



EPGs

This chapter contains the following sections:


- [About Endpoint Groups, on page 1](#)
- [Deploying an EPG on a Specific Port, on page 7](#)
- [Creating Domains, Attach Entity Profiles, and VLANs to Deploy an EPG on a Specific Port, on page 9](#)
- [Deploying EPGs to Multiple Interfaces Through Attached Entity Profiles, on page 13](#)
- [Intra-EPG Isolation, on page 15](#)
- [Configuring Intra-EPG Isolation for Cisco ACI Virtual Edge, on page 26](#)
- [Troubleshooting, on page 30](#)
- [Troubleshooting Endpoint Connectivity, on page 30](#)
- [Verifying IP-Based EPG Configurations, on page 34](#)

About Endpoint Groups

Endpoint Groups

The endpoint group (EPG) is the most important object in the policy model. The following figure shows where application EPGs are located in the management information tree (MIT) and their relation to other objects in the tenant.

Figure 1: Endpoint Groups



349872

An EPG is a managed object that is a named logical entity that contains a collection of endpoints. Endpoints are devices that are connected to the network directly or indirectly. They have an address (identity), a location, attributes (such as version or patch level), and can be physical or virtual. Knowing the address of an endpoint also enables access to all its other identity details. EPGs are fully decoupled from the physical and logical topology. Endpoint examples include servers, virtual machines, network-attached storage, or clients on the Internet. Endpoint membership in an EPG can be dynamic or static.

The ACI fabric can contain the following types of EPGs:

- Application endpoint group ($f_{v}AE_{Pg}$)
- Layer 2 external outside network instance endpoint group ($l2_{extInstP}$)
- Layer 3 external outside network instance endpoint group ($l3_{extInstP}$)
- Management endpoint groups for out-of-band ($mgmtOoB$) or in-band ($mgmtInB$) access.

EPGs contain endpoints that have common policy requirements such as security, virtual machine mobility (VMM), QoS, or Layer 4 to Layer 7 services. Rather than configure and manage endpoints individually, they are placed in an EPG and are managed as a group.

Policies apply to EPGs, never to individual endpoints. An EPG can be statically configured by an administrator in the APIC, or dynamically configured by an automated system such as vCenter or OpenStack.



Note When an EPG uses a static binding path, the encapsulation VLAN associated with this EPG must be part of a static VLAN pool. For IPv4/IPv6 dual-stack configurations, the IP address property is contained in the `fvStIp` child property of the `fvStCEP` MO. Multiple `fvStIp` objects supporting IPv4 and IPv6 addresses can be added under one `fvStCEP` object. When upgrading ACI from IPv4-only firmware to versions of firmware that support IPv6, the existing IP property is copied to an `fvStIp` MO.

Regardless of how an EPG is configured, EPG policies are applied to the endpoints they contain.

WAN router connectivity to the fabric is an example of a configuration that uses a static EPG. To configure WAN router connectivity to the fabric, an administrator configures an `l3extInstP` EPG that includes any endpoints within an associated WAN subnet. The fabric learns of the EPG endpoints through a discovery process as the endpoints progress through their connectivity life cycle. Upon learning of the endpoint, the fabric applies the `l3extInstP` EPG policies accordingly. For example, when a WAN connected client initiates a TCP session with a server within an application (`fvAEPg`) EPG, the `l3extInstP` EPG applies its policies to that client endpoint before the communication with the `fvAEPg` EPG web server begins. When the client server TCP session ends and communication between the client and server terminate, that endpoint no longer exists in the fabric.



Note If a leaf switch is configured for *static binding (leaf switches)* under an EPG, the following restrictions apply:

- The static binding cannot be overridden with a static path.
- Interfaces in that switch cannot be used for routed external network (L3out) configurations.
- Interfaces in that switch cannot be assigned IP addresses.

Virtual machine management connectivity to VMware vCenter is an example of a configuration that uses a dynamic EPG. Once the virtual machine management domain is configured in the fabric, vCenter triggers the dynamic configuration of EPGs that enable virtual machine endpoints to start up, move, and shut down as needed.

ACI Policy Configuration in EPG Shutdown

When the EPG is in shut down mode, the ACI policy configuration related to the EPG is removed from all the switches. The EPG is deleted from all the switches. While the EPG still exists in the ACI Data Store, it will be in inactive mode. In the APIC GUI you can check the box to remove the EPG from service.




Note Hosts attached to a EPG in shutdown mode cannot send or receive to/from the EPG.

Access Policies Automate Assigning VLANs to EPGs

While tenant network policies are configured separately from fabric access policies, tenant policies are not activated unless their underlying access policies are in place. Fabric access external-facing interfaces connect to external devices such as virtual machine controllers and hypervisors, hosts, routers, or Fabric Extenders (FEXs). Access policies enable an administrator to configure port channels and virtual port channels, protocols such as LLDP, CDP, or LACP, and features such as monitoring or diagnostics.

Figure 2: Association of Endpoint Groups with Access Policies



501289

In the policy model, EPGs are tightly coupled with VLANs. For traffic to flow, an EPG must be deployed on a leaf port with a VLAN in a physical, VMM, L2out, L3out, or Fibre Channel domain. For more information, see [Networking Domains](#).

In the policy model, the domain profile associated to the EPG contains the VLAN instance profile. The domain profile contains both the VLAN instance profile (VLAN pool) and the attachable Access Entity Profile (AEP), which are associated directly with application EPGs. The AEP deploys the associated application EPGs to all the ports to which it is attached, and automates the task of assigning VLANs. While a large data center could easily have thousands of active virtual machines provisioned on hundreds of VLANs, the ACI fabric can automatically assign VLAN IDs from VLAN pools. This saves a tremendous amount of time, compared with trunking down VLANs in a traditional data center.

VLAN Guidelines


Use the following guidelines to configure the VLANs where EPG traffic will flow.

- Multiple domains can share a VLAN pool, but a single domain can only use one VLAN pool.
- To deploy multiple EPGs with same VLAN encapsulation on a single leaf switch, see [Per Port VLAN, on page 4](#).

Per Port VLAN

In ACI versions prior to the v1.1 release, a given VLAN encapsulation maps to only a single EPG on a leaf switch. If there is a second EPG which has the same VLAN encapsulation on the same leaf switch, the ACI raises a fault.

Starting with the v1.1 release, you can deploy multiple EPGs with the same VLAN encapsulation on a given leaf switch (or FEX), in the Per Port VLAN configuration, similar to the following diagram:



To enable deploying multiple EPGs using the same encapsulation number, on a single leaf switch, use the following guidelines:

- EPGs must be associated with different bridge domains.
- EPGs must be deployed on different ports.
- Both the port and EPG must be associated with the same domain that is associated with a VLAN pool that contains the VLAN number.
- Ports must be configured with `portLocal` VLAN scope.

For example, with Per Port VLAN for the EPGs deployed on ports 3 and 9 in the diagram above, both using VLAN-5, port 3 and EPG1 are associated with Dom1 (pool 1) and port 9 and EPG2 are associated with Dom2 (pool 2).

Traffic coming from port 3 is associated with EPG1, and traffic coming from port 9 is associated with EPG2.

This does not apply to ports configured for Layer 3 external outside connectivity.

When an EPG has more than one physical domain with overlapping VLAN pools, avoid adding more than one domain to the AEP that is used to deploy the EPG on the ports. This avoids the risk of traffic forwarding issues.

When an EPG has only one physical domain with overlapping VLAN pool, you can associate multiple domains with single AEP.

Only ports that have the `vlanScope` set to `portLocal` allow allocation of separate (Port, VLAN) translation entries in both ingress and egress directions. For a given port with the `vlanScope` set to `portGlobal` (the default), each VLAN used by an EPG must be unique on a given leaf switch.



Note Per Port VLAN is not supported on interfaces configured with Multiple Spanning Tree (MST), which requires VLAN IDs to be unique on a single leaf switch, and the VLAN scope to be global.

Reusing VLAN Numbers Previously Used for EPGs on the Same Leaf Switch


If you have previously configured VLANs for EPGs that are deployed on a leaf switch port, and you want to reuse the same VLAN numbers for different EPGs on different ports on the same leaf switch, use a process, such as the following example, to set them up without disruption:

In this example, EPGs were previously deployed on a port associated with a domain including a VLAN pool with a range of 9-100. You want to configure EPGs using VLAN encapsulations from 9-20.

