conf` to allow glibc to return multiple addresses for a hostname that has multiple entries in `/etc/hosts` . Otherwise, glibc might return only the first matching address.
- Implement a fallback mechanism for applications that are not aware of both address families. Such applications might not attempt a connection by using the alternate address family.

IPv4 or IPv6 preference in a DNS profile

A DNS profile is a configuration object that controls the preferred IP version for network address resolution.

- Supports IPv4 or IPv6 as the preferred IP version.
- Uses IPv4 as the default preference.
- Supports configuration through the GUI or REST API.

IP version preference configuration

The following example uses Postman to configure IPv6 as the preferred IP version through the REST API:

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

    </dnsProfile>
```

The `gai.conf` settings control destination address selection through a label table, a precedence table, and an IPv4 scope table. To prioritize IPv4 or IPv6, modify the precedence table. Adding a precedence entry overrides the default precedence table.

The following `gai.conf` example assigns IPv4 a higher precedence than IPv6:

```
# Generated by APIC
    label ::1/128       0
    label ::/0          1
    label 2002::/16     2
    label ::/96         3
    label ::ffff:0:0/96 4
    precedence ::1/128       50
    precedence ::/0           40
    precedence 2002::/16      30
    precedence ::/96          20
    # Prefer IPv4 connections.
    precedence ::ffff:0:0/96  100
```

Configure a DNS service policy to connect to DNS providers by using the GUI

Before you begin

Configure Layer 2 or Layer 3 management connectivity.

Procedure

Step 1 On the menu bar, choose **FABRIC** > **Fabric Policies** . In the **Navigation** pane, expand **Global Policies** > **DNS Profiles** and click the **default** DNS profile.

Step 2 In the **Work** pane, choose the appropriate management EPG, such as **default (Out-of-Band)** , from the **Management EPG** drop-down list.

Step 3 Expand **DNS Providers** and configure the DNS provider.

a) In the **Address** field, enter the DNS provider address.

b) To designate this provider as the preferred provider, check the check box in the **Preferred** column.

You can configure only one preferred DNS provider.

c) Click **Update** .

d) (Optional) To add a secondary DNS provider, enter the provider address in the **Address** field, and click **Update** .

Step 4 Expand **DNS Domains** and configure the DNS domain.

a) In the **Name** field, enter the domain name, such as `cisco.com` .

b) To designate this domain as the default domain, check the check box in the **Default** column.

You can configure only one default DNS domain.

c) Click **Update** .

d) (Optional) To add a secondary DNS domain, enter the domain name in the **Name** field, and click **Update** .

Step 5 Click **Submit** .

The DNS service policy is configured.

Step 6 On the menu bar, choose **TENANTS** > **mgmt** .

Step 7 In the **Navigation** pane, expand **Networking** > **VRF** > **oob** and click **oob** .

Step 8 In the **Work** pane, under **Properties** , enter the appropriate DNS label, such as `default` , in the **DNS labels** field, and click **Submit** .

The DNS profile label is associated with the tenant and VRF.

Configuring Custom Certificates

Guidelines for configuring custom certificates

A custom certificate provides secure communication and authentication within the fabric and requires appropriate certificate lifecycle management.

- Certificate import and renewal procedures manage the certificate lifecycle.
- Root and intermediate certificate authority (CA) certificates establish the certificate chain of trust.

- Fabric deployment policies determine how certificates are applied to controllers and fabric switches.

Custom certificate guidelines and limitations

Follow these guidelines and limitations when you configure and maintain custom certificates:

- You cannot export a private key that was used to generate a certificate signing request (CSR) on the Cisco Application Policy Infrastructure Controller (APIC). To use the same certificate on multiple servers by specifying a wildcard in the Subject Alternative Name (SAN) field, generate the private key outside the Cisco Application Centric Infrastructure (ACI) fabric and import the key.
- Download and install the public intermediate and root CA certificates before you generate a CSR. The Cisco APIC requires the root CA certificate to prevent mismatches and verify that the configured CA signed the submitted certificate.
- To reuse the public and private keys for a renewed certificate, retain the original CSR, submit that CSR for renewal, and do not delete the original key ring.
- Cisco ACI Multi-Site , VCPlugin, VRA, and SCVMM do not support certificate-based authentication.
- A Cisco APIC cluster supports only one SSL certificate.
- Before you downgrade from a later release to Cisco APIC Release 4.0(1), disable certificate-based authentication.
- To terminate a certificate-based authentication session, log out and remove the Common Access Card (CAC).
- Custom certificates configured for the Cisco APIC are deployed to the leaf and spine switches. The certificate covers a fabric node when the URL or domain name that you use to connect to the node is included in the **Subject** or **Subject Alternative Name** field.
- The Cisco APIC GUI accepts certificates with a maximum size of 4 KB.
- Self-signed SSL certificates used for HTTPS access are renewed automatically when they expire.

Configure a custom certificate for Cisco ACI HTTPS access by using the GUI

Caution Perform this task only during a maintenance window because the operation can cause downtime.

The downtime affects access by external users and systems to the Cisco Application Policy Infrastructure Controller (APIC) cluster and fabric switches. It does not affect connectivity between Cisco APIC and the switches or the fabric data plane. The NGINX processes on the controllers and switches restart during this operation. As a result, external connectivity and access to configuration, management, and troubleshooting functions are temporarily unavailable.

Before you begin

Determine which certificate authority (CA) will issue the trusted certificate so that you can configure the appropriate CA.

Procedure

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

Step 2 In the **Navigation** pane, choose **Security** .

Step 3 In the **Work** pane, choose **Certificate Authorities** > **Actions** > **Create Certificate Authority** .

Step 4 In the **Create Certificate Authority** screen, enter a name for the CA in the **Name** field.

Step 5 (Optional) In the **Description** field, enter a description of the CA.

Step 6 In the **Certificate Chain** field, paste the intermediate and root certificates for the CA that will sign the certificate signing request (CSR) for Cisco APIC .

The certificates must use Base64-encoded X.509 CER format. Place the intermediate certificate before the root CA certificate, as shown in the following example:

```
-----BEGIN CERTIFICATE-----
<Intermediate Certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root CA Certificate>
-----END CERTIFICATE-----
```

Step 7 Click **Save** .

Step 8 In the **Work** pane, choose **Key Rings** > **Actions** > **Create Key Ring** .

Key rings enable you to manage the following items:

- Private keys that are imported from an external device or generated on Cisco APIC
- A CSR generated by the private key
- A certificate issued from the CSR

Step 9 In the **Create Key Ring** dialog box, enter a name in the **Name** field.

Step 10 (Optional) In the **Description** field, enter a description of the key ring.

Step 11 In the **Certificate Authority** field, click **Select Certificate Authority** and choose the CA that you created. To create a CA, click **Create Certificate Authority** .

Step 12 In the **Private Key** field, select one of the following options:

- **Generate New Key**
- **Import Existing Key**

Step 13 If you selected **Import Existing Key** , enter the private key in the **Private Key** field.

Step 14 If you selected **Generate New Key** , select one of the following **Key Type** options:

- **RSA** —Rivest-Shamir-Adleman
- **ECC** —Elliptic-curve cryptography, also known as the Elliptic Curve Digital Signature Algorithm (ECDSA)

Step 15 In the **Certificate** field, complete one of the following actions:

- To generate a CSR through the key ring on Cisco APIC , leave the field empty.

- If you generated the private key and CSR outside Cisco APIC and already have the CA-signed certificate, paste the certificate into the field.

Step 16 If you selected **Generate New Key** , configure the key strength.

a) If you selected **RSA** as the **Key Type** , choose a value from the **Modulus** drop-down list.
b) If you selected **ECC** as the **Key Type** , select the appropriate **ECC Curve** radio button.

Step 17 In the **Create Key Ring** screen, click **Save** .

Step 18 In the **Work** pane, choose **Key Rings** > ***key_ring_name*** and expand the selected key ring.

If you did not enter a signed certificate and private key, the **Admin State** of the key ring displays **Started** . Continue with the next step to generate a CSR.

If you entered both the signed certificate and private key, the **Admin State** of the key ring displays **Completed** . Continue with the steps for associating the key ring with the management access policy.

Note
Do not delete the key ring. Deleting the key ring also deletes the associated private key used with the CSR.

Step 19 In the **Certificate Request** pane, click **Create Certificate Request** .

The **Request Certificate** window appears.

a) In the **Subject** field, enter the common name (CN) for the CSR.

You can enter a wildcard fully qualified domain name (FQDN) for the Cisco APIC controllers. For modern certificates, Cisco recommends that you enter an identifiable certificate name in this field and enter the FQDN of every controller in the **Alternate Subject Name** field, also known as the Subject Alternative Name (SAN) field. Many modern browsers require the FQDN in the SAN field.

b) In the **Alternate Subject Name** field, enter the FQDN of every Cisco APIC controller, such as `DNS:apic1.example.com,DNS:apic2.example.com,DNS:apic3.example.com` or `DNS:*.example.com` .

To match an IP address in the SAN field, enter the Cisco APIC controller IP address in the following format:

