



Cisco ACI Virtual Edge Configuration Guide, Release 3.2(x)

First Published: 2021-06-08

Last Modified: 2023-07-28

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

New and Changed Information

This chapter contains the following sections:

- [New and Changed Information](#) , on page 1

New and Changed Information

The following table provides an overview of the significant changes to this guide up to this current release for Cisco ACI Virtual Edge.

Feature	Description	Where Documented
There are no new features in Cisco ACI Virtual Edge Release 3.2(x).	This document has no changes from the previous release.	N/A



CHAPTER 2

Cisco ACI Virtual Edge

This chapter contains the following sections:

- [What Cisco ACI Virtual Edge Is, on page 3](#)
- [Cisco ACI Virtual Edge and the VMware vCenter, on page 5](#)
- [Cisco ACI Virtual Edge in a Multipod Environment, on page 6](#)
- [Required Software, on page 7](#)
- [Cisco ACI vPod: Extending the Cisco ACI Fabric, on page 7](#)

What Cisco ACI Virtual Edge Is

Beginning with the Cisco APIC Release 3.1(1), the Cisco Application Centric Infrastructure (ACI) supports the Cisco ACI Virtual Edge. Cisco ACI Virtual Edge is the next generation of the Application Virtual Switch (AVS) for Cisco ACI environments. Cisco ACI Virtual Edge (AVE) is a hypervisor-independent distributed service VM that leverages the native distributed virtual switch that belongs to the hypervisor. Cisco ACI Virtual Edge runs in the user space, operates as a virtual leaf, and is managed by the Cisco Application Policy Infrastructure Controller (APIC).

If you use Cisco AVS, you can migrate to Cisco ACI Virtual Edge; if you use VMware VDS, you can run Cisco ACI Virtual Edge on top of it. Decoupling the Cisco ACI Virtual Edge from the kernel space makes the solution adaptable to different hypervisors. It also facilitates simple upgrades as Cisco ACI Virtual Edge is not tied to hypervisor upgrades. Cisco ACI Virtual Edge implements the OpFlex protocol for control plane communication. It supports two modes of traffic forwarding: local switching and no local switching.

Cisco ACI Virtual Edge Release 1.1(1a) supports only the VMware hypervisor. It leverages the vSphere Distributed Switch (VDS), which is configured in private VLAN (PVLAN) mode.

When network administrators create a Cisco ACI Virtual Edge VMM domain on Cisco APIC, they must associate the domain with a range of VLANs to be used for PVLAN pair association of port groups on the DVS. Server administrators do not need to associate PVLANS to port groups on vCenter because Cisco APIC automatically associates PVLAN pairs with the endpoint groups (EPGs).



Note EPGs in Cisco APIC are equivalent to port groups in vCenter.

Restrictions for Cisco AVE VMM Domains

Following are the restrictions while configuring Cisco AVE VMM domains:

- Fabric-wide multicast address cannot be in x.0.0.1 format.
- Multicast pool for EPGs cannot overlap with the fabric-wide multicast address configured for the domain.

VLANs which are used as internal VLANs on AVE (internal vlan pool), must not be allowed on physical interfaces assigned as uplinks for AVE-DVS. Such configuration(s) may lead to situations where AVE classifies ingress uplink traffic as internal, causing rouge endpoint registration and Opflex process instability.

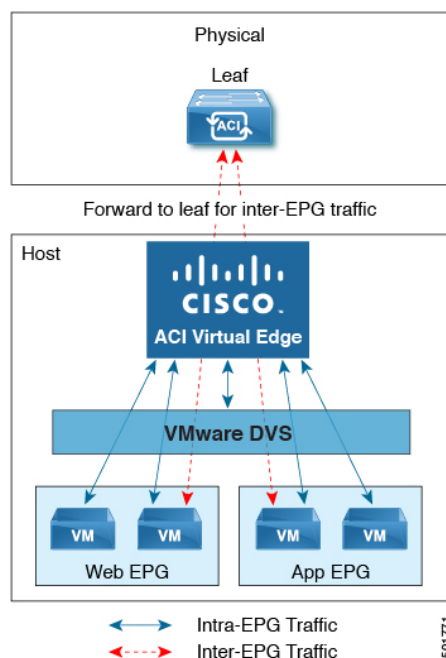
Local Switching Mode

In Local Switching mode, the Cisco ACI Virtual Edge locally forwards all intra-EPG traffic without involving the leaf. All inter-EPG traffic is forwarded through the leaf. In this mode, the Cisco ACI Virtual Edge can use either VLAN or VXLAN encapsulation—or both—for forwarding traffic to the leaf and back. You choose the encapsulation type during Cisco ACI Virtual Edge VMM domain creation.

You can configure a single VMM domain in Local Switching mode to use VLAN and VXLAN encapsulation.


If you choose VLAN encapsulation, a range of VLANs must be available for use by the Cisco ACI Virtual Edge. These VLANs have local scope in that they have significance only within the Layer 2 network between the Cisco ACI Virtual Edge and the leaf. If you choose VXLAN encapsulation, only the infra-VLAN must be available between the Cisco ACI Virtual Edge and the leaf. This results in a simplified configuration. It is the recommended encapsulation type if there are one or more switches between the Cisco ACI Virtual Edge and the physical leaf.

Figure 1: The Cisco ACI Virtual Edge in Local Switching Mode



No Local Switching Mode

In No Local Switching mode, the leaf forwards all traffic. In this mode, VXLAN is the only allowed encapsulation type.

Figure 2: The Cisco ACI Virtual Edge in No Local Switching Mode**Statistics Collection**

Statistics collection is enabled on Cisco ACI Virtual Edge by default. You may see Cisco ACI Virtual Edge faults within the Cisco APIC GUI relating to VM resource use.

Troubleshoot those faults in the VMware vCenter because the Cisco ACI only generates these faults based on information it receives from VMware vCenter.

Cisco ACI Virtual Edge and the VMware vCenter

The Cisco ACI Virtual Edge is a distributed virtual switch that extends across many virtualized hosts. It manages a data center defined by the vCenter Server.

The Cisco ACI Virtual Edge is compatible with any upstream physical access layer switch that complies with the Ethernet standard, including Cisco Nexus switches. The Cisco ACI Virtual Edge is compatible with any server hardware listed in the *VMware Hardware Compatibility List (HCL)*.


The Cisco ACI Virtual Edge is a distributed virtual switch solution that is fully integrated within the VMware virtual infrastructure, including VMware vCenter for the virtualization administrator. This solution allows the network administrator to configure virtual switch and port groups in order to establish a consistent data center network policy.



Note VMs deployed on Cisco ACI Virtual Edge support physical MAC addresses as indicated by vCenter; virtual MAC addresses are not supported.

The following figure shows a topology that includes the Cisco ACI Virtual Edge with the Cisco APIC and VMware vCenter.

Figure 3: Sample Cisco ACI Virtual Edge Topology



Note If there are multiple vCenters connected to a single Cisco ACI fabric, you should ensure that there are no overlapping MAC address allocation schema across the multiple vCenters while deploying the vCenters instead of the default OUI allocation. Overlaps can cause duplicate MAC address generation. For more information, see VMware documentation.

Cisco ACI Virtual Edge in a Multipod Environment

The Cisco ACI Virtual Edge can be part of a multipod environment. Multipod environments use a single Cisco APIC cluster for all the pods; all the pods act as a single fabric.

Multipod environments enable a more fault tolerant fabric comprising multiple pods with isolated control plane protocols. They also provide greater flexibility in full mesh cabling between leaf and spine switches.

Cisco ACI Virtual Edge does not require any additional configuration to operate in a multipod environment.

For detailed information about multipod environments, see the following documents on Cisco.com:

- *Cisco Application Centric Infrastructure Fundamentals*
- *Cisco APIC Getting Started Guide*
- *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*

The following features are not supported for Cisco ACI Virtual Edge with multipod in Cisco APIC releases 3.1(1) through 4.0(1):

- Storage vMotion with two separate NFS in two separate PODs
- ERSPAN destination in different PODs

- Distributed Firewall syslog server in different PODs

Required Software

The following table shows the versions of software required for Cisco ACI Virtual Edge to work with the Cisco APIC, VMware vCenter, and VMware ESXi hypervisor:

Component	Description
Cisco ACI Virtual Edge software	Cisco ACI Virtual Edge is supported beginning with Release 1.1(1).
Cisco APIC	Cisco ACI Virtual Edge is supported in Cisco APIC beginning with Release 3.1(1).
VMware vCenter	Cisco ACI Virtual Edge is compatible with release 6.0 and later versions of VMware vCenter Server.
VMware vSphere bare metal	Cisco ACI Virtual Edge is supported as a vLeaf for the Cisco APIC with release 6.0 and later releases of the VMware ESXi hypervisor.

Cisco ACI vPod: Extending the Cisco ACI Fabric

Organizations increasingly adopt hybrid data center models to meet infrastructure demands, flexibility, and reduce costs. They combine various technologies—including virtual private clouds and other internal IT resources—with remote locations. The remote locations can be hosted data centers, satellite data centers, or multicloud environments.

However, hybrid deployments require consistent management and policy for workloads regardless of their location. They also require support for disaster recovery and the ability to migrate workloads between data centers. Meanwhile, they can lack compatible hardware or space to add new equipment.

By deploying Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod), you can overcome these challenges and virtually extend the Cisco ACI fabric into various remote locations.

What Cisco ACI vPod Is

Cisco ACI vPod was introduced with general availability in Cisco APIC Release 4.0(2). It is a software-only solution that you can deploy wherever you have at least two servers on which you can run the VMware ESXi hypervisor. Cisco ACI vPod and its components—a virtual spine (vSpine), virtual leaf (vLeaf), and Cisco ACI Virtual Edge, run on the ESXi hypervisor.

Cisco ACI vPod allows you to use Cisco ACI Virtual Edge where you do not have a physical leaf. You can use up to eight instances of Cisco ACI Virtual Edge in each Cisco ACI vPod in the remote location as you would in your on-premises data center.

Cisco ACI vPod communicates with a physical, on-premises pod or multipod over an interpod network (IPN). You configure the physical pod or multipod, the IPN connection, and Cisco ACI vPod in Cisco Application Policy Infrastructure Controller (APIC). You then use the Cisco ACI vCenter plug-in, a Python script, or PowerCLI to deploy Cisco ACI vPod components.

Benefits of Cisco ACI vPod

Once Cisco ACI vPod is installed, you can use it with Cisco APIC to enforce Cisco ACI fabric policy in the remote location.

Cisco APIC provides central management of workloads in the on-premises data center and the remote location. It enables you to enforce policy easily and consistently in both on-premises and remote locations.

The flexibility, scalability, and central management of the Cisco ACI vPod solution enable you to take advantage of the following use case scenarios:

- Extension of the Cisco ACI fabric to the bare-metal cloud
- Extension of the Cisco ACI fabric to brownfield deployments
- Extension of the Cisco ACI fabric to colocation data centers
- Migration of workloads from non-Cisco hardware to the Cisco ACI fabric

Where to Find More Information

For general information, see the *Cisco ACI Virtual Pod Release Notes* on Cisco.com.



CHAPTER 3

VXLAN Load Balancing

- [VXLAN Load Balancing, on page 9](#)

VXLAN Load Balancing

Virtual extensible LAN (VXLAN) load balancing ensures that data moves efficiently between Cisco Application Centric Infrastructure (ACI) Virtual Edge and the leaf switch over multiple network interfaces under both of the following circumstances:

- When you have a MAC pinning policy and use VXLAN encapsulation
- When the MAC pinning policy and VXLAN encapsulation are enabled on the Cisco ACI Virtual Edge virtual machine manager (VMM) domain in the Cisco Application Policy Infrastructure Controller (APIC)

Beginning with this Cisco ACI Virtual Edge release, VXLAN load balancing is enabled by default. This Cisco ACI Virtual Edge release adds interfaces to accommodate VXLAN load-balancing and improve overall performance.



Note VXLAN load balancing is not supported for Cisco ACI Virtual Edge when Cisco ACI Virtual Edge is part of Cisco ACI Virtual Pod (vPod mode).

In previous releases, Cisco ACI Virtual Edge had three interfaces: one management, one internal, and one external. In VMware vCenter, there were two port groups, **internal** and **external**. Cisco ACI Virtual Edge now has the following:

- Two internal interfaces: Handle data traffic from the virtual machines (VMs). Traffic from private VLANs (PVLANS) is split evenly between two new internal port groups—**ave-internal-1** and **ave-internal-2** in VMware vCenter.
- Two external VXLAN interfaces: Load balance the VXLAN traffic. There are two new port groups in VMware vCenter—**ave-external-vxlan-1** and **ave-external-vxlan-2**, one for each interface. Infra VLAN, used by OpFlex, also uses the two external VXLAN interfaces.
- One external VLAN interface: Handles all VLAN-tagged traffic except for infra VLAN. It has its own VMware vCenter port group **ave-external-vlan**, which allows all Cisco ACI fabric VLANs, based on the VMM configuration.
- Management interface: Unchanged from previous releases.