1. Configure a new VLAN pool on a different port (with a range of, for example, 9-20).
2. Configure a new physical domain that includes leaf ports that are connected to firewalls.
3. Associate the physical domain to the VLAN pool you configured in step 1.
4. Configure the VLAN Scope as `portLocal` for the leaf port.
5. Associate the new EPGs (used by the firewall in this example) to the physical domain you created in step 2.
6. Deploy the EPGs on the leaf ports.

VLAN Guidelines for EPGs Deployed on vPCs

Figure 3: VLANs for Two Legs of a vPC



When an EPG is deployed on a vPC, it must be associated with the same domain (with the same VLAN pool) that is assigned to the leaf switch ports on the two legs of the vPC.

In this diagram, EPG A is deployed on a vPC that is deployed on ports on Leaf switch 1 and Leaf switch 2. The two leaf switch ports and the EPG are all associated with the same domain, containing the same VLAN pool.

Deploying an EPG on a Specific Port

Deploying an EPG on a Specific Node or Port Using the GUI

Before you begin

The tenant where you deploy the EPG is already created.

You can create an EPG on a specific node or a specific port on a node.

Procedure

- Step 1** Log in to the Cisco APIC.
- Step 2** Choose **Tenants > tenant**.
- Step 3** In the left navigation pane, expand *tenant*, **Application Profiles**, and the *application profile*.
- Step 4** Right-click **Application EPGs** and choose **Create Application EPG**.
- Step 5** In the **Create Application EPG STEP 1 > Identity** dialog box, complete the following steps:
- In the **Name** field, enter a name for the EPG.
 - From the **Bridge Domain** drop-down list, choose a bridge domain.
 - Check the **Statically Link with Leaves/Paths** check box.
This check box allows you to specify on which port you want to deploy the EPG.
 - Click **Next**.
 - From the **Path** drop-down list, choose the static path to the destination EPG.

- Step 6** In the **Create Application EPG STEP 2 > Leaves/Paths** dialog box, from the **Physical Domain** drop-down list, choose a physical domain.

- Step 7** Complete one of the following sets of steps:

Option	Description
If you want to deploy the EPG on...	Then
A node	<ol style="list-style-type: none"> Expand the Leaves area. From the Node drop-down list, choose a node. In the Encap field, enter the appropriate VLAN. (Optional) From the Deployment Immediacy drop-down list, accept the default On Demand or choose Immediate. (Optional) From the Mode drop-down list, accept the default Trunk or choose another mode.
A port on the node	<ol style="list-style-type: none"> Expand the Paths area. From the Path drop-down list, choose the appropriate node and port.

Option	Description
	<p>c. (Optional) In the Deployment Immediacy field drop-down list, accept the default On Demand or choose Immediate.</p> <p>d. (Optional) From the Mode drop-down list, accept the default Trunk or choose another mode.</p> <p>e. In the Port Encap field, enter the secondary VLAN to be deployed.</p> <p>f. (Optional) In the Primary Encap field, enter the primary VLAN to be deployed.</p>

Step 8 Click **Update** and click **Finish**.

Step 9 In the left navigation pane, expand the EPG that you created.

Step 10 Complete one of the following actions:

- If you created the EPG on a node, click **Static Leafs**, and in the work pane view details of the static binding paths.
- If you created the EPG on a port of the node, click **Static Ports**, and in the work pane view details of the static binding paths.

Deploying an EPG on a Specific Port with APIC Using the NX-OS Style CLI

This procedure deploys an EPG on a specific port with Cisco Application Policy Infrastructure Controller (APIC) using the NX-OS-style CLI.



Note Do not mix using the GUI and the CLI when configuring per-interface on the Cisco APIC. Configurations that you perform in the GUI might only partially work in the NX-OS-style CLI.

Procedure

Step 1 Configure a VLAN domain:

Example:

```
apic1(config)# vlan-domain dom1
apic1(config-vlan)# vlan 10-100
```

Step 2 Create a tenant:

Example:

```
apic1# configure
apic1(config)# tenant t1
```

Step 3 Create a private network/VRF instance:

Example:


```
apicl(config-tenant)# vrf context ctx1
apicl(config-tenant-vrf)# exit
```

Step 4 Create a bridge domain:

Example:

```
apicl(config-tenant)# bridge-domain bd1
apicl(config-tenant-bd)# vrf member ctx1
apicl(config-tenant-bd)# exit
```

Step 5 Create an application profile and an application EPG:

Example:

```
apicl(config-tenant)# application AP1
apicl(config-tenant-app)# epg EPG1
apicl(config-tenant-app-epg)# bridge-domain member bd1
apicl(config-tenant-app-epg)# exit
apicl(config-tenant-app)# exit
apicl(config-tenant)# exit
```

Step 6 Associate the EPG with a specific port:

Example:

```
apicl(config)# leaf 1017
apicl(config-leaf)# interface ethernet 1/13
apicl(config-leaf-if)# vlan-domain member dom1
apicl(config-leaf-if)# switchport trunk allowed vlan 20 tenant t1 application AP1 epg EPG1
```

Note

The vlan-domain and vlan-domain member commands in the example are a prerequisite for deploying an EPG on a port.

Creating Domains, Attach Entity Profiles, and VLANs to Deploy an EPG on a Specific Port

Creating Domains, Attach Entity Profiles, and VLANs to Deploy an EPG on a Specific Port

This topic provides a typical example of how to create physical domains, Attach Entity Profiles (AEP), and VLANs that are mandatory to deploy an EPG on a specific port.

All endpoint groups (EPGs) require a domain. Interface policy groups must also be associated with Attach Entity Profile (AEP), and the AEP must be associated with a domain, if the AEP and EPG have to be in same domain. Based on the association of EPGs to domains and of interface policy groups to domains, the ports and VLANs that the EPG uses are validated. The following domain types associate with EPGs:

- Application EPGs

- Layer 3 external outside network instance EPGs
- Layer 2 external outside network instance EPGs
- Management EPGs for out-of-band and in-band access

The APIC checks if an EPG is associated with one or more of these types of domains. If the EPG is not associated, the system accepts the configuration but raises a fault. The deployed configuration may not function properly if the domain association is not valid. For example, if the VLAN encapsulation is not valid for use with the EPG, the deployed configuration may not function properly.



Note EPG association with the AEP without static binding does not work in a scenario when you configure the EPG as **Trunk** under the AEP with one end point under the same EPG supporting Tagging and the other end point in the same EPG does not support VLAN tagging. While associating AEP under the EPG, you can configure it as Trunk, Access (Tagged) or Access (Untagged).

Creating Domains, and VLANs to Deploy an EPG on a Specific Port Using the GUI

Before you begin

- The tenant where you deploy the EPG is already created.
- An EPG is statically deployed on a specific port.

Procedure

- Step 1** On the menu bar, choose **Fabric > Access Policies**.
- Step 2** In the **Navigation** pane, choose **Quick Start**.
- Step 3** In the **Work** pane, click **Configure Interfaces**.
- Step 4** In the **Configure Interfaces** dialog, perform the following actions:
- For **Node Type**, click **Leaf**.
 - For **Port Type**, click **Access**.
 - For **Interface Type**, choose the desired type.
 - For **Interface Aggregation Type**, choose **Individual**.
 - For **Node**, click **Select Node**, put a check in the box for the desired node, then click **OK**. You can select multiple nodes.
 - For **Interfaces For All Switches**, enter the range of desired interfaces.
 - For **Leaf Access Port Policy Group**, click **Select Leaf Access Port Policy Group**.
 - In the **Select Leaf Access Port Policy Group** dialog, click **Create Leaf Access Port Policy Group**.
 - In the **Create Leaf Access Port Policy Group** dialog, for **Link Level Policy**, click **Select Link Level Policy**.
 - Choose a link level policy and click **Select**, or click **Create Link Level Policy**, fill out the fields as desired, and click **Save**.
 - Click **Save**.