```
IP:192.168.2.1
```

You can specify DNS names, IPv4 addresses, or both. IPv6 addresses are not supported.

c) In the **Locality** field, enter the city or town where the organization is located.
d) In the **State** field, enter the state where the organization is located.
e) In the **Country** field, enter the two-letter ISO country code for the organization.
f) In the **Organization Name** and **Organization Unit Name** fields, enter the organization and organizational unit.
g) In the **Email** field, enter the email address of the organizational contact.
h) Enter a password in the **Password** field, and enter it again in the **Confirm Password** field.
i) Click **OK** .

Step 20 Verify that the **Certificate Request Settings** pane displays the CSR information that you entered.

Step 21 After the CA issues the signed certificate, import the certificate into the key ring. In the **Work** pane, choose **Key Rings** > ***key_ring_name*** .

The selected key ring and its certificate details appear.

Note

A certificate that is not signed by the CA specified in the key ring, or that contains MS-DOS line endings, is not accepted. Remove any MS-DOS line endings before you import the certificate.

After the certificate is verified, the **Admin State** changes to **Completed** . The key ring is ready for use in the HTTP policy.

Step 22 On the menu bar, choose **Fabric** > **Fabric Policies** .

Step 23 In the **Navigation** pane, choose **Policies** > **Pod** > **Management Access** > **default** .

Step 24 In the **Work** pane, choose the key ring from the **Admin Key Ring** drop-down list.

Step 25 (Optional) To enable certificate-based authentication, choose the previously created local user policy from the **Client Certificate TP** drop-down list, and click **Enabled** for **Client Certificate Authentication state** .

Step 26 Click **Submit** .

All web servers restart, the certificate becomes active, and the nondefault key ring is associated with HTTPS access.

What to do next

Monitor the certificate expiration date and renew the certificate before it expires. To retain the same key pair, preserve and resubmit the original CSR. The CSR contains the public key that corresponds to the private key in the key ring. Do not delete the existing key ring or create a new key ring. Deleting the key ring also deletes the private key stored on Cisco APIC .

Provisioning Fabric Wide System Settings

Configure APIC in-band or out-of-band connectivity preferences

The **inband** option enables Cisco APIC to access external devices through leaf switches in the Cisco ACI fabric. The **ooband** option enables Cisco APIC to access external devices through connections outside the fabric.

Before you begin

Configure the in-band and out-of-band management networks. For more information, see *Management* in the *Cisco APIC Basic Configuration Guide, Release 3.x* .

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **APIC Connectivity Preferences** .

Step 3 Choose **inband** or **ooband** .

Step 4 Click **Submit** .

Configure quota management policies

Beginning with Cisco Application Policy Infrastructure Controller (APIC) Release 2.3(1), an administrator can limit the number of managed objects that tenant administrators can configure. The limits apply globally across tenants.

Quota management prevents a tenant or group of tenants from exceeding the per-leaf or per-fabric Cisco ACI maximums or consuming a disproportionate amount of the available resources, which could affect other tenants in the fabric.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 Right-click **Quota** and choose **Create Quota Configuration** .

Step 3 In the **Class** field, choose the managed object class to limit.

Step 4 In the **Container DN** field, enter the distinguished name (DN) of the container for the selected class.

Step 5 In the **Exceed Action** field, choose **Fail Transaction Action** or **Raise Fault Action** .

Step 6 In the **Max Number** field, enter the maximum number of managed objects that can be created before the selected exceed action is applied.

Step 7 Click **Submit** .

Create an enforced BD exception list

When the enforced bridge domain (BD) feature is configured, endpoints in a subject endpoint group (EPG) can ping only the subnet gateways in the associated bridge domain.

IP addresses in the exception subnets can ping all BD gateways across all VRFs.

An L3Out loopback interface does not enforce reachability restrictions for the IP address configured on that loopback interface.

If an external BGP (eBGP) peer IP address is in a different subnet from the L3Out interface, add the peer subnet to the exception list. Otherwise, the enforced BD policy blocks the eBGP traffic because its source IP address is not in the L3Out interface subnet.

Before you begin

Create an enforced bridge domain.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **BD Enforced Exception List** .

Step 3 In the **Exception List** table, click **+** .

Step 4 Enter the IP address and network mask of the subnet that can reach any subnet gateway.

Step 5 Repeat the preceding steps for each additional subnet that you want to add to the exception list.

Step 6 Click **Submit** .

Create a BGP route reflector policy and route reflector node endpoints

Cisco ACI fabric route reflectors use Multiprotocol Border Gateway Protocol (MP-BGP) to distribute external routes within the fabric. To enable route reflection, select the spine switches that will operate as route reflectors and specify the autonomous system (AS) number. After you enable the route reflectors, you can configure connectivity to external networks.

Before you begin

- To connect external routers to the Cisco ACI fabric, configure spine switches as BGP route reflectors.
- For redundancy, configure at least two spine switches as route reflector nodes: one primary route reflector and one secondary route reflector.

Procedure

Step 1 Create the BGP route reflector policy.

a) On the menu bar, choose **System** > **System Settings** .
b) In the **Navigation** pane, choose **BGP Route Reflector** .
c) In the **Autonomous System Number** field, enter the AS number.
d) In the **Route Reflector Nodes** table, click **+** .
e) Enter the node ID of the spine switch that will operate as a route reflector, and click **Submit** .

Step 2 Create the external route reflector node endpoints.

a) In the **External Route Reflector Nodes** table, click **+** .
b) Choose the spine switch that will operate as the external route reflector node endpoint.
c) (Optional) If Cisco ACI Multi-Site manages the site, specify an intersite spine route reflector.
d) Click **Submit** .

Configure the fabric-wide control-plane MTU policy

In a Cisco ACI Multi-Pod topology, the MTU configured on the fabric-external ports must be greater than or equal to the control-plane MTU. Otherwise, the fabric-external ports might drop control-plane packets.

Note

- If the L3Out interface profile inherits the MTU from the inter-pod network (IPN), the MTU is 9150 bytes. To use an MTU of 9216 bytes across the IPN, explicitly configure that value in the L3Out interface profile at **Tenants** > ***tenant-name*** > **Networking** > **External Routed Networks** > **Create Routed Outside** > **Nodes and Interface Protocol Profiles** > **Create Node Profile** > **Create Interface Profile** .
- Regardless of the control-plane MTU setting, Cisco APIC establishes TCP connections to fabric switches with an MTU of 1496 bytes and a TCP maximum segment size (MSS) of 1456 bytes. For fabric discovery, the IPN for remote pods and remote leaf switches must support an MTU of at least 1500 bytes.

If you change both the IPN MTU and the control-plane MTU, Cisco recommends that you change the control-plane MTU first and then change the MTU on the spine switches in the remote pod. This sequence reduces the risk of losing connectivity between pods because of an MTU mismatch.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Control Plane MTU** .
Step 3 Enter the MTU value for the fabric ports.
Step 4 Click **Submit** .

Configure endpoint loop protection

The endpoint loop protection policy detects loops by monitoring how frequently an endpoint moves between ports during a specified interval.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Endpoint Controls** .
Step 3 Click the **EP Loop Protection** tab.
Step 4 In the **Administrative State** field, click **Enabled** .
Step 5 (Optional) In the **Loop Detection Interval** field, enter the interval during which the policy monitors endpoint moves. The valid range is from 30 through 300 seconds. The default is 60 seconds.
Step 6 In the **Loop Detection Multiplication Factor** field, enter the number of times that a single endpoint can move between ports during the loop detection interval. The valid range is from 1 through 255. The default is 4.
Step 7 Choose the action that the fabric takes when it detects a loop.

Choose one of the following actions:

- **BD Learn Disable**

- **Port Disable**

The default is **Port Disable** .

Step 8 Click **Submit** .

Rogue Endpoint Control Policy

Rogue endpoint control policy

The rogue endpoint control policy is a global policy that monitors endpoint movement and mitigates fabric instability caused by rogue endpoints.

- Identifies rapidly moving MAC and IP endpoints that can cause high CPU utilization and EPM or EPMC process crashes.
- Quarantines an endpoint by temporarily marking it as static to stop its movement.
- Generates host-tracking packets and faults to support endpoint relearning and corrective administrative action.

Rogue endpoint control policy behavior

The policy operates on individual endpoints and does not distinguish between local and remote endpoint moves. The policy is disabled by default.

The policy handles endpoint movement as follows:

- Identifies rapidly moving MAC and IP endpoints.
- Stops endpoint movement by temporarily marking the endpoint as static, which quarantines the endpoint.
- Before Cisco APIC Release 3.2(6), keeps the endpoint static for the **Rogue EP Detection Interval** , drops traffic to and from the rogue endpoint, and then deletes the unauthorized MAC or IP address.
- Beginning with Cisco APIC Release 3.2(6), keeps the endpoint static for the **Rogue EP Detection Interval** without dropping traffic and then deletes the unauthorized MAC or IP address.
- Generates a host-tracking packet so that the system can relearn the affected MAC or IP address.
- Raises a fault so that an administrator can take corrective action.

Rogue endpoint control policy limitations

The rogue endpoint control policy detects and mitigates endpoints that can destabilize the fabric.

- Applies only to endpoints that connect directly to the fabric.
- Requires parameter values that are within the endpoint mapper limits.
- Operates within an individual site and does not track endpoint movement between sites.

Rogue endpoint control policy limitations

The following limitations apply to the rogue endpoint control policy:

- Changes to the policy parameters do not affect existing rogue endpoints.
- When rogue endpoint control is enabled, loop detection and bridge domain move-frequency settings do not take effect.
- Disabling rogue endpoint control clears all rogue endpoints.
- The endpoint mapper (EPM) enforces value limits for rogue endpoint parameters. If a parameter value is outside the supported range, the Cisco APIC raises a fault for each mismatched parameter.
- Rogue endpoint detection supports endpoints that connect directly to the fabric. It does not support endpoints that connect to remote leaf switches.
- In a Cisco ACI Multi-Site deployment, you can use rogue endpoint control within each site to detect endpoint movement caused by server misconfigurations. The feature is not designed to detect endpoint movement between sites.
- Disable rogue endpoint control before you upgrade to or from Cisco APIC Release 4.1.

Configure the rogue endpoint control policy by using the GUI

The **Rogue EP Control** policy detects and mitigates rogue endpoints in the fabric. You can also manually clear rogue endpoints from a leaf switch.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Endpoint Controls** .

Step 3 In the **Work** pane, click the **Rogue EP Control** tab.

Step 4 Set **Administrative State** to **Enabled** .

Step 5 Configure the rogue endpoint detection and hold settings.

- **Rogue EP Detection Interval** —Enter the interval during which the policy monitors endpoint movement. The valid range is from 0 through 65535 seconds. The default is 60 seconds.
- **Rogue EP Detection Multiplication Factor** —Enter the number of endpoint moves that the policy permits during the detection interval. If the number of moves exceeds this value, the endpoint is declared rogue. The valid range is from 2 through 10. The default is 6.
- **Hold Interval (sec)** —Enter the time that the endpoint remains static after it is declared rogue. During this interval, endpoint learning is prevented. Before Cisco APIC Release 3.2(6), traffic to and from the rogue endpoint is dropped. Beginning with Cisco APIC Release 3.2(6), the traffic is not dropped. When the interval expires, the endpoint is deleted.

  Before Cisco APIC Release 5.2(3), the valid range is from 1800 through 3600 seconds. Beginning with Cisco APIC Release 5.2(3), the valid range is from 300 through 3600 seconds. The default is 1800 seconds.

Step 6 (Optional) Clear rogue endpoints from a leaf switch.

a) On the Cisco APIC menu bar, choose **Fabric** > **Inventory** .
b) In the **Navigation** pane, expand the pod and choose the leaf switch from which you want to clear rogue endpoints.
c) After the leaf switch summary appears in the **Work** pane, right-click the leaf switch in the **Navigation** pane and choose **Clear Rogue Endpoints** .
d) Click **Yes** .

Configure the rogue endpoint control policy by using the NX-OS-style CLI

The rogue endpoint control policy detects and mitigates rapidly moving endpoints that could destabilize the fabric.

Procedure

Step 1 Enter global configuration mode.

Example:

```
apic1# configure
```

Step 2 Enable the global rogue endpoint control policy.

Example:

```
apic1(config)# endpoint rogue-detect enable
```

Step 3 Set the hold interval.

The hold interval specifies how long an endpoint remains static after it is declared rogue. During this interval, endpoint learning is prevented. Before Cisco APIC Release 3.2(6), traffic to and from the rogue endpoint is dropped. Beginning with Cisco APIC Release 3.2(6), the traffic is not dropped. When the interval expires, the endpoint is deleted.

In Cisco APIC Release 5.2(2) and earlier, the valid range is from 1800 through 3600 seconds. Beginning with Cisco APIC Release 5.2(3), the valid range is from 300 through 3600 seconds. The default is 1800 seconds.

Example:

```
apic1(config)# endpoint rogue-detect hold-interval 1800
```

Step 4 Set the detection interval.

The detection interval specifies the period during which rogue endpoint control counts endpoint moves. If the number of moves exceeds the detection multiplication factor during this interval, the endpoint is declared rogue. The valid range is from 0 through 65535 seconds. The default is 60 seconds.

Example:

```
apic1(config)# endpoint rogue-detect interval 60
```

Step 5 Set the detection multiplication factor.

The detection multiplication factor specifies the permitted number of endpoint moves during the detection interval. If the number of moves exceeds this value, the endpoint is declared rogue. The valid range is from 2 through 10. The default is 6.

Example:

```
apic1(config)# endpoint rogue-detect factor 6
```

Configure the rogue endpoint control policy by using the REST API

The rogue endpoint control policy detects and mitigates rapidly moving endpoints that could destabilize the fabric.

Procedure

Step 1 Send a POST request with an XML payload similar to the following example:

Example:

```
<polUni>
                        <infraInfra>
                        <epControlP name="default" adminSt="enabled" holdIntvl="1800"
rogueEpDetectIntvl="60" rogueEpDetectMult="6"/>
                        </infraInfra>
                        </polUni>
```

The payload includes the following attributes:

- `adminSt` —Specifies the administrative state of rogue endpoint control. Set the value to `enabled` to enable the policy.
- `holdIntvl` —Specifies how long an endpoint remains static after it is declared rogue. During this interval, endpoint learning is prevented. Before Cisco APIC Release 3.2(6), traffic to and from the rogue endpoint is dropped. Beginning with Cisco APIC Release 3.2(6), the traffic is not dropped by default. When the interval expires, the endpoint is deleted.

  In Cisco APIC Release 5.2(2) and earlier, the valid range is from 1800 through 3600 seconds. Beginning with Cisco APIC Release 5.2(3), the valid range is from 300 through 3600 seconds. The default is 1800 seconds.
- `rogueEpDetectIntvl` —Specifies the period during which rogue endpoint control counts endpoint moves. If the number of moves exceeds the detection multiplication factor during this interval, the endpoint is declared rogue. The valid range is from 0 through 65535 seconds. The default is 60 seconds.
- `rogueEpDetectMult` —Specifies the permitted number of endpoint moves during the detection interval. If the number of moves exceeds this value, the endpoint is declared rogue. The valid range is from 2 through 10. The default is 6.

Step 2 (Optional) To drop traffic to and from rogue endpoints, send a POST request with an XML payload similar to the following example:

Example:

```
<infraImplicitSetPol rogueModeAction="quarantine-fault-and-drop" infraDn="uni/infra"/>
```

Rogue/COOP exception list

The rogue/COOP exception list prevents specified endpoints from being marked as rogue at the standard rogue endpoint detection threshold.

- Provides a higher movement threshold for specified endpoints.
- Supports global exceptions for specific MAC addresses across all bridge domains.
- Supports exceptions for all MAC addresses in a specific bridge domain or L3Out.

Rogue/COOP exception list behavior

An endpoint in the exception list is marked as rogue only if it moves at least 3,000 times within 10 minutes. After the endpoint is marked as rogue, it remains static for 30 seconds to prevent further learning and is then deleted.

Beginning with Cisco Application Policy Infrastructure Controller (APIC) Release 6.0(3), you can configure the following exceptions:

- A global rogue/COOP exception list for specific MAC addresses
- A rogue/COOP exception list for an L3Out
- An exception for all MAC addresses in a bridge domain or L3Out

Guidelines and limitations for the rogue/COOP exception list

Use these guidelines when you configure a rogue/COOP exception list.

- The MAC address exception list applies to Layer 2 bridge domains, which have IP routing disabled. In a Layer 3 bridge domain, if IP addresses move with the MAC addresses, the IP addresses can be marked as rogue first. The IP and MAC addresses are then quarantined.
- For a Layer 3 bridge domain, disable data-plane IP address learning per subnet for the IP addresses that you want to exclude from rogue endpoint control. For information about data-plane IP address learning per subnet, see the *Cisco APIC Layer 3 Networking Configuration Guide* .
- Before you add a MAC address to the exception list, verify that excluding the address will not cause excessive endpoint movement across the fabric or on an individual leaf switch.
- You can configure up to 100 MAC addresses across the per-bridge-domain exception lists in the fabric. Beginning with Cisco APIC Release 6.0(3), you can also exclude all MAC addresses in a bridge domain from rogue endpoint control.
- Beginning with Cisco APIC Release 6.0(3), you can configure up to 100 external L3Out MAC addresses in the exception list across the fabric. You can also exclude all MAC addresses in an L3Out switched virtual interface (SVI) bridge domain from rogue endpoint control.
- Beginning with Cisco APIC Release 6.0(3), you can configure up to 6,000 MAC addresses in the global exception list if the total number of corresponding endpoints does not exceed 6,000. A MAC address that is discovered in multiple bridge domains counts as a separate endpoint in each bridge domain. For example, a configured MAC address that is discovered in 10 bridge domains counts as 10 endpoints.
- The exception from rogue endpoint control on leaf switches applies only when rogue endpoint control is enabled. When rogue endpoint control is disabled, only COOP dampening uses the MAC address exception list.

- The rogue/COOP exception list can contain MAC addresses in a bridge domain but cannot contain IP addresses in a VRF. An IP address-only move can still cause the IP address to be marked as rogue if it meets the standard rogue endpoint control criteria.
- To prevent rogue IP address detection and marking based on data-path traffic, disable bridge domain subnet learning. Disabling bridge domain subnet learning prevents Cisco ACI from learning the location of the IP address after each move.

Configure a rogue/COOP exception list while creating a bridge domain by using the GUI

Before you begin

- Create the tenant in which you will create the bridge domain.
- Enable rogue endpoint control. For instructions, see Configure the rogue endpoint control policy by using the GUI, on page 30 .

Use this procedure to configure a rogue/COOP exception list while you create a bridge domain.

Procedure

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

Step 2 In the **Navigation** pane, choose **Networking** > **Bridge Domains** .

Step 3 Right-click **Bridge Domains** and choose **Create Bridge Domain** .

Step 4 In the **Create Bridge Domain** dialog box, configure the applicable fields in **STEP 1 > MAIN** and **STEP 2 > L3 Configurations** .

Step 5 In **STEP 3 > Advanced/Troubleshooting** , under **Rogue/Coop Exception List** , click **+** , enter the endpoint MAC address, and click **Update** .

Enter the MAC address in the `AA:BB:CC:DD:EE:FF` format.

a) Repeat this substep for each MAC address that you want to add to the exception list.

Step 6 (Optional) Beginning with Cisco Application Policy Infrastructure Controller (APIC) Release 6.0(3), check the **Enable Rogue Exception MAC for BD** check box to exclude all MAC addresses in the bridge domain from rogue endpoint control.

Step 7 Configure the remaining fields in **STEP 3 > Advanced/Troubleshooting** as required.

Step 8 Click **Finish** .

Configure the rogue/COOP exception list for an existing bridge domain by using the GUI

Use this procedure to configure the rogue/COOP exception list for an existing bridge domain.

Before you begin

- Create a tenant and bridge domain.
- Enable rogue endpoint control. For instructions, see Configure the rogue endpoint control policy by using the GUI, on page 30 .

Procedure

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

Step 2 In the **Navigation** pane, choose **Networking** > **Bridge Domains** > ***bridge_domain_name*** .

Step 3 In the **Work** pane, choose **Policy** > **Advanced/Troubleshooting** .

Step 4 Under **Rogue/Coop Exception List** , click **+** , enter the endpoint MAC address, and click **Update** .

Enter the MAC address in the `AA:BB:CC:DD:EE:FF` format.

a) Repeat this substep for each MAC address that you want to add to the exception list.

Step 5 (Optional) Beginning with Cisco Application Policy Infrastructure Controller (APIC) Release 6.0(3), check the **Enable Rogue Exception MAC for BD** check box to exclude all MAC addresses in the bridge domain from rogue endpoint control.

Step 6 Click **Submit** .

Configure a rogue endpoint control exception list for an L3Out SVI by using the GUI

A rogue endpoint control exception on an L3Out switched virtual interface (SVI) applies to all SVIs that use the same VLAN encapsulation.

Before you begin

- Create a tenant and an L3Out.
- Enable rogue endpoint control. For instructions, see Configure the rogue endpoint control policy by using the GUI, on page 30 .

Procedure

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

Step 2 In the **Navigation** pane, choose **Networking** > **L3Outs** > ***L3Out_name*** > **Logical Node Profiles** > ***node_profile_name*** > **Logical Interface Profiles** > ***interface_profile_name*** .

Step 3 In the **Work** pane, choose **Policy** > **SVI** .

Step 4 In the **SVI** table, click **+** .

Step 5 In the **Select SVI** dialog box, configure one of the following exception options:

- For **Rogue Exception MAC Group** , choose an existing group or create a group. The group identifies the MAC addresses that rogue endpoint control ignores for all SVIs that use the same VLAN encapsulation.
- Check the **Exclude all MACs from Rogue EP Control** check box. This option excludes all MAC addresses from rogue endpoint control for all SVIs that use the same VLAN encapsulation.

Step 6 Configure the remaining fields as required.

Step 7 Click **Submit** .

Rogue/COOP exception list REST API requests for an existing bridge domain

Use these REST API requests to update the rogue/COOP exception list of an existing bridge domain.

Send the following POST request to add a specific MAC address to the rogue/COOP exception list:

```
POST URL: https://apic1.myDomain.com/api/node/mo/uni/tn-tenant1.xml
                POST payload:
                <fvBD name="bd1">
                <fvRogueExceptionMac annotation="" descr="" mac="00:16:04:00:00:01"/>
                </fvBD>
```

Send the following POST request to add a wildcard that excludes all MAC addresses in the bridge domain:

```
POST URL: https://apic1.myDomain.com/api/node/mo/uni/tn-tenant1.xml
                POST payload:
                <fvBD name="bd1">
                <fvRogueExceptionMac annotation="" descr="" mac="FF:FF:FF:FF:FF:FF"/>
                </fvBD>
```

Configure fabric-wide MAC address exceptions for rogue endpoint control by using the GUI

Beginning with Cisco APIC Release 6.0(3), you can configure fabric-wide MAC address exceptions for rogue endpoint control. Rogue endpoint control ignores each specified MAC address regardless of the bridge domain or L3Out switched virtual interface (SVI) to which the address belongs.

Before you begin

Enable rogue endpoint control. For instructions, see Configure the rogue endpoint control policy by using the GUI, on page 30 .

Procedure

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

Step 2 In the **Navigation** pane, choose **Policies** > **Global** > **Fabric Wildcard Rogue Exception** .

Step 3 In the **Work** pane, click **+** in the **Wildcard Rogue Exception** table.

a) Enter a MAC address that rogue endpoint control will ignore, and click **Update** .

b) Repeat the preceding substep for each MAC address that you want rogue endpoint control to ignore.

Maximum IP address flow control

Maximum IP address flow control identifies an endpoint as rogue when the number of IP addresses learned for its MAC address exceeds the supported limit.

- The Cisco Application Centric Infrastructure (ACI) fabric supports a maximum of 4,096 IP addresses for each MAC address.
- If an endpoint exceeds this limit, the leaf switch classifies the MAC address and all its associated IP addresses as rogue.
- The feature is enabled by default and does not require configuration.

Endpoint quarantine and rogue classification

When an endpoint exceeds the maximum number of IP addresses, the following actions occur:

- The endpoint is quarantined for 1 hour. If standard rogue endpoint control is enabled, the endpoint is quarantined for the configured hold interval instead.
- The APIC raises a fault.
- The leaf switch stops learning new IP addresses for the quarantined endpoint.

The ACI fabric identifies rogue endpoints based on the following criteria:

- The frequency of endpoint movement exceeds the configured threshold.
- The number of IP addresses learned for a single MAC address exceeds 4,096.

Configure COOP

COOP

The Council of Oracle Protocol (COOP) communicates endpoint identity and location information to the spine proxy.

- Uses Zero Message Queue (ZMQ) for communication between leaf and spine switches.
- Maintains a consistent copy of endpoint address and location information on all spine switches.
- Maintains the distributed hash table (DHT) repository that maps endpoint identities to locations.

COOP endpoint dampening

Malicious or erroneous behavior can generate excessive endpoint updates and prevent COOP from processing valid updates. When rogue endpoint detection on a leaf switch does not prevent these updates, COOP invokes endpoint dampening. The spine switch instructs all leaf switches to ignore updates from the misbehaving endpoint for a specified period. The endpoint enters the `Freeze` state, and the system raises a fault.

Table 2: Endpoint dampening penalty values

Event	Penalty value	Notes
Learn a new IP address	0	A new IP address is learned.
Learn an additional IP address	2	An additional IP address is learned for an existing endpoint MAC address.
Delete an IP address	50	A remote endpoint IP address is deleted after it is learned.
Learn a deleted IP address	50	A remote endpoint IP address is learned after it is deleted.
Delete an IP address	400	A local endpoint IP address is deleted after it is learned.

Event	Penalty value	Notes
Learn a deleted IP address	400	A local endpoint IP address is learned after it is deleted.
Endpoint move	200	An endpoint moves to a different interface.
IP address move	200	An IP address moves to a different MAC address. This event has a high penalty because it causes two BGP route updates.
URIB programming	50	The status of the spine switch tunnel interface for an endpoint changes between up and down.

The penalty value is calculated for each IP address and decreases by 50 percent every five minutes. The thresholds are not configurable.

Note COOP endpoint dampening was introduced and is enabled by default beginning with Cisco Application Policy Infrastructure Controller (APIC) Release 4.2(3).

COOP authentication

COOP data-path communication gives transport priority to secured connections. COOP supports the following ZMQ authentication modes:

- **Strict mode** —Allows only MD5-authenticated ZMQ connections.
- **Compatible mode** —Allows MD5-authenticated and unauthenticated ZMQ connections for message transport.

View COOP-dampened endpoints by using the GUI

COOP endpoint dampening causes leaf switches to ignore updates from an endpoint that generates frequent, unnecessary updates.

Procedure

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

Step 2 In the **Navigation** pane, expand the pod and spine switch.

Step 3 Expand **Protocols** > **COOP** and then expand the **COOP** instance.

Step 4 Click **Endpoint Database** .

Use the **Dampened State** column to identify dampened endpoints. The column displays one of the following states:

- **Normal** —The endpoint update activity is normal.
- **Critical** —The endpoint has generated enough updates that it might enter the **Freeze** state. If the endpoint remains in the **Critical** state for more than five minutes, its state changes to **Freeze** .
- **Freeze** —The leaf switches are ignoring updates from the endpoint because it generated frequent, unnecessary updates. The system raises a fault.

View COOP-dampened endpoints by using the switch CLI

Identify the endpoints that COOP has dampened on a spine or leaf switch.

Log in to the spine or leaf switch CLI and enter the following command:

```
show coop internal info repo ep dampening
```

Clear COOP-dampened endpoints by using the GUI

Perform this procedure on every spine switch and on the source leaf switch of the endpoint. If the dampened endpoint remains in the endpoint table on the leaf switch, the leaf switch publishes the endpoint to the COOP databases on the spine switches. Otherwise, the endpoint is deleted from the spine switch COOP databases after two minutes.

Procedure

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

Step 2 In the **Navigation** pane, expand the pod and choose the spine or leaf switch.

Step 3 Right-click the switch and choose **Clear Dampened Endpoints** .

Step 4 Click **Yes** to confirm the operation.

Clear a COOP-dampened endpoint by using the switch CLI

Recover a single COOP-dampened endpoint by clearing it from all spine switches and its source leaf switch.

On every spine switch and on the source leaf switch, enter the following command:

```
clear coop internal info repo ep dampening key <bd> <mac>
```

Replace `<bd>` with the bridge domain identifier and `<mac>` with the endpoint MAC address.

Disable COOP endpoint dampening by using the REST API

Configure the fabric-wide COOP endpoint dampening state.

COOP endpoint dampening is enabled by default. You might need to disable it when you expect frequent IP address updates for a single MAC address and ignoring those updates could disrupt the network.

Send the following POST request with `disableEpDampening="true"` to disable COOP endpoint dampening:

```
POST URL: https://apic-ip-address/api/policymgr/mo/.xml
    POST payload:
```

```
<!-- api/policymgr/mo/.xml -->
<polUni>
<infraInfra>
<infraSetPol disableEpDampening="true"></infraSetPol>
</infraInfra>
</polUni>
```

All fabric nodes disable COOP endpoint dampening and recover existing endpoints that are in the `Freeze` state.

To re-enable COOP endpoint dampening, send the same request with `disableEpDampening="false"`:

```
POST URL: https://apic-ip-address/api/policymgr/mo/.xml
    POST payload:
    <polUni>
    <infraInfra>
    <infraSetPol disableEpDampening="false"></infraSetPol>
    </infraInfra>
    </polUni>
```

Configure COOP authentication by using the Cisco APIC GUI

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **COOP Group** .

Step 3 In the **Work** pane, under **Policy Property** , choose **Compatible Type** or **Strict Type** from the **Type** field.

- **Compatible Type** —Allows MD5-authenticated and unauthenticated ZMQ connections.
- **Strict Type** —Allows only MD5-authenticated ZMQ connections.

Step 4 Click **Submit** .

The COOP authentication policy is configured.

Configure COOP authentication by using the NX-OS-style CLI

Procedure

Enter the following commands to configure COOP authentication in strict mode:

Example:

```
apic1# configure
                                                    apic1(config)# coop-fabric
                                                    apic1(config-coop-fabric)# authentication
type ?
                                                    compatible  Compatible type
                                                    strict      Strict type
```

```
                                                    apic101-apic1(config-coop-fabric)#
authentication type strict
```

Configure COOP authentication by using the REST API

Procedure

Send the following POST request to configure COOP authentication in strict mode:

Strict mode allows only MD5-authenticated ZMQ connections.

Example:

```
POST URL: https://172.23.53.xx/api/node/mo/uni/fabric/pol-default.xml
                        POST payload:
                        <coopPol type="strict">
                        </coopPol>
```

Endpoint Listen Policy

Endpoint listen policy

The endpoint listen policy identifies the MAC and IP addresses of endpoints that send untagged traffic to leaf switch ports without an enforced policy.

- Detects the MAC and IP addresses of anonymous endpoints.
- Is deployed automatically on ports without an enforced policy.
- Displays detected endpoints on the **Global Endpoints** configuration screen.

Endpoint listen policy behavior

By default, the fabric drops all endpoint traffic on ports that do not have a policy. The endpoint listen policy enables the fabric to identify the source of traffic on these ports so that you can assign the detected endpoints to endpoint groups (EPGs).

Caution The endpoint listen policy is a beta feature and might not operate as intended. Use this feature at your own risk.

Configure the endpoint listen policy by using the GUI

The endpoint listen policy detects untagged traffic from anonymous endpoints on Cisco Application Centric Infrastructure (ACI) leaf switch ports without an enforced policy.

Caution The endpoint listen policy is a beta feature and might not operate as intended. Use this feature at your own risk.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Global Endpoints** .

Step 3 In the **Work** pane, check the **End Point Listen Policy** check box.

Step 4 From the **End Point Listen Encap** drop-down list, choose **VLAN** .

Step 5 In the **End Point Listen Encap** text field, enter a reserved VLAN ID. The valid range is from 1 through 4094.

Do not use the reserved VLAN encapsulation for an endpoint group (EPG).

Step 6 Click **Submit** .

Configure IP aging

When IP aging is enabled, the policy sends ARP requests for IPv4 addresses and neighbor solicitations for IPv6 addresses to verify endpoint reachability. If an endpoint does not respond, the policy ages its unused IP addresses.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Endpoint Controls** .

Step 3 Click the **IP Aging** tab.

Step 4 In the **Administrative State** field, click **Enabled** .

What to do next

Create an endpoint retention policy to specify the timer that the fabric uses to track endpoint IP addresses. Navigate to **Tenants** > ***tenant-name*** > **Policies** > **Protocol** > **End Point Retention** .

Disable remote endpoint learning

This policy applies across the fabric. After you configure the policy, the fabric deploys it to each leaf switch when the switch comes online.

Enable this policy in fabrics that include N9K-C93128TX, N9K-C9396PX, or N9K-C9396TX switches with the N9K-M12PQ uplink module. Enable the policy only after all nodes are successfully upgraded to Cisco APIC Release 2.2(2x) or later.

You might need to manually clear previously learned IP endpoints after any of the following configuration changes:

- Remote IP endpoint learning is disabled.
- The VRF is configured for ingress policy enforcement.
- The VRF contains at least one Layer 3 interface.

To manually clear previously learned IP endpoints, enter the following command on both vPC peers: `vsh -c "clear system internal epm endpoint vrf <vrf-name> remote"`

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Fabric Wide Setting** .
Step 3 Check the **Disable Remote EP Learn** check box.
Step 4 Click **Submit** .

Globally enforce subnet checks

When subnet checks are enabled, each VRF learns IP addresses only within the subnets configured for that VRF.

This policy applies across the fabric. After you configure the policy, the fabric deploys it to each leaf switch when the switch comes online.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Fabric Wide Setting** .
Step 3 Check the **Enforce Subnet Check** check box.
Step 4 Click **Submit** .

Reallocate GIPos

Cisco ACI Multi-Site uses an updated GIPo allocation scheme that provides the following benefits:

- Minimizes the number of bridge domains that share the same GIPo.

- Prevents the GIPos assigned to stretched bridge domains from overlapping with the GIPos assigned to nonstretched bridge domains.

The allocation scheme uses different pools whose sizes change based on the number of stretched and nonstretched bridge domains.

For a new Cisco ACI installation, Cisco APIC provides both benefits. When you upgrade from a release earlier than Cisco APIC Release 2.3(1), Cisco APIC retains the previous allocation scheme to avoid disrupting the fabric because existing GIPos might already be assigned to nonstretched bridge domains. Therefore, Cisco APIC cannot guarantee that the GIPos assigned to stretched and nonstretched bridge domains do not overlap.

Enabling **Reallocate Gipo** causes Cisco APIC to reallocate the GIPos by using the updated allocation scheme. This action is a one-time operation. After the reallocation, the GIPos do not overlap.

This option applies only to a Cisco ACI Multi-Site Orchestrator deployment that was upgraded from a release earlier than Cisco APIC Release 2.3(1) to Cisco APIC Release 3.0(1) or later.

This policy applies across the fabric. After you configure the policy, the fabric deploys it to each leaf switch when the switch comes online.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Fabric Wide Setting** .
Step 3 Check the **Reallocate Gipo** check box.
Step 4 Click **Submit** .

Globally enforce domain validation

When domain validation is enabled, the fabric verifies that an endpoint group (EPG) is associated with a domain when a static path is added to the EPG.

This policy applies across the fabric. After you configure the policy, the fabric deploys it to each leaf switch when the switch comes online.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .
Step 2 In the **Navigation** pane, choose **Fabric Wide Setting** .
Step 3 Check the **Enforce Domain Validation** check box.
Step 4 Click **Submit** .

Enable OpFlex client authentication

In an environment where the network cannot guarantee the identity of a GOLF or Linux OpFlex client, certificate-based authentication dynamically validates the identity of the client.

Caution Enabling certificate enforcement disables connectivity for GOLF and Linux OpFlex clients that do not support client authentication.

This policy applies across the fabric. After you configure the policy, the fabric deploys it to each leaf switch when the switch comes online.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Fabric Wide Setting** .

Step 3 Check the **OpFlex Client Authentication** check box.

Step 4 Click **Submit** .

Fabric load balancing

Fabric load balancing distributes traffic flows across available uplinks.

- Static hash load balancing distributes flows by using a 5-tuple hash.
- Dynamic load balancing (DLB) adjusts traffic distribution based on current path congestion.
- GPRS Tunneling Protocol (GTP) tunnel load balancing uses tunnel endpoint identifier (TEID) values to prevent traffic polarization in wireless networks.

Fabric load-balancing mechanisms

The fabric supports the following traffic-distribution mechanisms:

- *Static hash* —Distributes flows evenly by using a 5-tuple hash. This mechanism is effective when many flows are present but might produce an uneven distribution when a small number of flows are significantly larger than the other flows.
- *Dynamic load balancing* —Measures congestion across the available paths and adjusts traffic distribution. DLB can operate at flow or flowlet granularity. A flowlet is a burst of packets separated from another burst by a time gap. Flowlets enable more granular load balancing while limiting packet reordering.
- *DLB modes* —Aggressive mode uses a short timeout to provide fine-grained load balancing. Conservative mode uses a longer timeout to prevent packet reordering.
- *Dynamic packet prioritization* —Prioritizes short flows of fewer than 15 packets to improve application performance. You cannot enable dynamic packet prioritization (DPP) and DLB simultaneously.

Table 3: Leaf and spine switch dynamic load balancing

Traffic type	Hashing data points
Leaf and spine IP unicast	Source MAC address, destination MAC address, source IP address, destination IP address, protocol type, source Layer 4 port, destination Layer 4 port, and segment ID (VXLAN VNID) or VLAN ID
Leaf and spine Layer 2	Source MAC address, destination MAC address, and segment ID (VXLAN VNID) or VLAN ID

Note Although all Cisco Nexus 9000 Series switches support DLB in hardware, current software releases do not enable DLB on second-generation platforms with EX, FX, or FX2 suffixes.

GTP tunnel load balancing

To prevent GTP traffic polarization, the fabric uses the 32-bit TEID in the GTP header instead of a UDP port number. The TEID enables the fabric to distribute traffic across multiple links in a bundle. The feature supports GTP with IPv4 or IPv6 transport headers on physical interfaces and GTP-U on UDP port 2152.

Configure the load-balancer policy by using the Cisco APIC GUI

The load-balancer policy distributes traffic across the available uplinks. Static hash load balancing assigns each flow to an uplink based on a hash of its 5-tuple. With many flows, this mechanism typically distributes flows and bandwidth evenly. However, the distribution might be uneven when a small number of flows are significantly larger than the other flows.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings** .

Step 2 In the **Navigation** pane, choose **Load Balancer** .

Step 3 Choose a **Dynamic Load Balancing Mode** .

Dynamic load balancing (DLB) measures congestion across the available paths and assigns flows to the least-congested paths. DLB can distribute traffic across the available uplinks at flow or flowlet granularity. A flowlet is a burst of packets separated from another burst in the same flow by a time interval.

Choose **Aggressive** , **Conservative** , or **Off** . The default is **Off** .

Step 4 Set **Dynamic Packet Prioritization** to **On** or **Off** . The default is **Off** .

Dynamic packet prioritization (DPP) gives short flows a higher priority than long flows. A short flow contains fewer than approximately 15 packets. Because short flows are more sensitive to latency, DPP can improve overall application performance.

Step 5 Set **Load Balancing Mode** to **Link Failure** or **Traditional** . The default is **Traditional** .

For static and dynamic load balancing, traffic is sent only through the uplinks or paths that qualify for equal-cost multipath (ECMP). From a routing perspective, these paths have an equal and lowest cost.

Step 6 Click **Submit** .

Creating a Load Balancer Policy Using the CLI

Configure a dynamic load-balancer policy by using the CLI

Dynamic load balancing supports the **dynamic-aggressive** and **dynamic-conservative** modes. The **dynamic-aggressive** mode uses a shorter flowlet timeout interval. The **dynamic-conservative** mode uses a longer flowlet timeout interval. For more information about these commands, see the *Cisco APIC NX-OS Style CLI Command Reference* .

Configure only the mode that you want to use.

Procedure

Step 1 To enable dynamic load balancing in aggressive mode, enter the following commands:

Example:

```
apic1# conf t
                        apic1(config)# system dynamic-load-balance mode dynamic-aggressive
```

Step 2 To enable dynamic load balancing in conservative mode, enter the following commands:

Example:

```
apic1# conf t
                      apic1(config)# system dynamic-load-balance mode dynamic-conservative
```

Enable dynamic packet prioritization by using the CLI

For more information about the command in this procedure, see the *Cisco APIC NX-OS Style CLI Command Reference*.

Procedure

Enable dynamic packet prioritization.

Example:

```
apic1# conf t
                     apic1(config)# system dynamic-load-balance mode packet-prioritization
```

Configure a GTP-U load-balancer policy by using the CLI

For more information about the command in this procedure, see the *Cisco APIC NX-OS Style CLI Command Reference*.

Procedure

Enable GTP-U load balancing.

Example:

```
apic1# conf t
                        apic1(config)# ip load-sharing address source_destination gtpu
```

Configure load-balancer policies by using the REST API

This procedure configures dynamic load balancing (DLB), dynamic packet prioritization (DPP), and GPRS Tunneling Protocol User Plane (GTP-U) load balancing. For a list of supported property values, see the *Cisco APIC Management Information Model Reference*.

Procedure

Send the following XML payload to the APIC:

Example:

```
POST URL: https://apic-ip-address/api/mo/uni.xml

                        POST payload:
                        <polUni>
                        <fabricInst>
                       <lbpPol name="default" hashGtp="yes" pri="on" dlbMode="aggressive">