- Two virtual tunnel endpoints (VTEPS) (kni interfaces) are created automatically and are pinned to each of the external-vxlan interfaces.



Note The names of the new VMware vCenter port groups are assigned automatically. Do not use these new port groups—which begin with **ave**—for tenant traffic.

In VMware vCenter, the **internal** and **external** port groups are still present. These port groups remain to accommodate Cisco ACI vPod and for upgrade and downgrade compatibility.

Verify VXLAN Load Balancing

You can verify that VXLAN load balancing is enabled by checking whether the Cisco Application Centric Infrastructure (ACI) Virtual Edge has received a DHCP IP address and whether OpFlex is up.

You run **ifconfig** commands to view kernel NIC information and check that Cisco ACI Virtual Edge has received a DHCP address. You run the **vemcmd show opflex** command to see if OpFlex is up.

View Kernel NIC Information

You can view information about the kernel NIC.

Enter the following commands: **ifconfig kni0** and **ifconfig kni2**, and check that kni0 and kni2 have been assigned IP addresses:

Example:

```
cisco-ave_198.51.100.62_AVE-FI:~$ ifconfig kni0
kni0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 198.51.100.111 netmask 255.255.0.0 broadcast 198.51.100.255
    inet6 fe80::250:56ff:feaf:807b prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:af:80:7b txqueuelen 1000 (Ethernet)
    RX packets 528552 bytes 50610919 (48.2 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 285294 bytes 44487029 (42.4 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
cisco-ave_198.51.100.62_AVE-FI:~$ ifconfig kni2

kni2: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 198.51.100.121 netmask 255.255.0.0 broadcast 198.51.100.255
    inet6 fe80::250:56ff:feaf:3dc9 prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:af:3d:c9 txqueuelen 1000 (Ethernet)
    RX packets 285152 bytes 17116682 (16.3 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 10873 bytes 2921194 (2.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
cisco-ave_198.51.100.62_AVE-FI:~$
```

The output displays information about the two virtual TEP internal uplinks.

Note You can enter the **ifconfig** command to view complete interface information, including information about ens160, the management interface, and kni1, the ave-ctrl interface for vMotion.

View OpFlex Information

You can see if OpFlex is online and view its runtime status.

Enter the following command: **vemcmd show opflex**.

Example:

```
cisco-ave_198.51.100.62_AVE-FI:~$ vemcmd show opflex
Status: 12 (Active)
Channel0: 12 (Active), Channel1: 12 (Active)
Dvs name: comp/prov-VMware/ctrlr-[AVE-FI]-vC-191/sw-dvs-413
Remote IP: 192.0.2.11 Port: 8000
Infra vlan: 5
FTEP IP: 192.0.2.20
Switching Mode: LS
Encap Type: VXLAN
NS GIPO: 225.10.10.1
cisco-ave_198.51.100.62_AVE-FI:~$
```



CHAPTER 4

Mixed-Mode Encapsulation

This chapter contains the following sections:

- [Mixed-Mode Encapsulation Configuration, on page 13](#)
- [Check or Change the VMM Domain Encapsulation Mode Using the APIC GUI, on page 14](#)
- [Check or Change the VMM Domain Encapsulation Mode Using the NX-OS CLI, on page 15](#)
- [Check or Change the VMM Domain Encapsulation Mode Using the REST API, on page 15](#)
- [Override the VMM Domain Encapsulation Mode for an EPG Using the APIC GUI, on page 16](#)
- [Override the VMM Domain Encapsulation Mode for an EPG Using the NX-OS Style CLI, on page 17](#)
- [Override the VMM Domain Encapsulation Mode for an EPG Using the REST API, on page 17](#)

Mixed-Mode Encapsulation Configuration

You can configure a single VMM domain to use VLAN and VXLAN encapsulation. Mixed-mode encapsulation enables you to have a single domain for all EPGs, regardless of encapsulation mode. That makes it easier to track and manage EPGs.

When you create a VMM domain, you can explicitly choose its encapsulation mode: VLAN or VXLAN.

When you create a new EPG for the VMM domain, each EPG for the domain by default uses the VMM domain's encapsulation mode. However, when you create a new EPG and associate it with a domain, you can configure the EPG to override the domain encapsulation mode and use another mode.

For example, you may choose VLAN configuration when you create a VMM domain. When you create a new EPG for the domain, you may configure it to use VLAN—the domain mode—or you may configure it to use VXLAN.

Ensure non overlapping VLAN blocks (internal for AVE and external for DVS) are used for AVE and DVS if AEP remains same for AVE-deployed EPGs in VLAN mode.



Note Mixed-mode encapsulation is available for Cisco ACI Virtual Edge in local switching mode only.

Encapsulation Pool Combinations

Your ability to add and delete VLAN and multicast pools for a VMM domain depends on whether EPGs are associated with the domain.

If no EPGs are associated with the VMM domain, you can add and delete VLAN and multicast pools. You can do this regardless of whether the VMM domain default encapsulation mode is VLAN or VXLAN.

If EPGs are associated with the VMM domain, you cannot delete existing VLAN or multicast pools.:

- **VLAN**—You can configure both VLAN and multicast pools.

VLAN is the default encapsulation mode for the domain. New EPGs created for this VMM domain use VLAN encapsulation by default. You can configure EPGs to use VXLAN encapsulation if multicast pools are configured in the VMM domain.



Note Configure the private VLAN as internal in the VLAN pool for internal switching.

- **VXLAN**—You can configure both VLAN and multicast pools. VXLAN is the default encapsulation mode for the domain. New EPGs created for the VMM domain use VXLAN encapsulation by default. You can configure EPGs to use VLAN encapsulation if VLAN pools are configured in the VMM domain.

Check or Change the VMM Domain Encapsulation Mode Using the APIC GUI

You can use the APIC GUI to discover and change the encapsulation mode of a VMM domain.



Note If EPGs are associated to the VMM domain, you cannot change its switching mode. If you want the domain to use a different switching mode, delete and re-create it. However, you can change the switching mode of the VMM domain if no EPGs are associated to it.

Step 1 Log in to the Cisco APIC.

Step 2 Go to **Virtual Networking > Inventory > VMM Domains > VMware > VMM domain**.

In the VMM domain work pane, in the **Properties** area, the **Default Encapsulation Mode** field highlights **VLAN** or **VXLAN** in blue.

Step 3 If you wish, change the mode in the **Default Encapsulation Mode** by clicking on the preferred mode.

Step 4 If needed, configure a VLAN or multicast pool in the work pane.

In order to change the default mode to **VLAN**, you must have a VLAN pool configured. In order to change the default mode to **VXLAN**, you must have a multicast address and multicast pool configured.

Note For both VLAN and VXLAN modes, if you have not already done so, you must configure an internal VLAN pool for the private VLAN, which is used for internal switching.

Step 5 Click **Submit**.

Check or Change the VMM Domain Encapsulation Mode Using the NX-OS CLI

You can use the NX-OS CLI to check or change the encapsulation mode of a VMM domain.



Note If EPGs are associated to the VMM domain, you cannot change its switching mode. If you want the domain to use a different switching mode, delete and re-create it. However, you can change the switching mode of the VMM domain if no EPGs are associated to it.

Step 1 Check the VMM domain encapsulation mode.

Example:

```
apic1(config-vmware-ave)# show run
# Command: show running-config vmware-domain mininet1 configure-ave
# Time: Tue Nov 21 07:07:58 2017
vmware-domain mininet1
  configure-ave
    switching mode vlan
    multicast-address 230.1.2.3
  exit
exit
apic1(config-vmware-ave)#
```

Step 2 Change the VMM domain encapsulation mode.

Example:

```
apic1# configure
apic1(config)# vmware-domain mininet
apic1(config-vmware)# configure-ave
apic1(config-vmware-ave)# switching mode ?
vlan          VLAN/SW Mode
vxlan         VXLAN/SW Mode
vxlan-ns      VXLAN/HW Mode
```

Check or Change the VMM Domain Encapsulation Mode Using the REST API

You can use the REST API to discover and change the encapsulation mode of a VMM domain.



Note If EPGs are associated to the VMM domain, you cannot change its switching mode. If you want the domain to use a different switching mode, delete and re-create it. However, you can change the switching mode of the VMM domain if no EPGs are associated to it.

Discover and change the encapsulation mode of a VMM domain.

Example:

```
<vmmProvP vendor="VMware">
  <vmmDomP name="mininet" enableAVE="true" enfPref="sw" mcastAddr="225.1.1.1" encapMode="vxlan"
  prefEncapMode="vxlan">
</vmmProvP>
```

Override the VMM Domain Encapsulation Mode for an EPG Using the APIC GUI

After you create an EPG and associate it with a VMM domain, you can change the encapsulation mode of the EPG. You can make the encapsulation mode different from or the same as that of the VMM domain.

Before you begin

You must already have created an EPG and have associated it with a VMM domain.

-
- Step 1** Log in to the Cisco APIC.
- Step 2** Go to **Tenants > tenant > Application Profiles > application profile > Application EPGs > EPG > Domains (VMs and Bare-Metals)**.
- Step 3** In the **Domains (VMs and Bare-Metals)** work pane, double-click a domain, ensure that the switching mode is **AVE**, and then choose a mode from the **Encap Mode** drop-down list.

You can choose one of the following encap modes:

- **VXLAN**: This overrides the domain's VLAN configuration, and the EPG uses VXLAN encapsulation. However, a fault is triggered for the EPG if a multicast pool is not configured on the domain.
- **VLAN**: This overrides the domain's VXLAN configuration, and the EPG uses VLAN encapsulation. However, a fault is triggered for the EPG if a VLAN pool is not configured on the domain.
- **Auto**: This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.

- Step 4** Click **Update**.
-

What to do next

Verify the configuration by checking the endpoints under the EPG for the encapsulation mode.

Override the VMM Domain Encapsulation Mode for an EPG Using the NX-OS Style CLI

After you create an EPG and associate it with a VMM domain, you can change the encapsulation mode of the EPG so it differs from or is the same of the VMM domain encapsulation mode.

Before you begin

You must already have created an EPG and have associated it with a VMM domain.

Specify the encapsulation mode for an EPG:

Example:

```
apicl(config)# tenant <tenant name>
apicl(config-tenant)# application <application name>
apicl(config-tenant-app)# epg <epg name>conf
apicl(config-tenant-app-epg)# vmware-domain member <vmm domain name>
apicl(config-tenant-app-epg-domain)# encap-mode auto | vlan | vxlan
apicl(config-tenant-app-epg-domain)# switching-mode AVE
```

You can choose one of the following encapsulation modes:

- **Auto**—This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
- **VLAN**—This overrides the domain's VXLAN configuration, and the EPG will use VLAN encapsulation. However, a fault will be triggered for the EPG if a VLAN pool is not configured on the domain.
- **VXLAN**—This overrides the domain's VLAN configuration, and the EPG will use VXLAN encapsulation. However, a fault will be triggered for the EPG if a multicast pool is not configured on the domain.

Override the VMM Domain Encapsulation Mode for an EPG Using the REST API

Override the VMM domain encapsulation mode for an EPG.

Example:

```
<polUni>
<fvTenant name="coke">
<fvAp name="sap">
<fvAEPg name="web1">
<fvRsDomAtt resImedcy="immediate"
tDn="uni/vmmp-VMware/dom-mininet"
switchingMode="AVE" encapMode="vxlan"/>
</fvAEPg>
</fvAp>
</fvTenant>
</polUni>
```

For **encapMode=**, you can enter one of the following:

- **auto**—This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
 - **vlan**—This overrides the domain's VXLAN configuration, and the EPG will use VLAN encapsulation. However, a fault will be triggered for the EPG if a VLAN pool is not configured on the domain.
 - **vxlan**—This overrides the domain's VLAN configuration, and the EPG will use VXLAN encapsulation. However, a fault will be triggered for the EPG if a multicast pool is not configured on the domain.
-



CHAPTER 5

Port Channel and Virtual Port Channel Configuration

This chapter contains the following sections:

- [Port Channel or Virtual Port Channel Configuration, on page 19](#)
- [Configure a Port Channel or Virtual Port Channel Using the GUI, on page 19](#)
- [Configure a Port Channel Policy, on page 20](#)
- [Enhanced LACP Policy Support, on page 21](#)

Port Channel or Virtual Port Channel Configuration

You can configure a port channel or virtual port channel or a port channel policy using the Cisco APIC GUI, NX-OS style CLI, or REST API.