- Step 5** Create a domain and VLAN pool by performing the following actions:
- In the **Navigation** pane, expand **Physical and External Domains**.
 - Right-click **Physical Domains** and choose the appropriate **Create Physical Domain**.
 - For **Name**, enter a name for the domain.
 - For **VLAN Pool**, choose **Create VLAN Pool**, fill out the fields as desired, then click **Submit**.
 - Fill out the remaining fields as desired.
 - Click **Submit**.
- Step 6** On the menu bar, choose **Tenants > All Tenants**.
- Step 7** In the **Work** pane, double click the desired tenant.
- Step 8** In the **Navigation** pane, expand *Tenant_name* > **Application Profiles** > *profile_name* > **Application EPGs** > *EPG_name* and perform the following actions:
- Right-click **Domains (VMs and Bare-Metals)** and choose **Add Physical Domain Association**.
 - In the **Add Physical Domain Association** dialog, from the **Physical Domain Profile** drop-down list, choose the domain that you created.
 - Click **Submit**.
The AEP is associated with a specific port on a node and with a domain. The physical domain is associated with the VLAN pool and the tenant is associated with this physical domain.
- The switch profile and the interface profile are created. The policy group is created in the port block under the interface profile. The AEP is automatically created, and it is associated with the port block and with the domain. The domain is associated with the VLAN pool and the tenant is associated with the domain.

Creating AEP, Domains, and VLANs to Deploy an EPG on a Specific Port Using the NX-OS Style CLI

Before you begin

- The tenant where you deploy the EPG is already created.
- An EPG is statically deployed on a specific port.

Procedure

- Step 1** Create a VLAN domain and assign VLAN ranges:

Example:

```
apicl(config)# vlan-domain domP
apicl(config-vlan)# vlan 10
apicl(config-vlan)# vlan 25
apicl(config-vlan)# vlan 50-60
apicl(config-vlan)# exit
```

- Step 2** Create an interface policy group and assign a VLAN domain to the policy group:

Example:

```
apic1(config)# template policy-group PortGroup
apic1(config-pol-grp-if)# vlan-domain member domP
```

- Step 3** Create a leaf interface profile, assign an interface policy group to the profile, and assign the interface IDs on which the profile will be applied:

Example:

```
apic1(config)# leaf-interface-profile InterfaceProfile1
apic1(config-leaf-if-profile)# leaf-interface-group range
apic1(config-leaf-if-group)# policy-group PortGroup
apic1(config-leaf-if-group)# interface ethernet 1/11-13
apic1(config-leaf-if-profile)# exit
```

- Step 4** Create a leaf profile, assign the leaf interface profile to the leaf profile, and assign the leaf IDs on which the profile will be applied:

Example:

```
apic1(config)# leaf-profile SwitchProfile-1019
apic1(config-leaf-profile)# leaf-interface-profile InterfaceProfile1
apic1(config-leaf-profile)# leaf-group range
apic1(config-leaf-group)# leaf 1019
apic1(config-leaf-group)#
```

Validating Overlapping VLANs

This global feature prevents association of overlapping VLAN pools on a single EPG. If there are any overlapping pools allocated with any EPG in APIC, then this feature cannot be enabled (an error is displayed if there is an attempt to enable it). If no existing overlapping pools are present, then this feature can be enabled. Once enabled, when an attempt to allocate a domain on an EPG is performed, and the domain contains a VLAN pool with a range overlapping with another domain already associated to the EPG, then the configuration is blocked.

When overlapping VLAN pools exist under an EPG, then the FD VNID allocated for the EPG by each switch is non-deterministic and different switches may allocate different VNIDs. This can cause EPM sync failures between leafs within a vPC domain (causing intermittent connectivity for all endpoints within the EPG). It can also cause bridging loops if user is extending STP between the EPG, as the BPDUs will be dropped between switches due to FD VNID mismatch.

Validating Overlapping VLANs Using the GUI

This procedure provides an example of using the APIC GUI to configure overlapping VLAN validation.

Procedure

- Step 1** On the menu bar, choose **System > System Settings**.
- Step 2** In the navigation pane, choose **Fabric Wide Setting**.
- Step 3** In the work pane, locate and check **Enforce EPG VLAN Validation**.

Note

If overlapping VLAN pools already exist and this parameter is checked, the system returns an error. You must assign VLAN pools that are not overlapping to the EPGs before choosing this feature.

If this parameter is checked and an attempt is made to add an overlapping VLAN pool to an EPG, the system returns an error.

Step 4 Click **Submit**.

Deploying EPGs to Multiple Interfaces Through Attached Entity Profiles

Deploying an Application EPG through an AEP or Interface Policy Group to Multiple Ports

Through the APIC Advanced GUI and REST API, you can associate attached entity profiles directly with application EPGs. By doing so, you deploy the associated application EPGs to all those ports associated with the attached entity profile in a single configuration.

Through the APIC REST API or the NX-OS style CLI, you can deploy an application EPG to multiple ports through an Interface Policy Group.

Deploying an EPG through an AEP to Multiple Interfaces Using the APIC GUI

You can quickly associate an application with an attached entity profile to quickly deploy that EPG over all the ports associated with that attached entity profile.

Before you begin

- The target application EPG is created.
- The VLAN pools has been created containing the range of VLANs you wish to use for EPG Deployment on the AEP.
- The physical domain has been created and linked to the VLAN Pool and AEP.
- The target attached entity profile is created and is associated with the ports on which you want to deploy the application EPG.

Procedure

- Step 1** Navigate to the target attached entity profile.
- a) Open the page for the attached entity profile to use. In the GUI, click **Fabric > Access Policies > Policies > Global > Attachable Access Entity Profiles**.

b) Click the target attached entity profile to open its Attachable Access Entity Profile window.

Step 2 Click the **Show Usage** button to view the leaf switches and interfaces associated with this attached entity profile.

the application EPGs associated with this attached entity profile are deployed to all the ports on all the switches associated with this attached entity profile.

Step 3 Use the **Application EPGs** table to associate the target application EPG with this attached entity profile. Click + to add an application EPG entry. Each entry contains the following fields:

Field	Action
Application EPGs	Use the drop down to choose the associated Tenant, Application Profile, and target application EPG.
Encap	Enter the name of the VLAN over which the target application EPG will communicate.
Primary Encap	If the application EPG requires a primary VLAN, enter the name of the primary VLAN.
Mode	Use the drop down to specify the mode in which data is transmitted: <ul style="list-style-type: none"> • Trunk -- Choose if traffic from the host is tagged with a VLAN ID. • Access -- Choose if traffic from the host is tagged with an 802.1p tag. • Access Untagged -- Choose if the traffic from the host is untagged.

Step 4 Click **Submit**.

the application EPGs associated with this attached entity profile are deployed to all the ports on all the switches associated with this attached entity profile.

Deploying an EPG through an Interface Policy Group to Multiple Interfaces Using the NX-OS Style CLI

In the NX-OS CLI, an attached entity profile is not explicitly defined to associate with an EPG for rapid deployment; instead the interface policy group is defined, assigned a domain, applied to all the ports associated with a VLAN and configured to include the application EPG to be deployed over that VLAN.

Before you begin

- The target application EPG is created.
- The VLAN pools has been created containing the range of VLANs you wish to use for EPG Deployment on the AEP.
- The physical domain has been created and linked to the VLAN Pool and AEP.
- The target attached entity profile is created and is associated with the ports on which you want to deploy the application EPG.

Procedure

Step 1 Associate the target EPG with the interface policy group.

The sample command sequence specifies an interface policy group **pg3** associated with VLAN domain, **domain1**, and with VLAN **1261**. The application EPG, **epg47** is deployed to all interfaces associated with this policy group.

Example:

```
apic1# configure terminal
apic1(config)# template policy-group pg3
apic1(config-pol-grp-if)# vlan-domain member domain1
apic1(config-pol-grp-if)# switchport trunk allowed vlan 1261 tenant tn10 application pod1-AP
epg epg47
```

Step 2 Check the target ports to ensure deployment of the policies of the interface policy group associated with application EPG.

The output of the sample **show** command sequence indicates that policy group **pg3** is deployed on Ethernet port **1/20** on leaf switch **1017**.

Example:

```
apic1# show run leaf 1017 int eth 1/20
# Command: show running-config leaf 1017 int eth 1/20
# Time: Mon Jun 27 22:12:10 2016
leaf 1017
  interface ethernet 1/20
    policy-group pg3
  exit
exit
ifav28-ifc1#
```

Intra-EPG Isolation

Intra-EPG Endpoint Isolation

Intra-EPG endpoint isolation policies provide full isolation for virtual or physical endpoints; no communication is allowed between endpoints in an EPG that is operating with isolation enforced. Isolation enforced EPGs reduce the number of EPG encapsulations required when many clients access a common service but are not allowed to communicate with each other.