                        </lbpPol>
                        </fabricInst>
                        </polUni>
```

Enable Precision Time Protocol by using the GUI

PTP is a distributed protocol that synchronizes real-time clocks throughout a network. The clocks form a synchronization hierarchy in which the grandmaster clock provides the reference time for the system. The clocks exchange PTP timing messages, and each member uses the timing information to synchronize with its master clock. PTP operates within a logical scope called a PTP domain.

Before release 4.2(5), PTP was supported only within the Cisco ACI fabric. Beginning with release 4.2(5), PTP is also supported outside the Cisco ACI fabric. You can connect an external grandmaster and external clients to the front-panel ports of switches in the Cisco ACI fabric.

Note You can connect only one external grandmaster switch to a switch in the Cisco ACI fabric that is configured as a client of that grandmaster. Connecting multiple external grandmaster switches to one client switch is not supported.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings**.

Step 2 In the **Navigation** pane, choose **Precision Time Protocol**.

Step 3 In the **Work** pane, choose **Enabled** or **Disabled**.

If you disable PTP, the fabric uses NTP for time synchronization. If you enable PTP, the system automatically selects a spine switch as the master clock and synchronizes the entire site with that clock.

Step 4 Click **Submit**.

Enable the infra GIPo as the system GIPo by using the GUI

A Cisco ACI Multi-Pod deployment requires the system Global IP Outside (GIPo) address, 239.255.255.240, to be configured as a Protocol Independent Multicast bidirectional (PIM BIDIR) range on the inter-pod network (IPN). You can avoid configuring this PIM BIDIR range on the IPN devices by using the `infra` tenant GIPo as the system GIPo.

Before you begin

Upgrade all leaf and spine switches in the Cisco ACI fabric to the latest release that is compatible with the Cisco APIC release.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings**.

Step 2 For **Use Infra GIPo as System GIPo**, choose **Enabled** or **Disabled**.

The default is **Disabled**.

Step 3 Click **Submit**.

Configure a fabric port tracking policy by using the GUI

The fabric port tracking policy monitors links between leaf and spine switches. When the policy is enabled and its configured threshold is triggered, the leaf switch brings down all access interfaces on which endpoint groups (EPGs) are deployed. For more information, see the *Cisco APIC Layer 2 Networking Configuration Guide*.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings**.
Step 2 In the **Navigation** pane, choose **Port Tracking**.
Step 3 Set **Port tracking state** to **on**.
Step 4 (Optional) Change the **Daily restore timer** value.
Step 5 Configure the **Number of active spine links that triggers port tracking** value.
Step 6 Click **Submit**.

Provisioning Global Fabric Access Policies

Create a global attachable access entity profile by using the GUI

An attachable access entity profile (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies include physical interface policies that configure protocol options such as Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), and Link Aggregation Control Protocol (LACP).

An AEP is required to deploy VLAN pools on leaf switches. Encapsulation blocks and their associated VLANs can be reused across leaf switches. An AEP implicitly defines the scope of the VLAN pool within the physical infrastructure.

Consider the following AEP requirements and dependencies for network connectivity, virtual machine manager (VMM) domains, and Cisco ACI Multi-Pod configurations:

- An AEP defines the range of allowed VLANs but does not provision them. Traffic does not flow unless an endpoint group (EPG) is deployed on the port. If a VLAN pool is not defined in the AEP, the VLAN is not enabled on the leaf port even if an EPG is provisioned.
- A VLAN is provisioned on a leaf port in response to an EPG event, such as a static EPG binding on the port or a virtual machine event from an external controller such as VMware vCenter or Microsoft System Center Virtual Machine Manager (SCVMM).
- You can associate application EPGs directly with an AEP. The associated application EPGs are deployed on all ports associated with the AEP. The AEP contains a configurable generic function, `infraGeneric`, with a relationship to an EPG, `infraRsFuncToEpg`. The EPG is deployed on all interfaces that are part of the selectors associated with the AEP.

A VMM domain automatically derives physical interface policies from the interface policy groups associated with an AEP.

Before you begin

Create the tenant, VRF, application profiles, and EPGs that you want to associate with the AEP.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.
Step 2 In the **Navigation** pane, expand **Policies** > **Global**.
Step 3 Right-click **Attachable Access Entity Profile** and choose **Create Attachable Access Entity Profile**.
Step 4 Enter a name for the profile.
Step 5 In the **Domains** table, click the + icon.
Step 6 Choose an existing physical, Layer 2, Layer 3, or Fibre Channel domain, or create a domain.
Step 7 Enter the encapsulation for the domain and click **Update**.
Step 8 In the **EPG DEPLOYMENT** table, click the + icon.
Step 9 Configure the EPG deployment.

Specify the tenant, application profile, EPG, encapsulation such as `vlan-1`, primary encapsulation number, and interface mode. The interface modes are **Trunk**, **Access (802.1P)**, and **Access (Untagged)**.

Step 10 Click **Update**.
Step 11 Click **Next**.
Step 12 Choose the interfaces to associate with the AEP.
Step 13 Click **Finish**.

Configure the global QoS class policy by using the GUI

Use the global quality of service (QoS) class policy to perform the following actions:

- Preserve the Class of Service (CoS) priority so that the CoS value of IEEE 802.1p packets is retained as the packets enter and traverse the Cisco ACI fabric. IEEE 802.1p CoS preservation is supported in single-pod and Multi-Pod topologies. In a Multi-Pod topology, you can preserve the QoS priority of IEEE 802.1p traffic that enters pod 1 and exits pod 2 without preserving the CoS or Differentiated Services Code Point (DSCP) settings as the traffic traverses the inter-pod network (IPN). To preserve the CoS or DSCP settings while Multi-Pod traffic traverses the IPN, configure a DSCP policy under **Tenants** > **infra** > **Policies** > **Protocol** > **DSCP class-cos translation policy for L3 traffic**.
- Change the properties of the default QoS classes, such as **MTU**, **Queue Limit**, and **Scheduling Algorithm**.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.
Step 2 In the **Navigation** pane, expand **Policies** > **Global**.
Step 3 Click **QOS Class**.
Step 4 To enable IEEE 802.1p CoS preservation, check the **Preserve COS** check box.
Step 5 To change the default settings of a QoS class, double-click the class, configure the settings, and click **Submit**.

Create a global DHCP relay policy by using the GUI

A global DHCP relay policy identifies a DHCP server that bridge domains in different tenants can reference. A tenant-specific DHCP relay policy must be configured in each tenant, even when the provider EPG belongs to the `common` tenant. A global DHCP relay policy can be referenced from any tenant and eliminates the need to configure the policy separately in each tenant.

Procedure

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

Step 2 In the **Navigation** pane, expand **Policies** > **Global**.

Step 3 Right-click **DHCP Relay** and choose **Create DHCP Relay Policy**.

Step 4 In the **Create DHCP Relay Policy** dialog, configure the policy.

a) In the **Name** field, enter a name for the DHCP relay policy, such as `DhcpRelayP`.

The name can contain up to 64 alphanumeric characters.

b) (Optional) In the **Description** field, enter a description of the DHCP relay policy.

The description can contain up to 128 alphanumeric characters.

c) Expand **Providers**.

The **Create DHCP Provider** dialog appears.

d) In the **Create DHCP Provider** dialog, select the appropriate **EPG Type** option based on where the DHCP server is connected.

The available fields vary according to the selected EPG type:

- If you select **Application EPG**, configure the following fields in the **Application EPG** area:
  - From the **Tenant** drop-down list, choose the tenant, such as `infra`.
  - From the **Application Profile** drop-down list, choose the application profile, such as `access`.
  - From the **EPG** drop-down list, choose the EPG, such as `default`.
- If you select **L2 External Network**, configure the following fields in the **L2 External Network** area:
  - From the **Tenant** drop-down list, choose the tenant.
  - From the **L2 Out** drop-down list, choose the Layer 2 Outside.
  - From the **External Network** drop-down list, choose the external network.
- If you select **L3 External Network**, configure the following fields in the **L3 External Network** area:
  - From the **Tenant** drop-down list, choose the tenant.
  - From the **L3 Out** drop-down list, choose the Layer 3 Outside.
  - From the **External Network** drop-down list, choose the external network.

- If you select **DN**, enter the distinguished name of the target EPG.

e) In the **DHCP Server Address** field, enter the IP address of the DHCP server.

f) In the **DHCP Server Preference** field, choose **None** or **Use Server VRF**.

The **DHCP Server Preference** field is available beginning with release 5.2(4). This setting determines whether the leaf switch routes DHCP relay packets through the client VRF or the server VRF. For more information, see DHCP relay policy, on page 11.

- **None**: This default option provides the behavior used before release 5.2(4). The switch routes DHCP relay packets through the client VRF. For inter-VRF DHCP relay, a shared-services contract is required to leak the server VRF network into the client VRF.
- **Use Server VRF**: This option provides the behavior introduced in release 5.2(4). The switch routes DHCP relay packets through the server VRF, regardless of whether a contract exists between the client EPG and the server EPG or the Layer 3 external network through which the DHCP server is reachable.

  In an inter-VRF configuration, selecting **Use Server VRF** causes the server subnet route to be programmed in the server VRF on the client leaf switch for route lookup. The DHCP process on the client leaf switch then sends the DHCP relay packets through the server VRF. The server VRF must therefore be deployed with at least one IP address on every leaf switch where a client bridge domain is deployed.