Configure a Port Channel or Virtual Port Channel Using the GUI

Use the Cisco APIC GUI to configure a port channel or virtual port channel.

- Step 1** Log in to the Cisco APIC.
- Step 2** Choose **Fabric > Access Policies**.
- Step 3** Expand the **Interface** and **Leaf Interfaces** folders.
- Step 4** Right-click the **Profiles** folder and choose **Create Leaf Interface Profile**.
- Step 5** In the **Create Leaf Interface Policy** dialog box, enter a name for the policy in the **Name** field.
- Step 6** In the **Interface Selectors** area, click + to add an access port selector.
- Step 7** In the **Create Access Port Selector** dialog box, complete the following steps:
- a) In the **Name** field, enter a name for the access port.
 - b) In the **Interface IDs** field, enter the interface IDs where the host is located.
 - c) From the **Interface Policy Group** drop-down list, choose **Create PC Interface Policy Group** or **Create VPC Interface Policy Group**.
- Step 8** In the **Create PC Interface Policy Group** dialog box or the **Create VPC Interface Policy Group** dialog box, complete the following steps:

- a) In the **Name** field, enter a name for the port channel.
- b) From the **Port Channel Policy** drop-down list, choose **Create Port Channel Policy**.

Step 9

In the **Create Port Channel Policy** dialog box, complete the following actions:

- a) In the **Name** field, enter a name for the policy.
- b) In the **Mode** field, choose one of the following options appropriate to your setup:
 - **Static Channel - Mode On**
 - **LACP Active**
 - **LACP Passive**
 - **MAC Pinning**
 - **MAC Pinning-Physical-NIC-load**

Note LACP Passive mode is not supported for directly connected hosts. Ports using LACP Passive mode do not initiate an LACP handshake. We recommend that you always use LACP Active instead of LACP Passive. LACP Passive can be used only with Cisco ACI Virtual Edge/TOR policy groups when there is an intermediate Layer 2 device and the Layer 2 device ports are using LACP Active mode.

Note MAC Pinning-Physical-NIC-load mode is not supported for Cisco ACI Virtual Edge.

- c) Click **Submit**.

Step 10

In the **Create PC Interface Policy Group** or **Create VPC Interface Policy Group** dialog box, from the **Attached Entity Profile** drop-down list, choose or create an attached entity profile, and then click **Submit**.

Step 11

In the **Create Access Port Selector** dialog box, click **OK**.

Step 12

In the **Create Leaf Interface Policy** dialog box, click **Submit**.

Configure a Port Channel Policy

You can configure one of several types of port channel policies on the Cisco ACI Virtual Edge:

- Link Aggregation Control Policy (LACP) in active mode
- Link Aggregation Control Policy (LACP) in passive mode
- Static mode
- MAC Pinning

You can configure port channel policies through the Cisco APIC GUI or the REST API. However, you can configure port channel mode using the NX-OS Style CLI.



Note When an LACP policy is applied as a vSwitch policy for the VMM domain, the LACP policy is applied only to the VMware vSphere Distributed Switch (VDS) uplinks. However, it is not applied to the Cisco ACI Virtual Edge port channel. This is expected behavior. Cisco ACI Virtual Edge does not support LACP on its uplinks because VDS does not support it for its virtual Ethernet (vEth) interfaces. So the VMM port channel policy is applied only for the VDS uplinks.

Enhanced LACP Policy Support

In Cisco Application Policy Infrastructure Controller (APIC) Release 3.2(7), you can improve uplink load balancing by applying different Link Aggregation Control Protocol (LACP) policies to different distributed virtual switch (DVS) uplink port groups.

Cisco APIC now supports VMware's Enhanced LACP feature, which is available for DVS 5.5 and later. Previously, the same LACP policy applied to all DVS uplink port groups. Before Cisco APIC Release 3.2(7), it was not possible to manage VMware link aggregation groups (LAGs) with Cisco APIC.

When you enable Enhanced LACP policy on the ACI side, it will push the configuration to DVS. Later, even if you remove the policy on the ACI side, enhanced LACP is still available on the DVS side, because after an enhanced LACP policy is enabled, it can not be reverted.



Note Enhanced LACP can be enabled either on the ACI or DVS side.

You can choose from up to 20 different load-balancing algorithms when you create a VMware vCenter virtual machine manager (VMM) domain for Cisco Application Centric Infrastructure (ACI) Virtual Edge or VMware VDS. You apply different policies to different uplink portgroups.

You have eight DVS uplink portgroups, and you must configure at least two uplinks in the same policy. So you can have up to four different LACP policies for each DVS. Enhanced LACP supports only active and passive LACP modes.



Note For Cisco ACI Virtual Edge VXLAN mode, it is mandatory to use a load-balancing algorithm having a UDP port. We recommend the algorithm **Source and Destination TCP/UDP Port**. In VLXAN mode, traffic is always sent between VTEP to the FTEP IP. So communication is always between one pair of IP address. So for VXLAN traffic, the only way to distinguish traffic is using the UDP port number.

Beginning with Release 5.2, Enhanced LACP policy is supported on interfaces of L4-L7 service devices used in service graphs. See *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide* for details about L4-L7 services.

The following sections provide instructions for configuring multiple LACP policies for DVS uplinks using the Cisco APIC GUI, NX-OS style CLI, or REST API.

Enhanced LACP Limitations

Be aware of the following limitations when using enhanced Link Aggregation Control Protocol (LACP) policies.

- You cannot fall back to the previous version of LACP after upgrading to enhanced LACP.
- You cannot downgrade to a version of Cisco Application Policy Infrastructure Controller (APIC) earlier than 3.2(7) without removing the enhanced LACP configuration. See the procedure [Remove the Enhanced LACP Configuration Before a Downgrade, on page 24](#) in this guide.

- For Cisco Application Centric Infrastructure (ACI) Virtual Edge, VXLAN mode traffic always uses the source IP address as the TEP IP address. To ensure proper load balancing, we recommend the algorithm **Source and Destination TCP/UDP Port**.
- If traffic is present for a Cisco ACI Virtual Edge domain over enhanced LACP, and you increase or reduce the number of uplinks, a traffic loss of 5 or 10 seconds occurs.
- Traffic is disrupted when an enhanced LACP LAG policy name conflicts with the name of a previous enhanced LACP link aggregation group (LAG) policy uplink. If you have an enhanced LACP LAG policy that is named ELACP-DVS for a DVS domain, its uplink is automatically named ELACP-DVS-1, ELACP-DVS-2, ELACP-DVS-3, and so on, depending on the number uplinks configured in the policy.

Traffic loss occurs if you then try to configure or add another enhanced LAG policy with a name that conflicts with a previous policy uplink name. To remedy the issue, delete the LAG policy and re-create it with a different name.

Create LAGs for DVS Uplink Port Groups Using the Cisco APIC GUI

Improve distributed virtual switch (DVS) uplink port group load balancing by putting the port groups into link aggregation groups (LAGs) and associating them with specific load-balancing algorithms. You can perform this task using the Cisco Application Policy Infrastructure Controller (APIC) GUI.

Before you begin

- You must have created a VMware vCenter virtual machine manager (VMM) domain for VMware VDS or Cisco Application Centric Infrastructure (ACI) Virtual Edge.
- If a vSwitch policy container does not exist, create one.



Note

You must configure a port channel policy before you create an enhanced LAG policy. You can create a port channel policy when you create a vCenter domain profile.

- Step 1** Log into the Cisco APIC.
- Step 2** Go to **Virtual Networking > Inventory > VMM Domains > VMware > domain**.
- Step 3** In the work pane, choose **Policy > VSwitch Policy**.
- Step 4** If you have not already done so, in the **Properties** area, choose a policy.
- Step 5** In the **Enhanced LAG Policy** area, click the + (plus) icon and then complete the following steps:
 - a) In the **Name** field, enter the name of the LAG.
 - b) From the **Mode** drop-down list, choose **LACP Active** or **LACP Passive**.
 - c) From the **Load Balancing Mode** drop-down list, choose a load-balancing method.
 - d) In the **Number of Links** selector, choose how many DVS uplink port groups to include in the LAG.
You can put two to eight uplink port groups into a LAG.
 - e) Click **Update** and then click **Submit**.

Step 6 Repeat Step 5 to create other LAGs for the DVS.

What to do next

If you are using VMware VDS, associate endpoint groups (EPGs) to the domain with the enhanced LACP policy. If you are using Cisco Application Centric Infrastructure (ACI) Virtual Edge, associate internally created inside and outside port groups with the enhanced LACP policy, then associate EPGs to the domain with the policy.

Associate Internal Port Groups to VMware vCenter Domains with Enhanced LACP Policies Using the Cisco APIC GUI

Associate Cisco Application Centric Infrastructure (ACI) Virtual Edge internally created inside and outside port groups with a VMware vCenter domain with an enhanced LACP policy. You can perform this task using the Cisco Application Policy Infrastructure Controller (APIC) GUI.

Before you begin

You must have created link aggregation groups (LAGs) for distributed virtual switch (DVS) uplink port groups and associated a load-balancing algorithm to the LAGs.

- Step 1** Log into the Cisco APIC.
 - Step 2** Go to **Virtual Networking > Inventory > VMM Domains > VMware > domain**.
 - Step 3** In the work pane, choose **Policy > General**.
 - Step 4** From the **Enhanced LAG Policy** drop-down list, choose a policy.
 - Step 5** Click **Submit**.
-

What to do next

Associate endpoint groups (EPGs) with the VMware vCenter domain containing the enhanced LACP policy.

Associate Application EPGs to VMware vCenter Domains with Enhanced LACP Policies Using the Cisco APIC GUI

Associate application endpoint groups (EPGs) with the VMware vCenter domain with LAGs and a load-balancing algorithm. You can perform this task using the Cisco Application Policy Infrastructure Controller (APIC) GUI.

Before you begin

You must have created link aggregation groups (LAGs) for distributed virtual switch (DVS) uplink port groups and associated a load-balancing algorithm to the LAGs.



Note This procedure assumes that you have not yet associated an application EPG with a VMware vCenter domain. If you have already done so, you edit the domain association.

-
- Step 1** Log into Cisco APIC.
- Step 2** Go to **Tenants > tenant > Application Profiles > application_profile > Application EPGs > EPG > Domains (VMs and Bare-Metals)**.
- Step 3** Right-click **Domains (VMs and Bare-Metals)** and choose **Add VMM Domain Association**.
- Step 4** In the **Add VMM Domain Association** dialog box, complete the following steps:
- From the **VMM Domain Profile** drop-down list, choose the domain that you want to associate the EPG to.
 - From the **Enhanced Lag Policy**, choose the policy configured for the domain that you want to apply to the EPG.
 - (Optional) In the **Delimiter** field, enter one of the following: |, ~, !, @, ^, +, or =.
If you do not enter a symbol, the system default | delimiter will appear in the policy.
 - Add remaining values as desired for the domain association, and then click **Submit**.
- Step 5** Repeat Step 2 through Step 4 for other application EPGs in the tenant as desired.
-

Remove the Enhanced LACP Configuration Before a Downgrade

Before you downgrade Cisco Application Policy Infrastructure Controller (APIC) to a release earlier than 3.2(7), you must remove the enhanced LACP configuration. Complete the steps in this procedure to remove the configuration.



Note Before downgrading, see [Enhanced LACP Limitations, on page 21](#) section for the required action, based on LAG support.

-
- Step 1** Reassign uplinks on all ESXi hosts from link aggregation groups (LAGs) to normal uplinks.
- Step 2** Remove LAG associations from all EPGs and interfaces of L4-L7 service devices used in service graphs, associated with the distributed virtual switch (DVS).
You can expect traffic loss while performing this step.
- Step 3** Change port channel settings to static channel or MAC pinning, which will cause traffic to recover once the port channel is up.
- Step 4** Remove all LAG-related configuration from the virtual machine manager (VMM).
- Step 5** Verify that all LAG-related policies are deleted from VMware vCenter.
-

What to do next

Downgrade to a Cisco APIC release earlier than 3.2(7).



CHAPTER 6

SPAN Features

This chapter contains the following sections:

- [About SPAN Feature Configuration, on page 25](#)
- [Configure SPAN Features Using the GUI, on page 26](#)

About SPAN Feature Configuration

The Cisco ACI Virtual Edge supports Switched Port Analyzer (SPAN) features, including local SPAN and Encapsulated remote SPAN (ERSPAN).