An EPG is isolation enforced for all Cisco Application Centric Infrastructure (ACI) network domains or none. While the Cisco ACI fabric implements isolation directly to connected endpoints, switches connected to the fabric are made aware of isolation rules according to a primary VLAN (PVLAN) tag.



Note If an EPG is configured with intra-EPG endpoint isolation enforced, these restrictions apply:

- All Layer 2 endpoint communication across an isolation enforced EPG is dropped within a bridge domain.
- All Layer 3 endpoint communication across an isolation enforced EPG is dropped within the same subnet.
- Preserving QoS CoS priority settings is not supported when traffic is flowing from an EPG with isolation enforced to an EPG without isolation enforced.

BPDUs are not forwarded through EPGs with intra-EPG isolation enabled. Therefore, when you connect an external Layer 2 network that runs spanning tree in a VLAN that maps to an isolated EPG on Cisco ACI, Cisco ACI might prevent spanning tree in the external network from detecting a Layer 2 loop. You can avoid this issue by ensuring that there is only a single logical link between Cisco ACI and the external network in these VLANs.

Intra-EPG Isolation for Bare Metal Servers

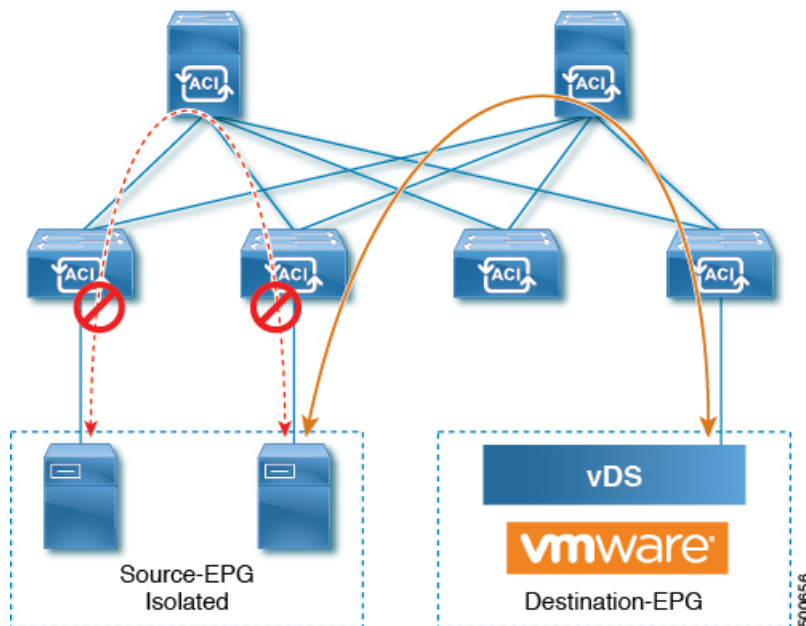
Intra-EPG Isolation for Bare Metal Servers

Intra-EPG endpoint isolation policies can be applied to directly connected endpoints such as bare metal servers.

Examples use cases include the following:

- Backup clients have the same communication requirements for accessing the backup service, but they don't need to communicate with each other.
- Servers behind a load balancer have the same communication requirements, but isolating them from each other protects against a server that is compromised or infected.

Figure 4: Intra-EPG Isolation for Bare Metal Servers



Bare metal EPG isolation is enforced at the leaf switch. Bare metal servers use VLAN encapsulation. All unicast, multicast and broadcast traffic is dropped (denied) within isolation enforced EPGs. ACI bridge-domains can have a mix of isolated and regular EPGs. Each Isolated EPG can have multiple VLANs where intra-vlan traffic is denied.

Configuring Intra-EPG Isolation for Bare Metal Servers Using the GUI

The port the EPG uses must be associated with a bare metal server interface in the physical domain that is used to connect the bare metal servers directly to leaf switches.

SUMMARY STEPS

1. In a tenant, right click on an **Application Profile**, and open the **Create Application EPG** dialog box to perform the following actions:
2. In the **Leaves/Paths** dialog box, perform the following actions:

DETAILED STEPS

Procedure

Step 1 In a tenant, right click on an **Application Profile**, and open the **Create Application EPG** dialog box to perform the following actions:

- a) In the **Name** field, add the EPG name (intra_EPG-deny).
- b) For **Intra EPG Isolation**, click **Enforced**.
- c) In the **Bridge Domain** field, choose the bridge domain from the drop-down list (bd1).
- d) Check the **Statically Link with Leaves/Paths** check box.
- e) Click **Next**.

Step 2 In the **Leaves/Paths** dialog box, perform the following actions:

- a) In the **Path** section, choose a path from the drop-down list (Node-107/eth1/16) in Trunk Mode.
Specify the **Port Encap** (vlan-102) for the secondary VLAN.

Note

If the bare metal server is directly connected to a leaf switch, only the Port Encap secondary VLAN is specified.

Specify the **Primary Encap** (vlan-103) for the primary VLAN.

- b) Click **Update**.
 - c) Click **Finish**.
-

Configuring Intra-EPG Isolation for Bare Metal Servers Using the NX-OS Style CLI

SUMMARY STEPS

1. In the CLI, create an intra-EPG isolation EPG:
2. Verify the configuration:

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	<p>In the CLI, create an intra-EPG isolation EPG:</p> <p>Example:</p> <p>The VMM case is below.</p> <pre> ifav19-ifc1(config)# tenant Test_Isolation ifav19-ifc1(config-tenant)# application PVLAN ifav19-ifc1(config-tenant-app)# epg EPG1 ifav19-ifc1(config-tenant-app-epg)# show running-config # Command: show running-config tenant Test_Isolation application PVLAN epg EPG1 tenant Test_Isolation application PVLAN epg EPG1 bridge-domain member BD1 contract consumer bare-metal contract consumer default contract provider Isolate_EPG isolation enforce <---- This enables EPG isolation mode. exit exit exit exit ifav19-ifc1(config)# leaf ifav19-leaf3 ifav19-ifc1(config-leaf)# interface ethernet 1/16 ifav19-ifc1(config-leaf-if)# show running-config ifav19-ifc1(config-leaf-if)# switchport trunk native vlan 101 tenant Test_Isolation application PVLAN epg StaticEPG primary-vlan 100 exit </pre>	
Step 2	<p>Verify the configuration:</p> <p>Example:</p> <pre> show epg StaticEPG detail Application EPg Data: Tenant : Test_Isolation Application : PVLAN AEPg : StaticEPG BD : BD1 uSeg EPG : no Intra EPG Isolation : enforced Vlan Domains : phys Consumed Contracts : bare-metal Provided Contracts : default,Isolate_EPG Denied Contracts : Qos Class : unspecified Tag List : VMM Domains: Domain Type Deployment Immediacy Resolution Immediacy State Encap Primary Encap ----- ----- </pre>	

	Command or Action	Purpose
	<pre> DVS1 VMware On Demand immediate formed auto auto Static Leaves: Node Encap Deployment Immediacy Mode Modification Time ----- ----- Static Paths: Node Interface Encap Modification Time ----- ----- 1018 eth101/1/1 vlan-100 2016-02-11T18:39:02.337-08:00 1019 eth1/16 vlan-101 2016-02-11T18:39:02.337-08:00 Static Endpoints: Node Interface Encap End Point MAC End Point IP Address Modification Time ----- ----- </pre>	

Intra-EPG Isolation for VMWare vDS

Intra-EPG Isolation for VMware VDS or Microsoft Hyper-V Virtual Switch

Intra-EPG Isolation is an option to prevent physical or virtual endpoint devices that are in the same base EPG or microsegmented (uSeg) EPG from communicating with each other. By default, endpoint devices included in the same EPG are allowed to communicate with one another. However, conditions exist in which total isolation of the endpoint devices from one another within an EPG is desirable. For example, you may want to enforce intra-EPG isolation if the endpoint VMs in the same EPG belong to multiple tenants, or to prevent the possible spread of a virus.

A Cisco Application Centric Infrastructure (ACI) virtual machine manager (VMM) domain creates an isolated PVLAN port group at the VMware VDS or Microsoft Hyper-V Virtual Switch for each EPG that has intra-EPG isolation enabled. A fabric administrator specifies primary encapsulation or the fabric dynamically specifies primary encapsulation at the time of EPG-to-VMM domain association. When the fabric administrator selects the VLAN-pri and VLAN-sec values statically, the VMM domain validates that the VLAN-pri and VLAN-sec are part of a static block in the domain pool.