g) Click **OK**.

The **Create DHCP Relay Policy** dialog appears.

h) Click **Submit**.

The global DHCP relay policy is created.

Enable the global MCP instance policy by using the GUI

Only one Mis-Cabling Protocol (MCP) instance runs in the system.

Procedure

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

Step 2 In the **Navigation** pane, expand **Policies** > **Global**.

Step 3 Click **MCP Instance Policy default**.

Step 4 Set **Admin State** to **Enabled**.

Step 5 (Optional) Configure the other policy properties as required for your fabric.

Step 6 Click **Submit**.

Create an error-disabled recovery policy by using the GUI

An error-disabled recovery policy automatically re-enables a port that was disabled because of one or more predefined error conditions.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.
Step 2 In the **Navigation** pane, expand **Policies** > **Global**.
Step 3 Click **Error Disabled Recovery Policy**.
Step 4 Double-click an event that you want to enable for automatic recovery.
Step 5 Check the event check box and click **Update**.
Step 6 (Optional) Repeat the preceding two steps to enable additional events.
Step 7 (Optional) Change the **Error disable recovery interval (sec)** value.
Step 8 Click **Submit**.

Per Port Policies

Per-port policies

A per-port policy is an implicit policy that you use to configure a leaf switch interface through the Cisco Application Policy Infrastructure Controller (APIC) GUI.

- Provides a simplified alternative to the standard policy-based configuration model.
- Creates a new set of policies for each interface instead of adding ports to existing policies.
- Uses the NX-OS-style CLI in the background to create implicit and explicit objects.

Per-port policy configuration

You can create a per-port policy only from the following location in the GUI:

Fabric > **Inventory** > ***Pod-#*** > ***leaf-switch-name*** > **Interface**

Note The **Interface** tab is in the **Work** pane. It is not the **Interfaces** folder in the **Navigation** pane.

Policy interaction behavior

Creating a port channel generates an explicit port channel policy group and an implicit override. Changes to the explicit policy group do not apply to the port until you remove the implicit policy group.

We recommend that you do not use both the CLI and the GUI to configure the same port. Use the per-port policy wizard to learn the Cisco Application Centric Infrastructure (ACI) policy model. When you are ready to use reusable policies for advanced configurations, remove the port configuration by using the same wizard.

Configure a per-port policy by using the GUI

A per-port policy provides a simplified method for configuring individual leaf switch interfaces.

Procedure

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

Step 2 In the **Navigation** pane, choose ***pod-#*** > ***leaf-switch-name*** .

Step 3 In the **Work** pane, click the **Interface** tab.

Step 4 From the **Mode** drop-down list, choose **Configuration**.

Step 5 Select one or more interface numbers.

The buttons directly below the tabs in the **Work** pane become active for the components that you can configure on the selected interfaces.

Step 6 Click the button for the component that you want to configure.

The **Work** pane displays the properties of the selected component.

Step 7 Configure the component properties as required.

Step 8 Click **Submit**.

Step 9 (Optional) Configure additional components for the selected interfaces, or select different interfaces and configure their components.

Verify a per-port policy by using the GUI

Use the Cisco Application Policy Infrastructure Controller (APIC) GUI to verify a per-port policy.

Before you begin

Configure the Cisco APIC to show hidden policies. Per-port policies are hidden by default.

Procedure

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

Step 2 In the **Navigation** pane, choose ***pod-#*** > ***leaf-switch-name*** .

Step 3 In the **Work** pane, click the **Interface** tab.

Step 4 From the **Mode** drop-down list, choose **Configuration**.

Step 5 Select an interface number.

The buttons directly below the tabs in the **Work** pane become active for the components that you can configure on the selected interface.

Step 6 Click the button for the component whose properties you want to verify.

The **Work** pane displays the properties of the selected component.

Step 7 Verify the component properties and correct any values that do not match the intended configuration.

Step 8 If you changed any values, click **Submit**. Otherwise, click **Cancel**.

Show hidden policies by using the GUI

By default, the Cisco Application Policy Infrastructure Controller (APIC) hides some policies, such as per-port policies. Change the application settings to make these policies visible.

Procedure

Step 1 In the upper-right corner of the GUI, choose **Manage My Profile** > **Settings**.

The **Application Settings** dialog appears.

Step 2 Check the **Show Hidden Policies** check box.

Step 3 Click **OK**.

Mis-cabling Protocol Interface policy

The mis-cabling protocol (MCP) was designed to handle misconfigurations that Link Layer Discovery Protocol (LLDP) and Spanning Tree Protocol (STP) are unable to detect. MCP has a Layer 2 packet that it uses, and MCP disables ports that form a loop within the Fabric. Cisco Application Centric Infrastructure (ACI) fabric leaf switches do not participate in spanning tree protocol (STP) and act as hub with respect to STP. The MCP packet is sent, and if the fabric sees that the packet comes back in, then the fabric knows that there is a loop and the fabric takes action based on that event. Faults and events are generated when this happens. MCP can be enabled globally and per-interface. By default, MCP is disabled globally and is enabled on each port. For MCP to work, it must be enabled globally, regardless of the per-interface configuration.

Starting in release 6.2(3), you can selectively suspend only the specific VLAN on which a loop is detected, rather than disabling the entire port. This minimizes the impact of loop-related outages by maintaining connectivity for unaffected VLANs on the same port.

Create a Mis-Cabling Protocol interface policy using the GUI

Before you begin

If enabling **Per-VLAN Disable**, meet the following requirements:

- Enable MCP PDU per-VLAN mode in the global MCP instance policy and the applicable MCP interface policy.

- In the Global MCP instance policy, check the box in the **Loop Protection Action** field so that **Port Disable / Per-VLAN Disable** are selected (This is a common checkbox). Go to MCP interface policy to configure **Per-VLAN Disable**.
- Configure the MCP interface to operate in strict or non-strict mode. Beginning with the 6.2(3) release, the new MCP Support: Per-VLAN Disable Loop Protection Action feature requires non-strict mode in the interface configuration.
- Beginning with the 6.2(3) release, set the **Loop Protection Action** to **Port-Disable** or **Per-VLAN Disable** in the MCP interface policy.

Follow these steps to create a mis-cabling protocol interface policy using the GUI.

Procedure

Step 1 On the mcenu bar, choose **Fabric** > **Access Policies**.

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

Step 3 In the **Work** pane, choose **Actions** > **Create Mis-cabling Protocol Interface Policy**.

Step 4 In the **Create Mis-cabling Protocol Interface Policy** dialog, configure the policy.

a) Enter a name for the policy.
b) Enter a description of the policy.
c) For **Admin State**, choose **Enable** or **Disable**.
d) In the 6.2(3) release and later, for **MCP PDU per VLAN**, choose **Enable** to enable the policy, or **Disable** to disable the policy.
e) In the 6.2(3) release and later, in the **Loop protection action**, choose **Per-VLAN Disable** or **Port Disable**.
f) Choose **Strict** or **Non-strict** as the MCP operating mode.

The following fields appear when you choose **Strict**:

- **Initial Delay Time (sec)**: Specifies the time allowed for Spanning Tree Protocol (STP) to converge on the external Layer 2 network. The default is 0 seconds when STP is disabled. When STP is enabled, use an initial delay from 45 to 60 seconds, depending on the network scale and topology.
- **Transmission Frequency (sec, msec)**: Specifies how frequently each Layer 2 interface transmits MCP packets during the grace period. The default is 500 milliseconds.
- **Grace Period (sec, msec)**: Specifies the period during which early loop detection occurs. During this period, the port transmits MCP packets aggressively for loop detection. The default is 3 seconds.

Step 5 Click **Submit**.

Restore a suspended VLAN using the GUI

This procedure applies to the MCP loop protection per-vlan disable feature in the 6.2(3) release and later.

Follow these steps to restore a suspended VLAN.

Procedure

Step 1 Identify the loop-detected VLAN(s) from the fault description.

Step 2 Use one of these methods to restore a suspended VLAN.

Follow these substeps to restore the VLAN by removing and re-adding the VLAN from the EPG:

a) Navigate to the corresponding Endpoint Group (EPG) in the APIC.
b) Remove and re-add the static port mapping for that specific VLAN to clear the suspension and restore forwarding.

Follow these substeps to restore VLAN by disabling and re-enabling the port:

a) Navigate to **Fabric** > **Inventory** > **Pod** > ***yourpodnumber*** > **Interfaces**.
b) Choose **Interface**.
c) Right click to disable, then re-enable.

Limitations for mis-cabling protocol interface policy

Theses limitations apply for the mis-cabling protocol interface policy.

- MCP is not supported on fabric ports, FEX fabric ports, and SPAN destination ports.
    - MCP is not supported on fabrix extender (FEX) host interface (HIF) ports when PDU per VLAN MCP is enabled.
- MCP Support: Per-VLAN Disable Loop Protection Action feature is not supported in mixed mode; it is Fabric scope. All leaf switches must run 16.2(3) release or later.
- If you downgrade to a release prior to 16.2(3), MCP loop protection action will be **Port Disable**.
- MCP Support: Per-VLAN Disable Loop Protection Action does not support automatic recovery. When a VLAN is suspended due to a loop, you must manually restore connectivity.
- **Per VLAN disable** is applicable only for non-strict mode.
- In strict mode, MCP supports 256 VLANs in **Maximum Number of VLANs** by default.
- If loop detection exceeds 256 (Number of Remote Ports x VLANs), then MCP Loop Protection Action will be **Port Disable** operationally. For example, if MCP detects loop by receiving BPDUs on 32 different VLANs from 8 ports, then Loop Protection Action will be **Port Disable**.