You cannot use the Cisco ACI Virtual Edge inside or outside interface uplinks as the source or destination of a SPAN sessions. The Cisco ACI Virtual Edge supports 64 SPAN sessions per DVS (local SPAN and ERSPAN). A source can be a member of a maximum of four SPAN sessions.

Guidelines for Configuring SPAN

Follow these guidelines when you configure local SPAN sessions on the Cisco ACI Virtual Edge:

- You can have only a single vLeaf per session.
- Sessions are defined by a client end point (CEP). EPG as a destination is not supported.
- Sessions are deployed on the vLeaf when a destination CEP is defined.
- No regular traffic is allowed from or to the destination CEP.
- A separate EPG with promiscuous mode enabled must be created for LSPAN destination CEP.

Guidelines for Configuring ERSPAN

Follow these guidelines when you configure ERSPAN sessions on the Cisco ACI Virtual Edge:

- Sessions are defined based on an IP address with other optional parameters.
- Sessions can be deployed on multiple vLeafs.
- Sessions are deployed to a vLeaf when a source CEP or endpoint group (EPG) is defined.
- The destination for an ERSPAN session should always be in overlay-1 (infraVRF [virtual routing and forwarding]). If the destination is a VM behind the Cisco ACI Virtual Edge, bring it up in the infra EPG.

The ERSPAN destination should always be remote. ERSPAN from a Cisco ACI Virtual Edge to a destination hosted behind the same Cisco ACI Virtual Edge is not supported.

- If the ERSPAN destination is a VM, make sure that vMotion is disabled on it. If the ERSPAN destination VM is moved to another host for any reason, make sure that the static CEP is configured accordingly. See Step 21 through Step 24 in the section [Configure SPAN Features Using the GUI, on page 26](#).
- The IP address for the destination can be obtained using DHCP (Option 61 is needed during DHCP) or static configuration. Make sure that the IP address is in the same subnet as the other VTEPs in overlay-1 (infra VRF).



Note Not all operating systems for VMs and devices support Option 61 for DHCP. In those cases, use a static IP address on infra VLAN. Choose a static IP address for ERSPAN carefully because it might lead to an IP conflict with the leased DHCP IPs on infra VLAN.

Guidelines for Configuring SPAN or ERSPAN with a UCS B Series Server

If you want to configure SPAN or ERSPAN on Cisco ACI Virtual Edge, and the Cisco ACI Virtual Edge hosts are running on a UCS B Series server, you must configure a port channel (PC) interface policy group with MAC pinning for the interfaces connecting to the fabric interconnects. This is because the virtual source (vsource) and virtual destination (vdestination) groups are specified only on PC policy groups.

Configure SPAN Features Using the GUI

Before you begin

If you are configuring LSPAN, you must have a new EPG configured with **Promiscuous** mode to capture local traffic on the same host. This EPG should be used on the VM that captures the traffic. Complete the following steps:

1. Create a new EPG and associate it to the VMM domain, choosing **AVE** as the switching mode and **Auto** as the encapsulation mode.
2. Enable **Promiscuous** mode on the EPG.

In Cisco APIC, expand the EPG, click **Domains(VMs and Bare-metals)**, right-click the VMM already associated with the EPG, and then click **Edit VMM Domain association**, set **Allow Promiscuous to Accept**, and then click **OK**.

-
- Step 1** Log in to the Cisco APIC.
 - Step 2** On the menu bar, choose **Fabric** > **Access Policies**.
 - Step 3** In the **Policies** navigation pane, open the **Policies** and the **Troubleshooting** folders.
 - Step 4** Expand the **VSPAN** folder.
 - Step 5** Right-click the **VSPAN Destinations Groups** folder and choose **Create VSPAN Destination Group**.
 - Step 6** In the **Create VSPAN Destination Group** dialog box, complete the following steps:

- a) In the **Name** field, enter a name.
- b) In the **Create Destinations** area, click the + icon.

Step 7

In the **Create VSPAN VDestination** dialog box, complete the following steps:

- a) In the **Destination Type** field, choose **ERSPAN** or **LSPAN** (for local SPAN).
- b) Complete one of the following series of steps:

If in Step 7 a you chose...	Then...
ERSPAN	<p>Enter the following values:</p> <ul style="list-style-type: none"> • Name—Enter a name for the VSPAN destination (Destination1). • Description—(Optional) Enter a description for the VSPAN destination. • Destination Type—Choose ERSPAN. • Destination IP—Specify a destination IP address. • Flow ID—Specify a flow ID value. • TTL—Specify a TTL value (64). • MTU—Specify an MTU value (1510). • DSCP—Enter a QoS DSCP value.
LSPAN	<p>Enter the following values:</p> <ul style="list-style-type: none"> • Name—Enter a name for the VSPAN destination (Destination1). • Description—(Optional) Enter a description for the VSPAN destination. • Destination Type—Choose LSPAN. • Destination CEP—(Optional) Choose a Tenant (1), Application Profile (a1), and EPG (e1), and CEP MAC address for the destination. <p>You see the destination CEP MAC address if you fulfilled the prerequisites for LSPAN.</p> <p>Note When you configure the destination CEP, choose the EPG that you created in the "Before You Begin" section with Promiscuous mode enabled.</p>

- c) Click **OK** to save the VSPAN destination.

Step 8

In the **Create VSPAN Destination** dialog box, click **Submit** to save the VSPAN destination group.

Step 9

In the **Policies** navigation pane, right-click the **VSPAN Sessions** folder and choose **Create VSPAN Session**.

- Step 10** In the **Create VSPAN Session** dialog box, in the **Name** field, enter a name for the source group.
- Step 11** In the **Admin State** field, ensure that **Start** is chosen.
- Step 12** From the **Destination Group** drop-down list, choose the new destination group.
- Step 13** In the **Create Sources** area, click the + icon.
- Step 14** In the **Create VSPAN VSource** dialog box, complete the following steps:
- In the **Name** field, enter a name for the source.
 - In the **Direction** area, choose a direction for the source (**Both**, **Incoming**, or **Outgoing**).
 - In the **Source type** area, choose **EPG** or **CEP**.
 - In the **Source EPG** or **Source CEP** area, choose a tenant, an application profile, and an EPG from the drop-down lists.
 - If you choose CEP as the source type, also choose a CEP from the drop-down list.
 - Disregard the **Add Source Access Paths** area.
 - Click **OK** to save the VSPAN VSource.
- Step 15** Click **Submit** to save the VSPAN VSource Group.
- Step 16** On the menu bar, choose **Fabric > Access Policies**.
- Step 17** In the **Policies** navigation pane, expand the **Interfaces**, **Leaf Interfaces**, and **Policy Groups** folders.
- Step 18** Expand the **VPC Interface** folder and click the policy group through which the SPAN source or destination is to be connected.
- Step 19** In the **PC/VPC Interface Policy Group** work pane for the policy group, complete the following steps:
- From the **Attached Entity Profile** drop-down list, choose or create an attached entity profile.
- Note** You may need to scroll down the page to complete the next steps.
- In the **VSource Groups** area, click the + icon, choose the desired SPAN source group, and then click **Update**.
This is the name of the source you that you created in Step 14 a.
 - In the **VDestination Group** area, choose the SPAN destination group, and then click **Update**.
This is the name of the destination you that you created in Step 7 b.
 - Click **Submit**.
These steps associate the SPAN source and SPAN destination groups with the selected policy groups.
- Step 20** To verify the configuration, open an SSH session on Cisco ACI Virtual Edge and enter the **vemcmd show span** command to display active SPAN sessions. Verify that the new session is running.
- Note** Step 21 through Step 24 are for ERSPAN only.
- Step 21** In the APIC GUI, on the menu bar, choose **Tenants > infra**
- Step 22** In the **Tenant infra** navigation pane, expand the following: **Application Profiles > access > Application EPGs > EPG default**.
- Step 23** Right-click the **Static EndPoint** folder and then choose **Create Static EndPoint**.
- Step 24** In the **Create Static Endpoint** dialog box, complete the following steps:
- In the **MAC** field, enter the ERSPAN destination's MAC address.
 - In the **Type** area, choose **tep**.
 - In the **Path Type** area, choose the appropriate path type.
If you choose Port as the path type, choose a node from the Node drop-down list.

The path type determines how the leaf is connected to the ERSPAN destination. The leaf can be connected by port, direct port channel, or virtual port channel.

- d) In the **Path** field, enter the appropriate path.

The path determines the policy group where the ERSPAN destination is attached.

- e) In the **IP Address** field, enter the ERSPAN destination IP address.
- f) In the **Encap** field, enter the overlay-1 VLAN.
- g) Click **Submit**.
- h) From the ERSPAN destination, ping any overlay- IP address.

This step ensures that the fabric learns the ERSPAN destination IP address.



CHAPTER 7

BPDU Features

This chapter contains the following sections:

- [Understanding Bridge Protocol Data Unit Features, on page 31](#)
- [Configure BPDU Features Using the GUI, on page 32](#)

Understanding Bridge Protocol Data Unit Features

The following sections describe supported bridge protocol data unit (BPDU) features on the Cisco ACI Virtual Edge with the Cisco APIC. BPDU Guard and BPDU filtering are switch-wide features, and they are applicable only for VM virtual Ethernet (vEth) ports.

BPDU Guard

BPDU Guard prevents loops by moving a nontrunking port into an errdisable state when a BPDU is received on that port. When you enable BPDU Guard on the switch, the interface is moved to blocking state on receiving a BPDU.

BPDU Guard provides a secure response to invalid configurations because the administrator must manually put the interface back in service. To put the interface back in service, disconnect the VM port and then reconnect it to the Cisco ACI Virtual Edge or an EPG port group through vCenter.

BPDU Filtering

BPDU filtering prevents sending and receiving of BPDUs on ports. Any BPDU that is received is dropped when filtering is enabled. BPDU filtering is enabled on VM vEth ports by default. When you enable this feature, Cisco ACI Virtual Edge drops all BPDUs received on uplink ports.



Note

We recommend that you configure BPDU policy in a single policy interface group. Configuring BPDU in multiple policy interface groups leads to inconsistent behavior.



Note In an L2 switch extended topology, we recommend that you configure BPDUs policy through an attached entity profile vSwitch policy override. If the interface policy group is used for configuration, then BPDUs Guard or filter is enabled on the Leaf ports. This causes those ports to become error-disabled when they receive BPDUs packets from an L2 switch.

For information about configuring BPDUs policy through an override policy, see the section "Modifying the Interface Policy Group to Override the vSwitch-Side Policies" in the *Cisco Application Virtual Edge Installation Guide*.

Configure BPDUs Features Using the GUI

- Step 1** Log in to the Cisco APIC.
- Step 2** On the menu bar, choose **Fabric > Access Policies**.
- Step 3** In the **Policies** navigation pane, expand the **Policies** and the **Interface** folders.
- Step 4** Right-click the **Spanning Tree Interface** folder and choose **Create Spanning Tree Interface Policy**.
- Step 5** In the **Create Spanning Tree Interface Policy** dialog, complete the following actions:
 - a) In the **Name** field, enter a name for the policy.
 - b) (Optional) In the **Description** field, enter a description of the policy.
 - c) In the **Interface controls** area, check the **BPDUs Guard enabled** check box or the **BPDUs filter enabled** check box.
 - d) Click **Submit** to save the policy.
- Step 6** Attach the spanning tree interface policy that you created in Step 5 by completing the following steps:
 - a) Go to **Virtual Networking > Inventory** and then expand the **VMM Domains** and **VMware** folders.
 - b) Click the VMM domain where you want to attach the policy.
 - c) Click the **VSwitch Policy** tab on the right side of the work pane.
 - d) From the **STP Policy** drop-down list, choose the policy that you created in Step 5.
 - e) Click **Submit**.
- Step 7** Verify the configuration by opening an ESXi CLI session to the ESXi hypervisor and entering the **vemcmd show card** command.

Example:

```
cisco-ave# vemcmd show card
Global BPDUs Guard: Enabled && Global BPDUs Filter: Enabled
```

The output indicates that BPDUs filtering and BPDUs Guard are enabled.



CHAPTER 8

IGMP Querier and Snooping

This chapter contains the following sections:

- [Guidelines and Limitations for Configuring IGMP Snooping and Querier, on page 33](#)
- [Configure IGMP Querier Using the GUI, on page 34](#)
- [Configure IGMP Snooping to Take Effect Immediately Using the GUI, on page 35](#)
- [Configure IGMP Snooping to Take Effect Later Using the GUI, on page 36](#)

Guidelines and Limitations for Configuring IGMP Snooping and Querier

Depending on your setup, you may need to configure IGMP on Layer 2 switches or on infra tenant or administrator-created tenant bridge domains. This section provides guidelines for two common scenarios when you must configure IGMP protocol snooping and querier.