Primary encapsulation is defined per EPG VLAN. In order to use primary encapsulation for Intra-EPG isolation, you must deploy it in one of the following ways:

- Segregate primary and secondary VLAN defined ports on different switches. EPG VLAN is created per switch. If you have port encapsulation, and only static ports on a switch for an EPG, primary encapsulation is not associated.
- Use a different encapsulation for static ports that use only port encapsulation. This creates a second EPG VLAN that does not have primary encapsulation associated with it.

In the example below, consider egress traffic on two interfaces (Eth1/1, Eth1/3) with primary VLAN-1103. Eth1/1 port encap was changed to VLAN-1132 (from VLAN-1130), so that it does not share the secondary VLAN with Eth1/3.

Port encap with VLAN-1130 on Eth1/1

Eth1/1: Port Encap only VLAN-1130

Eth1/6: Primary VLAN-1103 and Secondary VLAN-1130

```
fab2-leaf3# show vlan id 53 ext
```

VLAN Name	Encap	Ports
53 JT:jt-ap:EPG1-1	vlan-1130	Eth1/1, Eth1/3

```
module-1# show sys int eltmc info vlan access_encap_vlan 1130
```

```

vlan_id:          53  :::      isEpg:          1
bd_vlan_id:       52  :::      hwEpgId:         11278
srcpolicyincom:   0   :::      data_mode:       0
accencaptype:     0   :::      fabencaptype:    2
accencapval:    1130 :::      fabencapval:     12192
sclass:           49154 :::      sqlabel:         12
sclassprio:       1   :::      floodmetptr:     13
maclearnen:       1   :::      iplearnen:       1
sclasslrnen:      1   :::      bypselfwdchk:    0
qosusetc:         0   :::      qosuseexp:       0
isolated:         1   :::      primary_encap: 1103
proxy_arp:        0   :::      qinq core:       0
ivxlan_dl:        0   :::      dtag_mode:       0
is_service_epg:   0

```

Port encap changed to VLAN-1132 on Eth1/1

```
fab2-leaf3# show vlan id 62 ext
```

VLAN Name	Encap	Ports
62 JT:jt-ap:EPG1-1	vlan-1132	Eth1/1

```
module-1# show sys int eltmc info vlan access_encap_vlan 1132
```

```

[SDK Info]:
vlan_id:          62  :::      isEpg:          1
bd_vlan_id:       52  :::      hwEpgId:         11289
srcpolicyincom:   0   :::      data_mode:       0
accencaptype:     0   :::      fabencaptype:    2
accencapval:    1132 :::      fabencapval:     11224
sclass:           49154 :::      sqlabel:         12
sclassprio:       1   :::      floodmetptr:     13
maclearnen:       1   :::      iplearnen:       1
sclasslrnen:      1   :::      bypselfwdchk:    0
qosusetc:         0   :::      qosuseexp:       0
isolated:         1   :::      primary_encap: 0
proxy_arp:        0   :::      qinq core:       0

```

```

        ivxlan_dl:          0   :::          dtag_mode:          0
        is_service_epg:     0
fab2-leaf3# show vlan id 53 ext

VLAN Name                Encap                Ports
-----
53   JT:jt-ap:EPG1-1       vlan-1130            Eth1/3

module-1# show sys int eltmc info vlan access_encap_vlan 1130
[SDK Info]:
    vlan_id:          53   :::          isEpg:          1
    bd_vlan_id:       52   :::          hwEpgId:       11278
    srcpolicyincom:   0   :::          data_mode:     0
    accencaptype:     0   :::          fabencaptype:  2
    accencapval:     1130  :::          fabencapval:   12192
    sclass:           49154  :::          sglabel:       12
    sclassprio:       1   :::          floodmetptr:   13
    maclearnen:       1   :::          iplearnen:     1
    sclasslrnen:      1   :::          bypselffwdchk: 0
    qosusetc:         0   :::          qosuseexp:     0
    isolated:         1   :::          primary_encap: 1103
    proxy_arp:        0   :::          qinq core:     0
    ivxlan_dl:        0   :::          dtag_mode:     0

```

**Note**


- When intra-EPG isolation is not enforced, the VLAN-pri value is ignored even if it is specified in the configuration.
- A VMware distributed virtual switch (DVS) domain with EDM UCSM integration may fail. The domain fails if you configure intra-EPG isolation on the endpoint group (EPG) attached to the domain and you use UCSM Mini 6324, which does not support private VLANs.

BPDUs are not forwarded through EPGs with intra-EPG isolation enabled. Therefore, when you connect an external Layer 2 network that runs spanning tree in a VLAN that maps to an isolated EPG on Cisco ACI, Cisco ACI might prevent spanning tree in the external network from detecting a Layer 2 loop. You can avoid this issue by ensuring that there is only a single logical link between Cisco ACI and the external network in these VLANs.

VLAN-pri/VLAN-sec pairs for the VMware VDS or Microsoft Hyper-V Virtual Switch are selected per VMM domain during the EPG-to-domain association. The port group created for the intra-EPG isolation EPGs uses the VLAN-sec tagged with type set to `PVLAN`. The VMware VDS or the Microsoft Hyper-V Virtual Switch and fabric swap the VLAN-pri/VLAN-sec encapsulation:

- Communication from the Cisco ACI fabric to the VMware VDS or Microsoft Hyper-V Virtual Switch uses VLAN-pri.
- Communication from the VMware VDS or Microsoft Hyper-V Virtual Switch to the Cisco ACI fabric uses VLAN-sec.

Figure 5: Intra-EPG Isolation for VMware VDS or Microsoft Hyper-V Virtual Switch



Note these details regarding this illustration:

1. EPG-DB sends VLAN traffic to the Cisco ACI leaf switch. The Cisco ACI egress leaf switch encapsulates traffic with a primary VLAN (PVLAN) tag and forwards it to the Web-EPG endpoint.
2. The VMware VDS or Microsoft Hyper-V Virtual Switch sends traffic to the Cisco ACI leaf switch using VLAN-sec. The Cisco ACI leaf switch drops all intra-EPG traffic because isolation is enforced for all intra VLAN-sec traffic within the Web-EPG.
3. The VMware VDS or Microsoft Hyper-V Virtual Switch VLAN-sec uplink to the Cisco ACI leaf switch is in isolated trunk mode. The Cisco ACI leaf switch uses VLAN-pri for downlink traffic to the VMware VDS or Microsoft Hyper-V Virtual Switch.
4. The PVLAN map is configured in the VMware VDS or Microsoft Hyper-V Virtual Switch and Cisco ACI leaf switches. VM traffic from WEB-EPG is encapsulated in VLAN-sec. The VMware VDS or Microsoft Hyper-V Virtual Switch denies local intra-WEB EPG VM traffic according to the PVLAN tag. All intra-ESXi host or Microsoft Hyper-V host VM traffic is sent to the Cisco ACI leaf switch using VLAN-Sec.

Configuring Intra-EPG Isolation for VMware VDS or Microsoft Hyper-V Virtual Switch using the GUI

SUMMARY STEPS

1. Log into Cisco APIC.
2. Choose **Tenants** > *tenant*.
3. In the left navigation pane expand the **Application Profiles** folder and appropriate application profile.
4. Right-click the **Application EPGs** folder and then choose **Create Application EPG**.
5. In the **Create Application EPG** dialog box, complete the following steps:
6. Click **Update** and click **Finish**.

