



Common Pervasive Gateway

This chapter contains the following sections:

- [Overview, on page 1](#)
- [Configuring Common Pervasive Gateway Using the GUI, on page 2](#)
- [Configuring Common Pervasive Gateway Using the NX-OS Style CLI, on page 4](#)
- [Configuring Common Pervasive Gateway Using the REST API, on page 4](#)

Overview




Note The Common Pervasive Gateway feature is being deprecated and is not actively maintained anymore.

When operating more than one Cisco ACI fabric, we highly recommend that you deploy Multi-Site instead of interconnecting multiple individual ACI fabrics to each other through leaf switches using the Common Pervasive Gateway feature. The Common Pervasive Gateway feature is currently not supported because no validations and quality assurance tests are performed in this topology for many other new features, such as L3 multicast. Hence, although Cisco ACI had the Common Pervasive Gateway feature for interconnecting ACI fabrics prior to Multi-Site, we highly recommend that you design a new ACI fabric with Multi-Site instead when there is a requirement to interconnect separate APIC domains.

This example shows how to configure Common Pervasive Gateway for IPv4 when using the Cisco APIC.

Two ACI fabrics can be configured with an IPv4 common gateway on a per bridge domain basis. Doing so enables moving one or more virtual machine (VM) or conventional hosts across the fabrics while the host retains its IP address. VM host moves across fabrics can be done automatically by the VM hypervisor. The ACI fabrics can be co-located, or provisioned across multiple sites. The Layer 2 connection between the ACI fabrics can be a local link, or can be across a bridged network. The following figure illustrates the basic common pervasive gateway topology.



Note Depending upon the topology used to interconnect two Cisco ACI fabrics, it is required that the interconnecting devices filter out the traffic source with the Virtual MAC address of the gateway switch virtual interface (SVI).

Configuring Common Pervasive Gateway Using the GUI

Before you begin

- The tenant and VRF are created.
- The bridge domain virtual MAC address and the subnet virtual IP address must be the same across all ACI fabrics for that bridge domain. Multiple bridge domains can be configured to communicate across connected ACI fabrics. The virtual MAC address and the virtual IP address can be shared across bridge domains.
- The Bridge domain that is configured to communicate across ACI fabrics must be in **flood** mode
- Only one EPG from a bridge domain (If the BD has multiple EPGs) should be configured on a border Leaf on the port which is connected to the second Fabric.
- Do not connect hosts directly to an inter-connected Layer 2 network that enables a pervasive common gateway among the two ACI fabrics.

Procedure

- Step 1** On the menu bar, click **Tenants**.
- Step 2** In the **Navigation** pane, expand the *Tenant_name* > **Networking** > **Bridge Domains**.
- Step 3** Right-click **Bridge Domains**, and click **Create Bridge Domain**.
- Step 4** In the **Create Bridge Domain** dialog box, perform the required actions to choose the appropriate attributes:

- a) In the **Main** tab, in the **Name** field, enter a name for the bridge domain, and choose the desired values for the remaining fields.
- b) In the **L3 configurations** tab, expand **Subnets**, and in the **Create Subnets** dialog box, in the **Gateway IP** field, enter the IP address.

For example, 192.0.2.1/24.

- c) In the **Treat as virtual IP address** field, check the check box.
- d) In the **Make this IP address primary** field, check the check box to specify this IP address for DHCP relay.

Checking this check box affects DHCP relay only.

- e) Click **Ok**, then click **Next** to advance to the **Advanced/Troubleshooting** tab, then click **Finish**.

Step 5 Double click the **Bridge Domain** that you just created in the **Work** pane, and perform the following action:

- a) Click the **Policy** tab, then click the **L3 Configurations** subtab.
- b) Expand **Subnets** again, and in the **Create Subnets** dialog box, to create the physical IP address in the **Gateway IP** field, use the same subnet which is configured as the virtual IP address.

For example, if you used 192.0.2.1/24 for the virtual IP address, you might use 192.0.2.2/24 here for the physical IP address.

Note The Physical IP address must be unique across the ACI fabric.

- c) Click **Submit** to complete the configuration in the **Create Subnet** window.

Step 6 In the **L3 Configurations** tab for the same bridge domain that you just created, click the **Virtual MAC Address** field, and change **Not Configured** to the appropriate value, then click **Submit**.

Note The default BD MAC address values are the same for all ACI fabrics; this configuration requires the bridge domain MAC values to be unique for each ACI fabric.

Confirm that the bridge domain MAC (pmac) values for each fabric are unique.

Note This step essentially ties the virtual MAC address that you enter in this field with the virtual IP address that you entered in the previous step. If you were to delete the virtual MAC address at some point in the future, you should also remove the check from the **Treat as virtual IP address** field for the IP address that you entered in the previous step.

Step 7 To create an L2Out EPG to extend the BD to another fabric, in the Navigation pane, right-click **External Bridged Networks** and open the **Create Bridged Outside** dialog box, and perform the following actions:

- a) In the **Name** field, enter a name for the bridged outside.
- b) In the **Bridge Domain** field, select the bridge domain already previously created.
- c) In the **Encap** field, enter the VLAN encapsulation to match the other fabric l2out encapsulation.
- d) In the **Path Type** field, select **Port**, **PC**, or **VPC** to deploy the EPG and click **Next**.
- e) To create an External EPG network click in the **Name** field, enter a name for the network and you can specify the QoS class and click **Finish** to complete Common Pervasive configuration.

Configuring Common Pervasive Gateway Using the NX-OS Style CLI

Before you begin

- The tenant, VRF, and bridge domain are created.

Procedure

Configure Common Pervasive Gateway.

Example:

```
apicl#configure
apicl(config)#tenant demo
apicl(config-tenant)#bridge-domain test
apicl(config-tenant-bd)#l2-unknown-unicast flood
apicl(config-tenant-bd)#arp flooding
apicl(config-tenant-bd)#exit

apicl(config-tenant)#interface bridge-domain test
apicl(config-tenant-interface)#multi-site-mac-address 12:34:56:78:9a:bc
apicl(config-tenant-interface)#mac-address 00:CC:CC:CC:C1:01 (Should be unique for each ACI fabric)
apicl(config-tenant-interface)#ip address 192.168.10.1/24 multi-site
apicl(config-tenant-interface)#ip address 192.168.10.254/24 (Should be unique for each ACI fabric)
```

Configuring Common Pervasive Gateway Using the REST API

Before you begin

- The tenant, VRF, and bridge domain are created.

Procedure

Configure common pervasive gateway.

In the following example REST API XML, the bolded text is relevant to configuring a common pervasive gateway.

Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- api/policymgr/mo/.xml -->
<polUni>
```

```
<fvTenant name="test">
  <fvCtx name="test"/>

  <fvBD name="test" vmac="12:34:56:78:9a:bc">
    <fvRsCtx tnFvCtxName="test"/>
    <!-- Primary address -->
    <fvSubnet ip="192.168.15.254/24" preferred="yes"/>
    <!-- Virtual address -->
    <fvSubnet ip="192.168.15.1/24" virtual="yes"/>
  </fvBD>

  <fvAp name="test">
    <fvAEPg name="web">
      <fvRsBd tnFvBDName="test"/>
      <fvRsPathAtt tDn="topology/pod-1/paths-101/pathep-[eth1/3]" encap="vlan-1002"/>
    </fvAEPg>
  </fvAp>
</fvTenant>
</polUni>
```