Note Cisco ACI Virtual Edge does not support IGMP snooping. The guidelines and limitations and configuration procedures for IGMP snooping in this section are for configuring IGMP snooping on the leaf switch.

Multi-destination Flood for VXLAN-Encapsulated Traffic

To receive multi-destination flood on Cisco ACI Virtual Edge for VXLAN-encapsulated traffic and minimize multicast flooding traffic originating from and terminating on the Cisco ACI Virtual Edge if there is a Layer 2 device between the leaf and the Cisco ACI Virtual Edge, do the following:

- Apply IGMP snooping policy and enable IGMP querier on the infra tenant bridge domain subnet through the Cisco APIC. See the instructions in the section [Configure IGMP Querier Using the GUI, on page 34](#) in this guide.
- Enable IGMP snooping on each of any Layer 2 devices between the leaf and the Cisco ACI Virtual Edge. Follow the instructions that are specific to the device. For example, if the Layer 2 device is a Cisco Nexus 5000 Series switch, see the instructions in the configuration guide for that switch.

Sending or Receiving Multicast Streams with Virtual Machines

If you have virtual machines connected to the Cisco ACI Virtual Edge and want to send or receive multicast streams, do the following:

- Apply IGMP snoop policy and enable IGMP querier for administrator-created tenant bridge domain. If you have multiple administrator-created tenant bridge domains, you must apply IGMP snoop policy and configure IGMP querier on each administrator-created tenant bridge domain through the Cisco APIC. See the instructions in the section [Configure IGMP Querier Using the GUI, on page 34](#) in this guide.
- Enable IGMP snooping on each Layer 2 device between the leaf and the Cisco ACI Virtual Edge. Follow the instructions that are specific to the device. For example, if the Layer 2 device is a Cisco Nexus 5000 Series switch, see the instructions in the configuration guide for that switch.
- If the multicast traffic that originates from or terminates on the VMs is VXLAN-encapsulated, follow all the guidelines in the previous section as well as this one.

L3 Multicast

L3 Multicast is not supported on bridge domains with endpoints behind AVE. If L3 multicast is enabled, it alters the endpoint learning on the ACI physical leaf in such a way that it starts to learn endpoint IP addresses from IGMP reports coming from the endpoint. This limitation does not affect L2 multicast, as in the case of L2 multicast, the endpoint learning on the ACI physical leaf is not changed in such a way. In L2 multicast, the IGMP reports from the endpoints on the AVE will not impact the endpoint learning on the physical port.

Order of Configuration

Configure IGMP querier before you configure IGMP snooping.

Disabling IGMP Snoop

When Protocol Independent Multicast (PIM) is enabled on Cisco ACI, it is recommended to disable IGMP snoop on UCSM environment to avoid incorrect mac learn(s).

Configure IGMP Querier Using the GUI

Step 1 Log in to the Cisco APIC.

Step 2 Complete one of the following series of steps, depending on the type of tenant:

If you have ...	Then...
An infra tenant	<ol style="list-style-type: none"> Choose Tenants > infra. In the navigation pane, open the following folders: Networking > Bridge Domains > default > Subnets. Choose the subnet in the Subnets folder. In the Properties work pane, in the Subnet Control area, make sure that the Querier IP check box is checked. Click Submit.

If you have ...	Then...
An administrator-created tenant	<p>a. Choose Tenants and then choose the tenant on which you want to configure the IGMP querier.</p> <p>b. In the tenant navigation pane, open the Networking folder, the Bridge Domains folder, and then the folder for the bridge domain created earlier for the tenant.</p> <p>If the selected bridge domain already has a subnet with a gateway IP, you can use it to enable IGMP querier in the Subnet Control area. Or you can follow the remaining steps to create a new subnet to enable IGMP querier.</p> <p>c. Right-click the Subnets folder inside the bridge domain folder and choose Create Subnet.</p> <p>d. In the Create Subnet dialog box, complete the following steps:</p> <ol style="list-style-type: none"> Specify a gateway IP address. <p>Note You can configure any IP address except one from the 10.0.0.0/16 network because that network is reserved for Cisco APIC fabric devices.</p> <ol style="list-style-type: none"> In the Subnet Control area, make sure that the Querier IP check box is checked. Click Submit.

Configure IGMP Snooping to Take Effect Immediately Using the GUI

Step 1 Log in to the Cisco APIC.

Step 2 Take one of the following actions:

- If you have an infra tenant, choose **Tenants > infra**.
- If you have an administrator-created tenant, choose **Tenants** and then choose the tenant on which you want to configure the IGMP snooping.

Step 3 Take one of the following actions in the tenant navigation pane:

- If you have an infra tenant, open the **Networking** folder, open the **Bridge Domains** folder, and then choose the **default** folder.
- If you have an administrator-created tenant, open the **Networking** folder, open the **Bridge Domains** folder, and then choose the bridge domain created earlier for the tenant.

Step 4 In the **Properties** work pane, from the **IGMP Snoop Policy** drop-down list, choose **Create IGMP Snoop Policy**.

Step 5 In the **Create IGMP Snoop Policy** dialog box, complete the following steps:

- In the **Name** field, enter a name for the policy.
- In the **Control** area, check the **Enable querier** check box.
- (Optional) Configure any other relevant IGMP parameters.
- Click **Submit**.

Step 6 In the **Properties** pane, click **Submit**.

Configure IGMP Snooping to Take Effect Later Using the GUI

Step 1 Log in to the Cisco APIC.

Step 2 Take one of the following actions:

- If you have an infra tenant, choose **Tenants > infra**.
- If you have an administrator-created tenant, choose **Tenants** and then choose the tenant on which you want to configure the IGMP snooping.

Step 3 In the tenant navigation pane, open the **Policies** and **Protocol** folders.

Step 4 Right-click the **IGMP Snoop** folder and then choose **Create IGMP Snoop Policy**.

Step 5 In the **Create IGMP Snoop Policy** dialog box, complete the following steps:

- a) In the **Name** field, enter a name for the policy.
 - b) In the **Control** area, check the **Enable querier** check box.
 - c) (Optional) Configure any other relevant IGMP parameters.
 - d) Click **Submit**.
-

What to do next

Once you configure IGMP snooping, you can apply it at any time to a bridge domain by completing the following steps:

1. Take one of the following actions:
 - If you have an infra tenant, choose **Tenants > infra**.
 - If you have an administrator-created tenant, choose **Tenants** and then choose the tenant on which you want to configure the IGMP snooping.
2. Take one of the following actions in the **Tenant** navigation pane:
 - If you have an infra tenant, click the + icons to open the **Networking** and **Bridge Domain** folders, and then choose the **default** folder.
 - If you have an administrator-created tenant, open the **Networking** and **Bridge Domain** folders, and then choose the bridge domain created earlier for the tenant.
3. In the **Properties** pane, in the **IGMP Snoop Policy** drop-down list, choose the IGMP snooping policy that you want to apply.
4. Click **Submit** for the IGMP policy to go into effect for the bridge domain.



CHAPTER 9

vMotion with Cisco ACI Virtual Edge

This chapter contains the following sections:

- [Guidelines for Using VMware vMotion with Cisco ACI Virtual Edge, on page 37](#)

Guidelines for Using VMware vMotion with Cisco ACI Virtual Edge

You cannot move the Cisco ACI Virtual Edge VM with VMware vMotion, but you can move guest VMs on the same host with vMotion. Follow the guidelines in this section for using vMotion with guest VMs sharing the same host as Cisco ACI Virtual Edge.

vMotion Configuration

We recommend that you configure vMotion on a separate VMkernel NIC with a separate EPG that uses native switching mode.

Cross-VMware vCenter vMotion Support

Microsegmentation with Cisco ACI for Cisco ACI Virtual Edge is supported for cross-VMware vCenter and cross-VDS vMotion.



Note When you do a cross-VMware vCenter vMotion of endpoints, you may experience a few seconds of traffic loss.

Guidelines for Using Cross-VMware vCenter and Cross-VDS vMotion

- The source and destination VMware vCenter Server instances and ESXi hosts must be running version 6.0 or later.
- The source and destination vSphere Distributed Switch (VDS) version must be same.
- Refer to VMware documentation for prerequisites for cross-VDS and Cross-VMware VCenter vMotion.

vMotion support with Cisco ACI Virtual Edge

Cisco ACI Virtual Edge supports cross-VMware vCenter and cross-DVS when Distributed Firewall is not enabled on the Cisco ACI Virtual Edge domain. Be aware of the following vMotion limitations when Distributed Firewall is enabled:

Table 1: vMotion When Distributed Firewall is Enabled

Type of vMotion	Intra-VMM (Intra-DVS)	Inter-VMM (Cross-DVS)
Cross-vCenter	Supported	Supported
Single-vCenter	Supported	Supported
Cross Cisco ACI Multi-Site	Not supported	Not supported

Stale VM Entry After Cross-Data Center VMware vMotion

After you migrate VMs using cross-data center VMware vMotion in the same VMware vCenter, you may find a stale VM entry under the source DVS. This stale entry can cause problems, such as host removal failure. The workaround for this problem is to enable "Start monitoring port state" on the vNetwork DVS. See the KB topic "Refreshing port state information for a vNetwork Distributed Virtual Switch" on the VMware Web site for instructions.



CHAPTER 10

Intra-EPG Isolation Configuration

- [Intra-EPG Isolation Enforcement for Cisco ACI Virtual Edge, on page 39](#)
- [Configure Intra-EPG Isolation for Cisco ACI Virtual Edge Using the GUI, on page 40](#)
- [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge, on page 41](#)
- [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 41](#)
- [Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Virtual Networking Tab, on page 41](#)
- [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge, on page 42](#)
- [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 42](#)
- [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Virtual Networking Tab, on page 43](#)

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).

-
- 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.

- g) 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 41](#) and [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 42](#) in this guide.

Choose Statistics for Isolated Endpoints on Cisco ACI Virtual Edge

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.

-
- | | |
|---------------|--|
| Step 1 | Log in to Cisco APIC. |
| Step 2 | Choose Tenants > <i>tenant</i> . |
| Step 3 | In the tenant navigation pane, expand the Application Profiles , <i>profile</i> , 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 . |
-

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.

-
- 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

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 41](#) in this guide for instructions.

-
- 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.

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 41](#) in this guide for instructions.

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.

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



Distributed Firewall

This chapter contains the following sections:

- [About Distributed Firewall, on page 45](#)
- [Benefits of Distributed Firewall, on page 46](#)
- [Distributed Firewall Configuration, on page 47](#)
- [Distributed Firewall Flow Logging, on page 51](#)
- [Distributed Firewall Flow Counts, on page 57](#)

About Distributed Firewall

The Distributed Firewall is a hardware-assisted firewall. It supplements—but does not replace—other security features in the Cisco Application Centric Infrastructure (ACI) fabric such as Cisco Adaptive Security Virtual Appliance (ASAv) or secure zones created by Microsegmentation with Cisco ACI Virtual Edge.

No additional software is required for the Distributed Firewall to work. However, you must configure policies in the Cisco Application Policy Infrastructure Controller (APIC) to work with the Distributed Firewall.

The Distributed Firewall is supported on all Virtual Ethernet (vEth) ports but is disabled for kni-opflex, kni-ave-ctrl dpdk interfaces and for all uplink ports.

Key Features of the Distributed Firewall

Feature	Description
Provides dynamic packet filtering (also known as stateful inspection)	Tracks the state of TCP and FTP connections and blocks packets unless they match a known active connection. Traffic from the Internet and internal network is filtered based on policies that you configure in the APIC GUI.
Is distributed	Tracks connections even if you use vMotion to move virtual machines (VMs) to other servers.
Prevents SYN-ACK attacks	When the provider VM initiates SYN-ACK packets, the Distributed Firewall on the provider Cisco ACI Virtual Edge drops these packets because no corresponding flow (connection) is created.

Feature	Description
Supports TCP flow aging	Connections in ESTABLISHED state are maintained for 2 hours unless the per-port limit reaches the 75% threshold. Once that threshold is reached, any new connection can potentially replace the old connection (which has been inactive for at least 5 minutes). Connections in non-ESTABLISHED TCP state are retained for 5 minutes of idle or inactive time.
Is implemented at the flow level	Enables a flow between VMs over the TCP connection, eliminating the need to establish a TCP/IP connection for each packet.
Not dependent on any particular topology or configuration	Works with either Local Switching and No Local Switching modes and with either VLAN and VXLAN.
Is hardware-assisted	In the ACI fabric, Cisco Nexus 9000 leaf switches store the policies, avoiding impact on performance.
Bases implementation on 5-tuple values	Uses the source and destination IP addresses, the source and destination ports, and the protocol in implementing policies.
Is in learning mode by default	Facilitates upgrades. Distributed Firewall must be in learning mode if you migrate to Cisco ACI Virtual Edge from Cisco AVS if the version of Cisco AVS is earlier than Release 5.2(1)SV3(1.5). Those versions do not support Distributed Firewall.