DETAILED STEPS

Procedure

-
- Step 1** Log into Cisco APIC.
- Step 2** Choose **Tenants** > *tenant*.
- Step 3** In the left navigation pane expand the **Application Profiles** folder and appropriate application profile.
- Step 4** Right-click the **Application EPGs** folder and then choose **Create Application EPG**.
- Step 5** In the **Create Application EPG** dialog box, complete the following steps:
- a) In the **Name** field, add the EPG name.
 - b) In the **Intra EPG Isolation** area, click **Enforced**.
 - c) In the **Bridge Domain** field, choose the bridge domain from the drop-down list.
 - d) Associate the EPG with a bare metal/physical domain interface or with a VM Domain.
 - For the VM Domain case, check the **Associate to VM Domain Profiles** check box.
 - For the bare metal case, check the **Statically Link with Leaves/Paths** check box.
 - e) Click **Next**.
 - f) In the **Associated VM Domain Profiles** area, click the + icon.
 - g) From the **Domain Profile** drop-down list, choose the desired VMM domain.
- For the static case, in the **Port Encap (or Secondary VLAN for Micro-Seg)** field, specify the secondary VLAN, and in the **Primary VLAN for Micro-Seg** field, specify the primary VLAN. If the Encap fields are left blank, values will be allocated dynamically.
- Note**
For the static case, a static VLAN must be available in the VLAN pool.
- Step 6** Click **Update** and click **Finish**.
-

Configuring Intra-EPG Isolation for VMware VDS or Microsoft Hyper-V Virtual Switch using the NX-OS Style CLI

SUMMARY STEPS

1. In the CLI, create an intra-EPG isolation EPG:
2. Verify the configuration:

DETAILED STEPS

Procedure

Step 1 In the CLI, create an intra-EPG isolation EPG:

Example:

The following example is for VMware VDS:

```
apic1(config)# tenant Test_Isolation
apic1(config-tenant)# application PVLAN
apic1(config-tenant-app)# epg EPG1
apic1(config-tenant-app-epg)# show running-config
# Command: show running-config tenant Tenant_VMM application Web epg intraEPGDeny
tenant Tenant_VMM
  application Web
    epg intraEPGDeny
      bridge-domain member VMM_BD
      vmware-domain member PVLAN encap vlan-2001 primary-encap vlan-2002 push on-demand
      vmware-domain member mininet
      exit
    isolation enforce
    exit
  exit
exit
apic1(config-tenant-app-epg)#
```

Example:

The following example is for Microsoft Hyper-V Virtual Switch:

```
apic1(config)# tenant Test_Isolation
apic1(config-tenant)# application PVLAN
apic1(config-tenant-app)# epg EPG1
apic1(config-tenant-app-epg)# show running-config
# Command: show running-config tenant Tenant_VMM application Web epg intraEPGDeny
tenant Tenant_VMM
  application Web
    epg intraEPGDeny
      bridge-domain member VMM_BD
      microsoft-domain member domain1 encap vlan-2003 primary-encap vlan-2004
      microsoft-domain member domain2
      exit
    isolation enforce
    exit
  exit
exit
apic1(config-tenant-app-epg)#
```


Step 2 Verify the configuration:**Example:**

```

show epg StaticEPG detail
Application Epg Data:
Tenant           : Test_Isolation
Application      : PVLAN
AEPg            : StaticEPG
BD              : VMM_BD
uSeg EPG       : no
Intra EPG Isolation : enforced
Vlan Domains    : VMM
Consumed Contracts : VMware_vDS-Ext
Provided Contracts : default,Isolate_EPG
Denied Contracts :
Qos Class       : unspecified
Tag List        :
VMM Domains:
Domain          Type      Deployment Immediacy Resolution Immediacy State      Encap
-----
Primary
Encap
-----
DVS1           VMware    On Demand           immediate        formed        auto
  auto
-----

Static Leaves:
Node          Encap          Deployment Immediacy Mode          Modification Time
-----
-----

Static Paths:
Node          Interface          Encap          Modification Time
-----
1018         eth101/1/1         vlan-100       2016-02-11T18:39:02.337-08:00
1019         eth1/16            vlan-101       2016-02-11T18:39:02.337-08:00

Static Endpoints:
Node          Interface          Encap          End Point MAC      End Point IP Address
Modification Time
-----
-----

Dynamic Endpoints:
Encap: (P):Primary VLAN, (S):Secondary VLAN
Node          Interface          Encap          End Point MAC      End Point IP Address
Modification Time
-----
1017         eth1/3             vlan-943 (P)    00:50:56:B3:64:C4  ---
2016-02-17T18:35:32.224-08:00
vlan-944 (S)

```

Configuring Intra-EPG Isolation for Cisco ACI Virtual Edge

Intra-EPG Isolation Enforcement for Cisco ACI Virtual Edge

By default, endpoints with an EPG can communicate with each other without any contracts in place. However, you can isolate endpoints within an EPG from each other. For example, you may want to enforce endpoint isolation within an EPG to prevent a VM with a virus or other problem from affecting other VMs in the EPG.

You can configure isolation on all or none of the endpoints within an application EPG; you cannot configure isolation on some endpoints but not on others.

Isolating endpoints within an EPG does not affect any contracts that enable the endpoints to communicate with endpoints in another EPG.



Note Enforcing intra-EPG Isolation is not supported for the EPG that is associated with Cisco ACI Virtual Edge domains in VLAN mode. If you try to enforce intra-EPG isolation with such an EPG, a fault is triggered.



Note Using intra-EPG isolation on a Cisco ACI Virtual Edge microsegment (uSeg) EPG is not currently supported.



Note Proxy ARP is not supported for Cisco ACI Virtual Edge EPGs using VXLAN encapsulation and on which intra-EPG Isolation is enforced. Therefore, intra-subnet communication is not possible between intra-EPG isolated EPGs even though contracts are in place between those Cisco ACI Virtual Edge EPGs. (VXLAN).

Configure Intra-EPG Isolation for Cisco ACI Virtual Edge Using the GUI

Follow this procedure to create an EPG in which the endpoints of the EPG are isolated from each other.

The port that the EPG uses must belong to one of the VM Managers (VMMs).



Note This procedure assumes that you want to isolate endpoints within an EPG when you create the EPG. If you want to isolate endpoints within an existing EPG, select the EPG in Cisco APIC, and in the **Properties** pane, in the **Intra EPG Isolation** area, choose **Enforced**, and then click **SUBMIT**.

Before you begin

Make sure that VXLAN-related configuration is present on the Cisco ACI Virtual Edge VMM domain, particularly a Cisco ACI Virtual Edge fabric-wide multicast address and pool of multicast addresses (one per EPG).

Procedure

-
- Step 1** Log in to Cisco APIC.
- Step 2** Choose **Tenants**, expand the folder for the tenant, and then expand the **Application Profiles** folder.
- Step 3** Right-click an application profile, and choose **Create Application EPG**.
- Step 4** In the **Create Application EPG** dialog box, complete the following steps:
- In the **Name** field, enter the EPG name.
 - In the **Intra EPG Isolation** area, click **Enforced**.
 - From the **Bridge Domain** drop-down list, choose the bridge domain.
 - Check the **Associate to VM Domain Profiles** check box.
 - Click **Next**.
 - In the **Associate VM Domain Profiles** area, complete the following steps:
 - Click the + (plus) icon, and from the **Domain Profile** drop-down list, choose the desired Cisco ACI Virtual Edge VMM domain.
 - From the **Switching Mode** drop-down list, choose **AVE**.
 - From the **Encap Mode** drop-down list, choose **VXLAN** or **Auto**.
If you choose **Auto**, make sure that encapsulation mode of the Cisco ACI Virtual Edge VMM domain is **VXLAN**.
 - (Optional) Choose other configuration options appropriate to your setup.
 - Click **Update** and click **Finish**.
-

What to do next

You can select statistics and view them to help diagnose problems involving the endpoint. See the sections [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 27](#) and [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 28](#) in this guide.

Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab

If you configured intra-EPG isolation on a Cisco ACI Virtual Edge, choose statistics—such as denied connections, received packets, or transmitted multicast packets—for the endpoints. You can then view the statistics.

Procedure

-
- Step 1** Log in to Cisco APIC.
- Step 2** Choose **Tenants** > *tenant*.
- Step 3** In the tenant navigation pane, expand the **Application Profiles**, *profile*, and **Application EPGs** folders, and then choose the EPG containing the endpoint the statistics for which you want to view.

- Step 4** In the EPG **Properties** work pane, click the **Operational** tab to display the endpoints in the EPG.
 - Step 5** Double-click the endpoint.
 - Step 6** In the **Properties** dialog box for the endpoint, click the **Stats** tab and then click the check icon.
 - Step 7** In the **Select Stats** dialog box, in the **Available** pane, choose the statistics that you want to view for the endpoint, and then use the right-pointing arrow to move them into the **Selected** pane.
 - Step 8** Click **Submit**.
-

View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab

If you configured intra-EPG isolation on a Cisco ACI Virtual Edge, once you have chosen statistics for the endpoints, you can view them.

Before you begin

You must have chosen statistics to view for isolated endpoints. See [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 27](#) in this guide for instructions.