Benefits of Distributed Firewall

This section provides examples of how Distributed Firewall works with hardware in the Cisco ACI fabric to provide security.

Enhanced Security for Reflexive ACLs

An administrator creates a contract using subjects and filters in the Cisco APIC between consumer and provider EPGs to allow web traffic. The administrator creates a policy in Cisco APIC to allow traffic from any source port to destination port 80.

When the policy is configured in Cisco APIC, a reflexive access control list (ACL) entry from the provider to the consumer is automatically programmed in the ACI hardware. This reflexive ACL is created to allow the reverse traffic for the time when a connection remains established. This reflexive ACL entry is necessary to allow the reverse traffic to flow.

Because of the automatic reflexive ACL creation, the leaf switch allows the provider to connect to any client port when the connection is in the established state. But this may not be desirable for some data centers. That is because an endpoint in a provider EPG may initiate a SYN attack or a port-scan to the endpoints in the consumer EPGs using its source port 80.

However, the Distributed Firewall, with the help of the physical hardware, will not allow such attack. The physical leaf hardware evaluates the packet it receives from the hypervisor against the policy ternary content addressable memory (TCAM) entry.

Protecting Data When VMs Are Moved with VMotion

Every packet sent or received follows the flow-based entry in the Distributed Firewall in the Cisco ACI Virtual Edge and in the physical leaf. Since the flows are directly attached to a virtual machine (VM) virtual Ethernet (vEth) interface, even when VMs are moved by VMotion to a different hypervisor host, the flows and table entries move with it to the new hypervisor.

This movement is also reported back to physical leaf. The physical leaf allows the legitimate flow to continue and prevents attacks if they occur. So even when the VM is moved to the new hosts, VM is still communicating without losing protection.

Seamless FTP Traffic Handling

The behavior and inter-working of the FTP protocol is different than other TCP-based protocols. For this reason, it requires special treatment in the Distributed Firewall. FTP Server (Provider) listens on the Control port (TCP port 21) and a Data port (TCP port 20). When communication begins between FTP client (Consumer) and server (Provider), the control connection is set up initially between the FTP client and server. The data connection is set up on demand (only when there is data to be exchanged) and torn down immediately after the data transfer.

Distributed Firewall supports only Active-FTP mode handling. The data connections are not tracked for the Passive-FTP mode.

Distributed Firewall allows the FTP data connection only if it matches the FTP Client IP and Port information that was received during the control connection handshake. Distributed Firewall blocks the FTP data connections if there is no corresponding control connection; this is what prevents FTP attacks.

Distributed Firewall Configuration

You configure Distributed Firewall by setting it to one of its three modes:

- Enabled—Enforces the Distributed Firewall.
- Disabled—Does not enforce Distributed Firewall. Use this mode only if you do not want to use the Distributed Firewall. Disabling Distributed Firewall removes all flow information on the Cisco ACI Virtual Edge.
- Learning—Cisco ACI Virtual Edge monitors all TCP communication and creates flows in a flow table but does not enforce the firewall. Learning is the default mode.

Distributed Firewall works with policies created in Cisco APIC; unless you create the policies, Distributed Firewall cannot work effectively.



Note We recommend that you use vmxnet3 adapters for the VMs when using Distributed Firewall.

**Important**

When Distributed Firewall is enabled on a Cisco ACI Virtual Edge VMM domain, there are restrictions on vMotion. See the section [Guidelines for Using VMware vMotion with Cisco ACI Virtual Edge](#), on page 37 in this guide for more information.

Workflow for Distributed Firewall Configuration

This section provides a high-level description of the tasks that you perform to configure Distributed Firewall.

1. Create an interface policy group to enable the firewall policy in the Cisco APIC, or, if you already have an interface policy group, make sure that it contains a firewall policy.
2. Configure a stateful policy for Distributed Firewall.

Follow instructions in the section [Configure a Stateful Policy for Distributed Firewall Using the GUI](#), on page 48 in this guide.

3. Change the Distributed Firewall mode if necessary.

Distributed Firewall is in learning mode by default. If you have not previously enabled Distributed Firewall, follow the instructions in this guide to for changing the Distributed Firewall mode.

4. Cisco ACI Virtual Edge reports the flows that are permitted or denied by Distributed Firewall to the system log (syslog) server. You can configure parameters for the flows and view the denied flows on the syslog server. See the instructions in the section [Distributed Firewall Flow Logging](#), on page 51 in this guide.
5. Choose which Distributed Firewall flow count statistics that you want to view.

Cisco ACI Virtual Edge collects Distributed Firewall flow information, but you must choose which statistics you want to know about before you can view them. See the section [Distributed Firewall Flow Counts](#), on page 57 in this guide for instructions.

Configure a Stateful Policy for Distributed Firewall Using the GUI

Before you can configure a Distributed Firewall policy, configure a stateful policy for Distributed Firewall.

- Step 1** Log in to the Cisco APIC.
- Step 2** Choose **Tenants > tenant**.
- Step 3** In the navigation pane, expand the folder for the tenant.
- Step 4** Right-click the **Contracts** folder and then choose **Create Contract**.
- Step 5** In the **Create Contract** dialog box, in the **Name** field, type a name for the contract.
- Step 6** In the **Subjects** area, click the + icon.
- Step 7** In the **Create Contract Subject** dialog box, in the **Name** field, type a name for the subject.
- Step 8** In the **Filter Chain** area, click the + icon next to **Filters**.
- Step 9** Click the down arrow to display the **Name** drop-down filter list, and then click the + icon at the top of the **Name** list.
- Step 10** In the **Create Filter** dialog box, complete the following actions:
 - a) In the **Name** field, type a name for the filter.
 - b) In the **Entries** area, click the + icon to display more fields.

- c) In the **Name** field, type a name to further describe the filter.
- d) From the **Ether Type** drop-down list, choose **IP**.
- e) From the **IP Protocol** field, choose **tcp**.
- f) Check the **Stateful** check box.
- g) (Optional) In the **Source Port / Range** field, from the **To** and the **From** drop-down lists, choose **Unspecified**, the default.
- h) In the **Destination Port / Range** field, from the **To** and the **From** drop-down lists, choose **http**.
- i) Click **Update** and then click **Submit**.

Step 11 In the **Create Contract Subject** dialog box, in the **Filters** area, click **Update** and then click **OK**.

Step 12 In the **Create Contract** dialog box, click **Submit**.

What to do next

Create a Distributed Firewall policy.

Create a Distributed Firewall Policy Using the GUI

You can create a Distributed Firewall policy using the Cisco APIC GUI.

Before you begin

You must have done the following:

- Created an interface policy group to enable the Distributed Firewall policy in Cisco APIC.
- Created a stateful policy for Distributed Firewall.

Step 1 Log in to the Cisco APIC.

Step 2 Go to **Fabric > Access Policies**.

Step 3 In the **Policies** navigation pane, expand the **Policies** and **Interface** folders.

Step 4 Right-click the **Firewall** folder and choose **Create Firewall Policy**.

Step 5 In the **Create Firewall Policy** dialog box, in the **Name** field, type a name for the policy.

Step 6 In the **Mode** area, choose a mode.

The default mode is Learning to facilitate upgrades.

Distributed Firewall must be in Learning mode if you migrate to Cisco ACI Virtual Edge from Cisco AVS if the version of Cisco AVS is earlier than Release 5.2(1)SV3(1.5). Those versions do not support Distributed Firewall.

Otherwise, enable Distributed Firewall.

Note Do not change the mode from Disabled directly to Enabled. Doing so can lead to traffic loss. Instead, from Disabled mode, change the mode to Learning, wait 5 minutes, and then change the mode to Enabled. The **Create Firewall Policy** dialog box includes a **Syslog** area. This is where you can configure the source for Distributed Firewall flow information that is sent to the syslog server. See the section [Distributed Firewall Flow Logging, on page 51](#) in this guide for instructions.

Step 7 Click **Submit**.

Step 8 Associate the new policy with the VMM domain by completing the following steps:

- a) Go to **Virtual Networking > Inventory**.
- b) In the **Inventory** navigation pane, expand the **VMM Domains** folder and the **VMware** folder, and then choose the relevant VMM domain.
- c) In the VMM domain work pane, click the **VSwitch Policies** tab.
- d) In the **Properties** work pane, from the **Firewall Policy** drop-down list, choose the firewall policy that you created.
- e) Click **Submit**.

What to do next

Verify that the Distributed Firewall policy is created and is in the desired state by completing the following steps:

1. Go to **Fabric > Access Policies**.
2. In the **Policies** navigation pane, expand the **Policies**, **Interface** , and **Firewall** folders.
3. Choose the policy.
4. In the **Properties** work pane, verify that the policy appears and that the mode is correct.

Change Distributed Firewall Policy Mode Using the GUI

Use the following procedure to change the Distributed Firewall mode.



Note Enable Distributed Firewall if you migrated from Cisco AVS to Cisco ACI Virtual Edge and did not have Distributed Firewall enabled for Cisco AVS.

Before you begin

Ensure that your Distributed Firewall policy is associated with a VMM domain.

Step 1 Log in to the Cisco APIC.

Step 2 Go to **Fabric > Access Policies**.

Step 3 In the **Policies** navigation pane, expand the **Policies**, **Interface**, and **Firewall** folders.

Step 4 Click the policy that you want to modify.

Step 5 In the **Properties** work pane, in the **Mode** area, choose a mode, and then click **Submit**.

Note Do not change the mode from Disabled directly to Enabled. Doing so can lead to traffic loss. Instead, from Disabled mode, change the mode to Learning, wait 5 minutes, and then change the mode to Enabled. Changing to Learning mode allows Cisco ACI Virtual Edge to add flow table entries for existing flows.

Note The **Properties** work pane includes a **Syslog** area where you can configure the source for Distributed Firewall flow information that is sent to the syslog server. See the section [Distributed Firewall Flow Logging](#), on page 51 in this guide for instructions.

What to do next

Verify that the Distributed Firewall is in the desired state by completing the following steps:

1. In the **Policies** navigation pane, choose the policy in the **Firewall** folder.
2. In the **Properties** dialog box, verify that the mode is correct.

Distributed Firewall Flow Logging

You can view flow information for Distributed Firewall with the Cisco APIC to assist with auditing network security.

Cisco ACI Virtual Edge reports the flows that are denied and permitted by Distributed Firewall to the system log (syslog) server. When you enable Distributed Firewall, Cisco ACI Virtual Edge monitors TCP, UDP, and ICMP traffic by default. It also tracks, logs, and—depending on how you configure parameters—permits or denies TCP traffic. You can view the denied and permitted flows on the syslog server.

Parameters for Distributed Firewall Flow Information

Cisco ACI Virtual Edge reports the flows that are denied or permitted by Distributed Firewall as well UDP and ICMP flows to the system log (syslog) server.

You configure Distributed Firewall logging in two tasks: configuring up to three syslog servers, referred to as remote destinations in the GUI, and configuring the syslog policy. You can configure the following parameters:

- Syslog server parameters
 - Enable/disable



Note Distributed Firewall logging is disabled by default.

- Permitted flows, Denied flows, or both
- Polling interval

You can set the interval for exporting the flows from 60 seconds to 24 hours.



Note A polling interval of 125 seconds is required to send data at maximum scale. We recommend that you configure the syslog timer with a polling interval of at least 150 seconds.

- Log severity

You can set the severity level from 0-7.

- Syslog policy parameters

- IP address

- Port

- Log severity

You can set the severity level from 0-7.

- Log facility

Cisco ACI Virtual Edge reports up to 250,000 denied or permitted flows to the syslog server for each polling interval. If you choose to log denied and permitted flows, Cisco ACI Virtual Edge reports up to 500,000 flows. Cisco ACI Virtual Edge also reports up to 100,000 short-lived flows—flows that are shorter than the polling interval.

Syslog messages are sent only if the syslog destination log severity is at or below the same log severity for the syslog policy. Severity levels for the syslog server and syslog policy are as follows:

- 0: Emergency
- 1: Alert
- 2: Critical
- 3: Error
- 4: Warning
- 5: Notification
- 6: Information
- 7: Debug

Guidelines for Configuring the Syslog Server

Follow the guidelines in this section when configuring the syslog server for Cisco ACI Virtual Edge.

- The syslog server should always be reachable from the Cisco ACI Virtual Edge host management network or Cisco ACI Virtual Edge infra port group (overlay-1 vrf of tenant infra).

If the syslog server is behind the Cisco ACI Virtual Edge, bring up the VM VNIC in the infra port group.

- The syslog server should always be on a different host from Cisco ACI Virtual Edge.

Sending log messages from a Cisco ACI Virtual Edge to a syslog server hosted behind the same Cisco ACI Virtual Edge is not supported.

- If the syslog server destination is a VM, make sure that vMotion is disabled on it. If the syslog server destination VM is moved to another host for any reason, make sure that the static client endpoint (CEP) is configured accordingly. See the section [Configure a Static End Point Using the GUI](#), on page 54 in this guide.

The IP for the syslog server can be obtained using DHCP (Option 61 is needed during DHCP) or static configuration. Make sure that the IP address is in the same subnet as the other EPs in infra port group (overlay-1 VRF of tenant infra).

Distributed Firewall Flow Syslog Messages

This section provides the formats and examples of syslog messages for distributed Firewall flows.

- Denied flows

- Format

```
<Syslog Server timestamp> < PRI = Facility*8 + Severity > <syslog version> <Host
timestamp> <Host IP> <Application name (ave-dfwlog)> - AVE IP: <AVEIP> AVE Hostname
<hostname> DFWLOG-DENY_FLOW - <Deny Reason> AVE UUID: <UUID>, Source IP: <Source
IP address>, Destination IP: <Destination IP address> , Source Port: <Port number>,
Destination Port: <Port Number>, Source Interface: <Interface name>, Protocol:
"TCP"(6), Hit-Count = <Number of Occurrences>, EPG Name: <EPG Name>, EpP DN: <EpP
DN>
```

- Example

```
Thu Apr 21 14:36:45 2016 10.197.139.205 <62>1 2017-12-06T18:58:30.835 10.197.139.205
ave-dfwlog - AVE IP: 10.197.139.205 AVE Hostname localhost, DFWLOG-DENY_FLOW -
SYN ACK ingress AVE UUID: 42094298-4996-60EF-CE86-E2B7FC70C2EC, Source IP: 54.0.0.5,
Destination IP: 54.0.0.6, Source Port: 53535, Destination Port: 5555, Source
Interface: 00:50:56:89:4d:3e, Protocol: "TCP"(6), Hit-Count = 1, EPG Name =
Tenant1|AP-1|EPG-54, EpP DN: uni/epf/fv-[uni/tn-Tenant1/ap-AP-1/epg-EPG-54]
```

- Permitted flows

- Format

```
<Syslog server timestamp> < PRI = Facility*8 + Severity> <syslog version> <Host
timestamp> <Host IP> <Application name (ave-dfwlog)> - AVE IP: <AVEIP> AVE Hostname
<hostname> DFWLOG-PERMIT_FLOW -<flow status> AVE UUID: <UUID>, Source IP: <Source
IP address>, Destination IP: <Destination IP address>, Source Port: <Port Number>,
Destination Port: <Port Number>, Source Interface: <Interface name>, Protocol:
"TCP"(6), Age = <Age in seconds>, EPG Name: <EPG Name>, EpP DN: <EpP DN>
```

- Example

```
Tue Apr 19 19:31:21 2016 10.197.139.205 <62>1 2017-12-06T18:45:13.458 10.197.139.205
ave-dfwlog - AVE IP: 10.197.139.205 AVE Hostname localhost, DFWLOG-PERMIT_FLOW -
ESTABLISHED AVE UUID: 42094298-4996-60EF-CE86-E2B7FC70C2EC, Source IP: 54.0.0.5,
Destination IP: 54.0.0.6, Source Port: 59846, Destination Port: 5001, Source
Interface: 00:50:56:89:4d:3e, Protocol: "TCP"(6), Age = 0, EPG Name =
Tenant1|AP-1|EPG-54, EpP DN: uni/epf/fv-[uni/tn-Tenant1/ap-AP-1/epg-EPG-54]
```

- Short-lived permitted flows

- Format

```
<Syslog Server timestamp> < PRI = Facility*8 + Severity > <syslog version> <Host
timestamp> <Host IP> <Application name (ave-dfwlog)> - AVE IP: <AVEIP> AVE Hostname
<hostname> DFWLOG-PERMIT_SHORT_LIVED - <State of flow> AVE UUID: <UUID>, Source
IP: <Source IP address>, Destination IP: <Destination IP address>, Source Port:
<Port Number>, Destination Port: <Port Number>, Source Interface: <Interface Name>,
Protocol: "TCP"(6), Timestamp = <Host Timestamp>, EPG Name: <EPG Name>, EpP DN:
<EpP DN>
```

- Example

```
Thu Apr 21 14:46:38 2016 10.197.139.205 <62>1 2017-12-06T18:59:37.702 10.197.139.205
ave-dfwlog - AVE IP: 10.197.139.205 AVE Hostname localhost,
DFWLOG-PERMIT_SHORT_LIVED - CLOSED AVE UUID: 42094298-4996-60EF-CE86-E2B7FC70C2EC,
Source IP: 54.0.0.5, Destination IP: 54.0.0.6, Source Port: 59847, Destination
Port: 5001, Source Interface: 00:50:56:89:4d:3e, Protocol: "TCP"(6), Timestamp =
2017-12-06T18:59:37.702, EPG Name = Tenant1|AP-1|EPG-54, EpP DN:
uni/epp/fv-[uni/tn-Tenant1/ap-AP-1/epg-EPG-54]
```

- ICMP monitored flows

- Format

```
<Syslog server timestamp> < PRI = Facility*8 + Severity> <syslog version> <Host
timestamp> <Host IP> <Application name (ave-dfwlog)> - AVE IP: <AVEIP> AVE Hostname
<hostname>
DFWLOG-ICMP_TRACKING - AVE UUID: <UUID>, Source IP: <Source IP address>,
Destination IP: <Destination IP address>, Type:<ICMP type field>, Source Interface:
<Interface name>, Protocol: "ICMP"(1), Timestamp= <Host time stamp>, Direction:
<Egress/Ingress>, EPG Name:<EPG Name>, EpP DN: <EpP DN>
```

- Example

```
2016-11-28 11:02:43 News.Info 10.197.139.205 2017-12-06T19:01:05.061 10.197.139.205
ave-dfwlog - AVE IP: 10.197.139.205 AVE Hostname localhost, DFWLOG-ICMP_TRACKING
AVE UUID: 42094298-4996-60EF-CE86-E2B7FC70C2EC, Source IP: 54.0.0.5, Destination
IP: 54.0.0.6, Icmp type and code: Echo request (8,0) Source Interface:
00:50:56:89:4d:3e, Protocol: "ICMP"(1), Timestamp = 2017-12-06T19:01:05.061,
Direction: Ingress, EPG Name = Tenant1|AP-1|EPG-54, EpP DN:
uni/epp/fv-[uni/tn-Tenant1/ap-AP-1/epg-EPG-54]
```

- UDP monitored flows

- Format

```
UDP:
<Syslog server timestamp> < PRI = Facility*8 + Severity> <syslog version> <Host
timestamp> <Host IP> <Application name (ave-dfwlog)> - AVE IP: <AVEIP> AVE Hostname
<hostname> DFWLOG-UDP_TRACKING - AVE UUID: <UUID>, Source IP: <Source IP address>,
Destination IP: <Destination IP address>, Source Port: <Port Number>, Destination
Port: <Port Number>, Source Interface: <Interface name>, Protocol: "UDP"(17),
Timestamp=<Host timestamp>, Direction: <Egress/Ingress>, EPG Name: <EPG Name>
```

- Example

```
2016-11-28 11:00:23 News.Info 10.197.139.205 1 2017-12-06T19:01:46.785 10.197.139.205
ave-dfwlog - AVE IP: 10.197.139.205 AVE Hostname localhost, DFWLOG-UDP_TRACKING
AVE UUID: 42094298-4996-60EF-CE86-E2B7FC70C2EC, Source IP: 55.0.0.253, Destination
IP: 55.0.0.5, Source Port: 67, Destination Port: 68, Source Interface:
00:50:56:00:55:05, Protocol: "UDP"(17), Timestamp = 2017-12-06T19:01:46.785,
Direction: Egress, EPG Name = Tenant1|AP-1|EPG-55, EpP DN:
uni/epp/fv-[uni/tn-Tenant1/ap-AP-1/epg-EPG-55]
```

Configure a Static End Point Using the GUI

- Step 1** Log in to Cisco APIC.
- Step 2** In the **Tenant infra** navigation pane, open the following folders: **Application Profiles** > **access** > **Application EPGs** > **default**.
- Step 3** Right-click the **Static EndPoint** folder and then choose **Create Static EndPoint**.

- Step 4** In the **Create Static Endpoint** dialog box, complete the following steps:
- In the **MAC** field, enter the syslog server destination's MAC address.
 - In the **Type** area, choose **tep**.
 - In the **Path Type** area, choose the appropriate path type.
The path type determines how the leaf is connected to the syslog server destination. The leaf can be connected by port, direct port channel, or virtual port channel.
 - If you chose **Port** as the **Path Type**, choose a node from the **Node** drop-down list.
 - In the **Path** field, enter the appropriate path.
The path determines the policy group where the syslog server destination is attached.
 - In the **IP Address** field, enter the syslog server destination IP address.
 - In the **Encap** field, enter the overlay-1 VLAN (vlan-xxix).
 - Click **Submit**.
- Step 5** From the syslog server destination, ping any overlay-IP address—for example, 10.0.0.30.
This step ensures that the fabric learns the Syslog server destination IP address.
-

Configure Parameters for Distributed Firewall Flow Information Using the GUI

To configure parameters, you first configure the parameters for the syslog server or servers and then configure the parameters for the syslog policy. The syslog server is referred to as the *Remote Destination* in the GUI.

Before you begin

You must have Distributed Firewall enabled.

-
- Step 1** Log in to Cisco APIC.
- Step 2** Go to **Admin > External Data Collectors**.
- Step 3** In the **External Data Collectors** navigation pane, expand the **Monitoring Destinations** folder and then choose the **Syslog** folder.
- Step 4** In the **Syslog** work pane, click the **ACTIONS** down arrow and then choose **Create Syslog Monitoring Destination Group**.
- Step 5** In the **Create Syslog Monitoring Destination Group STEP 1 > Profile** dialog box, complete the following steps:
- In the **Define Group Name and Profile** area, enter a name in the **Name** field.
 - In the **Admin State** area, make sure that **enabled** is chosen from the drop-down list.
 - Accept the defaults in the rest of the dialog box and click **NEXT**.
- Step 6** In the **Create Syslog Monitoring Destination Group STEP 2 > Remote Destinations** dialog box, click the + icon.
- Step 7** In the **Create Syslog Remote Destination** dialog box, complete the following steps:
- In the **Host** field, enter the host IP address.
 - In the **Name** field, enter the host name.
 - In the **Admin State** area, make sure that **enabled** is chosen.
 - In the **Format** area, make sure that **aci** is chosen.
 - From the Severity drop-down list, choose a severity.

- f) From the **Port** drop-down list, accept the standard port unless you are using another port.
- g) From the **Forwarding Facility** drop-down list, choose a facility.
- h) Ignore the **Management EPG** drop-down list and click **OK**.

Step 8 (Optional) In the **Create Syslog Monitoring Destination Group STEP 2 > Remote Destinations** dialog box, create up to two additional remote destinations.

Step 9 In the **Create Syslog Monitoring Destination Group STEP 2 > Remote Destinations** dialog box, click **FINISH**. The newly created destination appears in the **Syslog** folder in the **External Data Collectors** navigation pane.

Step 10 Choose **Fabric > Access Policies**.

Step 11 In the **Policies** navigation pane, open the **Polices** and **Interface** folders.