Procedure

- Step 1** Log in to Cisco APIC.
- Step 2** Choose **Tenants > tenant**.
- Step 3** In the tenant navigation pane, expand the **Application Profiles, profile**, and **Application EPGs** folders, and then choose the EPG containing the endpoint with statistics that you want to view.
- Step 4** In the EPG **Properties** work pane, click the **Operational** tab to display the endpoints in the EPG.
- Step 5** Double-click the endpoint with statistics that you want to view.
- Step 6** In the **Properties** work pane for the endpoint, click the **Stats** tab.

The work pane displays the statistics that you chose earlier. You can change the view by clicking the table view or chart view icon on the upper left side of the work pane.

Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Virtual Networking Tab

If you configured intra-EPG isolation on a Cisco ACI Virtual Edge, choose statistics—such as denied connections, received packets, or transmitted multicast packets—for the endpoints. You can then view the statistics.

Procedure

- Step 1** Log in to Cisco APIC.
- Step 2** Choose **Virtual Networking > Inventory > VMM Domains > VMware > VMM domain > Controllers > controller instance name > DVS-VMM name > Portgroups > EPG name > Learned Point MAC address (node) > .**

- Step 3** Click the **Stats** tab.
 - Step 4** Click the tab with the check mark.
 - Step 5** In the **Select Stats** dialog box, click the statistics that you want to view in the **Available** pane, and then click the arrow pointing right to put them in the **Selected** pane.
 - Step 6** (Optional) Choose a sampling interval.
 - Step 7** Click **Submit**.
-

View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Virtual Networking Tab

If you configured intra-EPG isolation on a Cisco ACI Virtual Edge, once you have chosen statistics for the endpoints, you can view them.

Before you begin

You must have chosen statistics to view for isolated endpoints. See [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 27](#) in this guide for instructions.

Procedure

- Step 1** Log in to Cisco APIC.
 - Step 2** Choose **Virtual Networking > Inventory > VMM Domains > VMware > VMM name > Controllers > controller instance name > DVS-VMM name > Portgroups > EPG name > Learned Point MAC address (node)**
 - Step 3** Click the **Stats** tab.

The central pane displays the statistics that you chose earlier. You can change the view by clicking the table view or chart view icon on the upper left side of the work pane.
-

Configure Intra-EPG Isolation for Cisco ACI Virtual Edge Using the NX-OS Style CLI

Before you begin

Make sure that VXLAN-related configuration is present on the Cisco ACI Virtual Edge VMM domain, particularly a Cisco ACI Virtual Edge fabric-wide multicast address and pool of multicast addresses (one per EPG).

Procedure

In the CLI, create an intra-EPG isolation EPG:

Example:

```
# Command: show running-config tenant Tenant2 application AP-1 epg EPG-61
tenant Tenant2
  application AP-1
```

```
epg EPG-61
  bridge-domain member BD-61
  vmware-domain member D-AVE-SITE-2-3
  switching-mode AVE
  encap-mode vxlan
  exit
  isolation enforce          # This enables EPG into isolation mode.
  exit
exit
exit
```

What to do next

You can select statistics and view them to help diagnose problems involving the endpoint. See the sections [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 27](#) and [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 28](#) in this guide.

Troubleshooting

Troubleshooting Endpoint Connectivity

Procedure

- Step 1** Inspect the operational status of each endpoint.
- The operational status will reveal any fault or misconfiguration of the endpoints. See [Inspecting the Endpoint Status, on page 31](#).
- Step 2** Inspect the status of the tunnel interface.
- The operational status will reveal any fault or misconfiguration of the tunnel. See [Inspecting the Tunnel Interface Status, on page 31](#).
- Step 3** Perform a traceroute between the endpoint groups (EPGs).
- A traceroute will reveal any problems with intermediate nodes, such as spine nodes, between the endpoints. See [Performing a Traceroute Between Endpoints, on page 32](#).
- Step 4** Configure an atomic counter on an endpoint.
- The atomic counter will confirm whether the source endpoint is transmitting packets or the destination endpoint is receiving packets, and whether the number of packets received equals the number of packets sent. See [Configuring Atomic Counters, on page 33](#).
- Step 5** Inspect the contracts under each EPG.
- Inspect the contracts under each EPG to make sure they allow the traffic that should flow between the EPGs. As a test, you can temporarily open the contracts to allow unrestricted traffic.

Step 6 Configure a SPAN policy to forward source packets to a monitoring node.

A packet analyzer on the monitoring node will reveal any packet issues such as an incorrect address or protocol. See [Configuring a Tenant SPAN Session Using the Cisco APIC GUI, on page 33](#).

Inspecting the Endpoint Status

Procedure

Step 1 In the menu bar, click **Tenants**.

Step 2 In the submenu bar, click the tenant that contains the source endpoint.

Step 3 In the **Navigation** pane, expand the tenant, expand **Application Profiles**, and expand the application profile that contains the endpoint.

Step 4 Expand **Application EPGs** and click the EPG to be inspected.

Step 5 In the **Work** pane, from the list of endpoints in the **Endpoint** table, double-click the source endpoint to open the **Client End Point** dialog box.

Step 6 In the **Client End Point** dialog box, verify the endpoint properties and click the **Operational** tab.

Step 7 In the **Operational** tab, view the health, status, and fault information.

In the **Status** table, click any items with entries, such as changes, events, or faults.

Step 8 Close the **Client End Point** dialog box.

Step 9 In the **Endpoint** table, view the **Interface** entry for the endpoint and note the node and tunnel IDs.

Step 10 Repeat this procedure for the destination endpoint.

Note

Occasionally, bidirectional traffic is interrupted between IP addresses in two micro-segmented EPGs deployed behind two leaf switches in the fabric. This can occur when the IP addresses are transitioning because of a configuration change from micro-segment EPG to base EPG. Or conversely, this can occur on two different leaf switches at the same time while bidirectional traffic is running. In this case, the policy tag for each remote endpoint still points to its previous EPG.

Workaround: Manually clear the remote endpoints on the switches or wait for the remote endpoint to age out. To clear the endpoints, log on to the CLI on each switch and enter the **clear system internal epm endpoint** command with the appropriate option. For example, if your endpoints are based on the IP address, enter **clear system internal epm endpoint key vrf vrf_name{ip | ipv6} ip-address**. The endpoints are then relearned with the correct policy tag.

Inspecting the Tunnel Interface Status

This procedure shows how to inspect the operational status of the tunnel interface.

Procedure

-
- Step 1** In the menu bar, click **Fabric**.
- Step 2** In the submenu bar, click **Inventory**.
- Step 3** In the **Navigation** pane, expand the pod and expand the node ID of the source endpoint interface.
- Step 4** Under the node, expand **Interfaces**, expand **Tunnel Interfaces**, and click the tunnel ID of the source endpoint interface.
- Step 5** In the **Work** pane, verify the tunnel interface properties and click the **Operational** tab.
- Step 6** In the **Operational** tab, view the health, status, and fault information.
In the **Status** table, click any items with entries, such as changes, events, or faults.
- Step 7** Repeat this procedure for the destination endpoint interface.
-

Performing a Traceroute Between Endpoints

Procedure

-
- Step 1** In the menu bar, click **Tenants**.
- Step 2** In the submenu bar, click the tenant that contains the source endpoint.
- Step 3** In the **Navigation** pane, expand the tenant and expand **Policies > Troubleshoot**.
- Step 4** Under **Troubleshoot**, right-click on one of the following traceroute policies:
- **Endpoint-to-Endpoint Traceroute Policies** and choose **Create Endpoint-to-Endpoint Traceroute Policy**
 - **Endpoint-to-External-IP Traceroute Policies** and choose **Create Endpoint-to-External-IP Traceroute Policy**
 - **External-IP-to-Endpoint Traceroute Policies** and choose **Create External-IP-to-Endpoint Traceroute Policy**
 - **External-IP-to-External-IP Traceroute Policies** and choose **Create External-IP-to-External-IP Traceroute Policy**
- Step 5** Enter the appropriate values in the dialog box fields and click **Submit**.
- Note**
For the description of a field, click the help icon (?) in the top-right corner of the dialog box.
- Step 6** In the **Navigation** pane or the **Traceroute Policies** table, click the traceroute policy.
The traceroute policy is displayed in the **Work** pane.
- Step 7** In the **Work** pane, click the **Operational** tab, click the **Source Endpoints** tab, and click the **Results** tab.
- Step 8** In the **Traceroute Results** table, verify the path or paths that were used in the trace.
- Note**
- More than one path might have been traversed from the source node to the destination node.

- For readability, you can increase the width of one or more columns, such as the **Name** column.

Configuring Atomic Counters

Procedure

- Step 1** In the menu bar, click **Tenants**.
- Step 2** In the submenu bar, click the desired tenant.
- Step 3** In the **Navigation** pane, expand the tenant and expand **Policies** and then expand **Troubleshoot**.
- Step 4** Under **Troubleshoot**, expand **Atomic Counter Policy** and choose a traffic topology.
You can measure traffic between a combination of endpoints, endpoint groups, external interfaces, and IP addresses.
- Step 5** Right-click the desired topology and choose **Add topology Policy** to open an **Add Policy** dialog box.
- Step 6** In the **Add Policy** dialog box, perform the following actions:
- In the **Name** field, enter a name for the policy.
 - choose or enter the identifying information for the traffic source.
The required identifying information differs depending on the type of source (endpoint, endpoint group, external interface, or IP address).
 - choose or enter the identifying information for the traffic destination.
 - (Optional) (Optional) In the **Filters** table, click the + icon to specify filtering of the traffic to be counted.
In the resulting **Create Atomic Counter Filter** dialog box, you can specify filtering by the IP protocol number (TCP=6, for example) and by source and destination IP port numbers.
 - Click **Submit** to save the atomic counter policy.
- Step 7** In the **Navigation** pane, under the selected topology, choose the new atomic counter policy.
The policy configuration is displayed in the **Work** pane.
- Step 8** In the **Work** pane, click the **Operational** tab and click the **Traffic** subtab to view the atomic counter statistics.
-

Configuring a Tenant SPAN Session Using the Cisco APIC GUI

SPAN can be configured on a switch or on a tenant. This section guides you through the Cisco APIC GUI to configure a SPAN policy on a tenant to forward replicated source packets to a remote traffic analyzer. The configuration procedure requires entering values in the fields of one or more GUI dialog boxes. To understand a field and determine a valid value, view the help file by clicking the help icon (?) at the top-right corner of the dialog box.

Procedure

- Step 1** In the menu bar, click **Tenants**.

- Step 2** In the submenu bar, click the tenant that contains the source endpoint.
- Step 3** In the **Navigation** pane, expand the tenant, expand **Policies > Troubleshooting > SPAN**.
Two nodes appear under **SPAN**: **SPAN Destination Groups** and **SPAN Source Groups**.
- Step 4** From the **Navigation** pane, right-click **SPAN Source Groups** and choose **Create SPAN Source Group**.
The **Create SPAN Source Group** dialog appears.
- Step 5** Enter the appropriate values in the required fields of the **Create SPAN Source Group** dialog box.
- Step 6** Expand the **Create Sources** table to open the **Create SPAN Source** dialog box.
- Step 7** Enter the appropriate values in the **Create SPAN Source** dialog box fields.
- Step 8** When finished creating the SPAN source, click **OK**.
You return to the **Create SPAN Source Group** dialog box.
- Step 9** When finished entering values in the **Create SPAN Source Group** dialog box fields, click **Submit**.

What to do next

Using a traffic analyzer at the SPAN destination, you can observe the data packets from the SPAN source EPG to verify the packet format, addresses, protocols, and other information.

Verifying IP-Based EPG Configurations

There are two types of endpoint groups (EPGs) that you can create: application EPGs and IP-based EPGs. IP-based EPGs differ from regular application EPGs in that they are microsegment EPGs. This chapter explains how to verify that your IP-based EPG configurations are properly classified as IP-based using the GUI or using switch commands.

This chapter contains the following sections:

Verifying IP-Based EPG Configurations Using the GUI

This procedure explains how to verify that you have correctly configured an IP-based EPG using the GUI and Visore tool.

Procedure

-
- Step 1** Verify that the IP-based EPG you created is listed under the **uSeg EPGs** folder in the GUI (shown in the following screen capture).
Note that there is one IP-based EPG listed under uSeg EPGs named "IP" that was created using the REST API.
- Step 2** Verify that the information is correct in the EPG - IP properties screen (right side window pane) for each EPG IP (IP-based EPG).
Note the list of IP-based EPGs and IP addresses that are shown at the bottom of the screen.
- Step 3** From your web browser, enter the APIC IP address followed by "/visore.html." Visore is a tool that allows you to view all the objects in the system, such as EPGs. You can use Visore to verify that your IP-based EPGs have been properly

configured. For more information about Visore, see the *Application Policy Infrastructure Controller Visore Tool Introduction* document.

Step 4 Enter your username and password then click **Login** to log into Visore.

Step 5 Run a query for the IP-based EPGs that you verified in the GUI by entering the name of the class in the field next to **Class or DN** (for example, "fvAEPg").

Note

This is a view from the APIC point of view. You can see that the "Total objects shown" above is "3", meaning there are three EPGs that were downloaded to the switch. You can see that the IP-based EPG that was previously listed in the GUI as "IP" is now shown next to "dn". Also note that "yes" is displayed next to "isAttrBasedEPg", which means that this has been properly configured as an IP-based EPG. You can verify all the objects have been configured successfully using Visore, including both application EPGs and IP-based EPGs.

Step 6 This is a view from the switch point of view. On the switch, you can run a query for the fvEpP class to see the EPGs and check for the "crtmEnabled" attribute. It will be set to "yes" for IP-based EPGs.

Verify that under this EPG, the children of the EPG are shown with IP addresses to ensure a proper configuration. For each IP address configured, there is one object (named "I3IpCktEp") that the switch uses to classify the traffic. Once the configuration is there, when the packets arrive, the switch uses these objects to classify them.

Step 7 Verify that the pcTags for all the endpoints and IP addresses that you configured match. Every EPG has a pcTag. All the endpoints that match with the IP addresses you configured are classified into this pcTag. Every endpoint has an IP address that you can run a class query on. When you are troubleshooting, you want to verify whether these endpoints (servers) are properly getting classified into this IP-based EPG or not. (The pcTags should match for the IP-based EPG.)

Verifying IP-EPG Configurations Using Switch Commands

This procedure explains how to use switch commands to verify you IP-EPG ("IpCkt") configurations.

Procedure

Step 1 Log in to the leaf.

Step 2 Navigate to the /mit/sys directory.

Step 3 In the /mit/sys directory, find ctx (vrf context directory)

Step 4 In the VRF ctx directory, go to the specific BD directory where the IpCkt is configured. You should see the IpCkt.

Note

"IpCkt" and "IP-EPG" are used interchangeably in this document.

Step 5 Navigate to the directory and the "cat summary" gives you the information regarding IpCkt.

Step 6 Ensure that the summary's "operSt" does not say "unsupported".

Step 7 Find out the VLAN ID that corresponds to the BD where the IpCkt is configured.

Note

The VLAN ID can be found through any of the **show vlan internal bd-info** commands or through the **show system internal epm vlan all** command.

- Step 8** Once you find the VLAN ID of the BD, issue **show system internal epm <vlan-id> detail**. Here you should be able to see all the configured IpCkts with a specific sclass. (It should match that of what you see in the /mit/sys directory.)
- Step 9** Repeat the steps for vsh_lc that you followed for vsh.
- Step 10** Send the traffic with an IP matching the IpCtk in the BD, and through **show system internal epm endp ip <a.b.c.d>**, you can verify that the learned IP has the IP-flags for "sclass" and a specific sclass value.
- Step 11** Repeat the steps for vsh_lc that you followed for vsh.

List of the Switch Troubleshooting Commands Used in this Procedure:

```
Cd /mits/sys/ctx-vxlan.../bd-vxlan...
- cat summary
Vsh -c "show system internal epm vlan all" or
Vsh -c "show vlan internal bd-info"
Vsh -c "show system internal epm vlan <vlan-id> detail"
Vsh -c "show system internal epm endp ip <a.b.c.d>"
Vsh_lc -c "show system internal epm vlan all" or
Vsh_lc -c "show vlan internal bd-info"
Vsh_lc -c "show system internal epm vlan <vlan-id> detail"
vsh_lc -c "show system internal epm endp ip <a.b.c.d>"
vsh_lc -c "show system internal epm epg"
```