Step 12 Complete one of the following sets of steps:

If you want to...	Then...
Configure a syslog policy with a new Distributed Firewall policy	<ul style="list-style-type: none"> a. Right-click the Firewall folder and choose Create Firewall Policy. b. In the Create Firewall Policy dialog box, in the Specify the Firewall Policy Properties area, type a name for the policy in the Name field. c. In the Mode area, choose a mode. Distributed Firewall must be in Learning mode if you migrate to Cisco ACI Virtual Edge from Cisco AVS if the version of Cisco AVS is earlier than Release 5.2(1)SV3(1.5). Those versions do not support Distributed Firewall. d. In the Syslog area, make sure that enabled is chosen from the Administrative State drop-down list. e. From the Included Flows area, choose Permitted flows, Denied flows, or both. f. In the Polling Interval (seconds) area, choosing an interval from 60 seconds to 24 hours. g. From the Log Level drop-down list, choose a severity level. The logging severity level should be greater than or equal to severity level defined for the syslog server. See the section Parameters for Distributed Firewall Flow Information, on page 51 in this guide for information about severity. h. From the Dest Group drop-down list, choose the destination group that you just created. i. Click Submit. j. Go to the section "What To Do Next" and associate the new Distributed Firewall policy with a VMM domain.
Configure a syslog policy with an existing Distributed Firewall policy	<ul style="list-style-type: none"> a. Expand the Firewall folder and choose the Distributed Firewall policy that you want to modify. b. In the policy work pane, change the Mode if desired. Distributed Firewall must be in Learning mode if you migrate to Cisco ACI Virtual Edge from Cisco AVS if the version of Cisco AVS is earlier than Release 5.2(1)SV3(1.5). Those versions do not support Distributed Firewall. c. In the Syslog area, make sure that enabled is chosen from the Administrative State drop-down list.

If you want to...	Then...
	<p>d. From the Included Flows area, choose Permitted flows, Denied flows, or both.</p> <p>e. In the Polling Interval (seconds) area, choosing an interval from 60 seconds to 24 hours.</p> <p>f. From the Log Level drop-down list, choose a severity level.</p> <p>The logging severity level should be greater than or equal to severity level defined for the syslog server. See the section Parameters for Distributed Firewall Flow Information, on page 51 in this guide for information about severity.</p> <p>g. From the Dest Group drop-down list, choose the destination group that you just created.</p> <p>h. Click Submit.</p> <p>i. If you see the Policy Usage Warning dialog box, click SUBMIT CHANGES.</p>

What to do next

If you configured a syslog policy with a new Distributed Firewall policy, you must associate the Distributed Firewall policy with a VMM domain.

1. In Cisco APIC, choose **Virtual Networking > Inventory**.
2. In the navigation pane, expand the **VMM Domains** folder and the **VMware** folder, and then choose the relevant VMM domain.
3. In the work pane, click the **VSwitch Policy** tab under the **Policy** tab.
4. In the **Create VSwitch Policy Container** dialog box, click **Yes**.
5. In the work pane, from the **Firewall Policy** drop-down list, choose the policy.
6. Click **Submit**.
7. If you see the **Policy Usage Warning** dialog box, click **SUBMIT CHANGES**.

Distributed Firewall Flow Counts

You can view Distributed Firewall flow counts with the Cisco APIC.

Cisco ACI Virtual Edge collects Distributed Firewall flow information, but you must choose which statistics you want to know about before you can view them. You can choose a sampling interval with choices ranging from 10 seconds to 1 year; however, the default is 5 minutes.

You can choose statistics and view them from two different places in Cisco APIC: one beginning with **Virtual Networking** and one beginning with **Tenants**. However, the steps for choosing and viewing statistics are the same.

When you choose statistics in Cisco APIC, you see a list of different kinds of statistics, but only nine are relevant to Distributed Firewall:

- **aged connections (connections)**

- **created connections (connections)**
- **destroyed connections (connections)**
- **denied global input connections (connections)**
- **denied per port limit connections (connections)**
- **invalid SYN ACK packets (packets)**
- **invalid SYN packets (packets)**
- **invalid connection packets (packets)**
- **invalid ftp SYN packets (packets)**

Choose Statistics to View for Distributed Firewall

Before you begin

You must have Distributed Firewall enabled.

-
- Step 1** 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 2** Click the **Stats** tab.
- Step 3** Click the tab with the check mark.
- Step 4** 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 5** (Optional) Choose a sampling interval.
- Step 6** Click **Submit**.
-

View Statistics for Distributed Firewall

Once you have chosen statistics for Distributed Firewall, you can view them.

Before you begin

You must have chosen statistics to view for Distributed Firewall.

-
- Step 1** Choose **Virtual Networking > Inventory > VMware > VMM Domains > VMM_name > Controllers > controller instance name > DVS-VMM name > Portgroups > EPG_name > Learned Point MAC address (Node)**.
- Step 2** 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.
-



CHAPTER 12

Microsegmentation with Cisco ACI

- [Microsegmentation with Cisco ACI, on page 59](#)

Microsegmentation with Cisco ACI

You can use Cisco APIC to configure Microsegmentation with Cisco ACI. Microsegmentation gives you the ability to assign endpoints to special endpoint groups, or EPGs, based on various attributes. These attribute-based EPGs are called microsegments and function as logical security zones because you can apply filtering and forwarding policies to them.

See the chapter "Microsegmentation with Cisco ACI" in the [Cisco ACI Virtualization Guide](#) for information about using and configuring microsegmentation.



CHAPTER 13

Attachable Entity Profile Configuration

This chapter contains the following sections:

- [Configuring an Attachable Entity Profile Using the GUI, on page 61](#)

Configuring an Attachable Entity Profile Using the GUI

The Cisco ACI fabric provides multiple attachment points that connect through leaf ports to various external entities such as bare metal servers, virtual machine hypervisors, Layer 2 switches, or Layer 3 routers. These attachment points can be physical ports, FEX ports, port channels, or a virtual port channel on leaf switches.

An Attachable Entity Profile (AEP) represents a group of external entities with similar infrastructure policy requirements. The infrastructure policies consist of physical interface policies that configure various protocol options.

An AEP is required to deploy VLAN pools on leaf switches. Encapsulation blocks (and associated VLANs) are reusable across leaf switches. An AEP implicitly provides the scope of the VLAN pool to the physical infrastructure. See the [Cisco ACI Fundamentals Guide](#) for detailed information about AEPs.

-
- Step 1** Log in to the Cisco APIC.
- Step 2** On the menu bar, choose **Fabric > Access Policies**.
- Step 3** In the left **Policies** navigation pane, expand the **Global Policies** folder.
- Step 4** Right-click the **Attachable Access Entity Profiles** folder and choose **Create Attachable Access Entity Profile**.
- Step 5** In the **Create Attachable Access Entity Profile** dialog box, complete the following steps:
- In the **Name** field, enter a name.
 - Check the **Enable Infrastructure VLAN** check box.
 - In the **Domains (VMM, Physical or External) To Be Associated To Interfaces** area, click the + icon.
 - From the **Domain Profile** drop-down list, choose a domain profile (VMM domain).
 - Click **Update** to update the domains.
 - Click **Next**.
- Step 6** In the **Create Attachable Access Entity Profile** dialog box, complete the following steps:
- Choose the interface policy groups that you created for your hosts.
 - For each interface policy group that you choose, choose **All** or **Specific**.

If you choose **All**, the attached entity applies to all interfaces associated with the policy group. If you choose **Specific**, you choose a switch ID from the **Switch IDs** drop-down list that appears to the right of the interface policy group list.

- c) Click **Finish**.
-



CHAPTER 14

Layer 4 to Layer 7 Services

- [Layer 4 to Layer 7 Services, on page 63](#)
- [Guidelines and Limitations for Layer 4 to Layer 7 Configuration, on page 63](#)
- [Qualified Service Devices, on page 66](#)
- [Supported Deployments, on page 66](#)
- [Bridge Domain Configuration for Cisco ASAV, Citrix NetScaler, or F5 BIG-IP ADC, on page 67](#)

Layer 4 to Layer 7 Services

Cisco Application Centric Infrastructure (ACI) treats services as a key part of an application. Any services that are required are treated as a service graph that is instantiated on the Cisco ACI fabric from the Cisco Application Policy Infrastructure Controller (APIC). You define the service for the application, while service graphs identify the set of network or service functions that the application requires.

Beginning with Cisco ACI Virtual Edge Release 1.2(1), Layer 4 to Layer 7 service graphs are supported for Cisco ACI Virtual Edge.

Beginning with the Cisco ACI Virtual Edge Release 2.2(1), support for Layer 4 to Layer 7 service graphs is extended to Cisco ACI Virtual Edge when it is part of Cisco ACI Virtual Pod.

For information about configuring Layer 4 to Layer 7 services on Cisco ACI Virtual Edge, see the *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*. However, you first must follow the guidelines and understand the limitations in the next section of this chapter.

When you follow instructions in the *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*, instead of configuring services on the VMware Distributed Virtual Switch (DVS) VMM domain, configure the services on the Cisco ACI Virtual Edge VMM domain with **AVE** as the switching mode.

Guidelines and Limitations for Layer 4 to Layer 7 Configuration

Follow the guidelines and note the limitations in this section when preparing to configure Layer 4 to Layer 7 service graphs for Cisco Application Centric Infrastructure (ACI) Virtual Edge.



Note

The guidelines and limitations differ for Cisco ACI Virtual Edge when it is part of Cisco ACI Virtual Pod (vPod) and when it is not part of Cisco ACI vPod.

Guidelines and Limitations When Cisco ACI Virtual Edge Is Not Part of Cisco ACI vPod

Follow the guidelines in this section when preparing to configure Layer 4 to Layer 7 service graphs for Cisco Application Centric Infrastructure (ACI) Virtual Edge when it is not part of Cisco ACI Virtual Pod (vPod).



Note For information about Layer 4 to Layer 7 guidelines and limitations for Cisco ACI Virtual Edge when it is part of Cisco ACI Virtual Pod, see the section [Guidelines and Limitations for Cisco ACI Virtual Edge When It Is Part of Cisco ACI vPod, on page 65](#) in this guide.

- Layer 4 to Layer 7 services is supported in routed mode with policy-based redirect (PBR); there is no support for transparent mode.

- Do not deploy both service VMs of an HA pair behind the same Cisco ACI Virtual Edge.

To ensure that both service VMs of an HA pair do not end up behind the same Cisco ACI Virtual Edge after deployment, create a VM-host affinity rule. The rule ensures that each service VM of an HA pair runs on different hosts.

When creating VM-host affinity rule, for **Type**, choose **Virtual Machines to Hosts** and in DRS groups, choose **Must run on hosts in group**. For more information about creating a VM-host-affinity rule, refer to VMware documentation for the corresponding vSphere version.

- Do not manually associate non-service VMs to a service EPG. At any point on a single host, only one endpoint for each service EPG is supported.
- Do not tag service VM interfaces deployed on Cisco ACI Virtual Edge; Cisco ACI Virtual Edge does not support trunk port groups.
- Virtual MAC-based service VM deployment is not supported for Cisco ACI Virtual Edge when it is not part of Cisco ACI Virtual Pod.
- The supported modes of service VM deployment on Cisco ACI Virtual Edge are standalone and HA mode (active/standby).

If you use Citrix NetScaler LoadBalancer, 1-Arm mode is also supported.

- Cisco ACI Virtual Edge supports vMotion of service VMs.



Note Refer to the corresponding vendor documentation for support of vMotion of service VMs on the VMware environment. The vMotion support is vendor-specific and may have certain guidelines and limitations.

- Only service-graph based deployments are supported on Cisco ACI Virtual Edge.
- Cisco ACI Virtual Edge does not support Route-Peering and Trunking Port.
- You cannot migrate Layer 4 to Layer 7 services deployed on a Cisco Application Virtual Switch (AVS) domain to Cisco ACI Virtual Edge.

To proceed with migration, undeploy services on Cisco AVS. Also, while migrating from a VMware VDS Domain to Cisco ACI Virtual Edge, you can move the consumer and provider EPGs to Cisco ACI

Virtual Edge. However, Layers 4 to Layer 7 service EPGs still belong to the VMware VDS. For more information, see the chapter "Migration from VMware VDS to Cisco ACI Virtual Edge" of the [Cisco ACI Virtual Edge Installation Guide](#).

- Ensure that the management and HA interfaces of service VMs are not connected to the Cisco ACI Virtual Edge port group.
- When you configure the Cisco ACI Virtual Edge VMM domain, it is mandatory to associate a VLAN pool with the domain.

You must associate a VLAN pool with the domain because service VMs are deployed on the Cisco ACI Virtual Edge VMM domain with VLAN encapsulation mode. Configure both internal and external ranges for the VLAN pool. See the chapter [Mixed-Mode Encapsulation, on page 13](#) in this guide for information.

- You can deploy compute VMs (providers and consumers) in the Cisco ACI Virtual Edge VMM domain with VXLAN or VLAN encapsulation mode.

To support compute VMs in either mode, configure the Cisco ACI Virtual Edge VMM domain with mixed-mode encapsulation. See the chapter [Mixed-Mode Encapsulation, on page 13](#) in this guide for information.

Guidelines and Limitations for Cisco ACI Virtual Edge When It Is Part of Cisco ACI vPod

Follow the guidelines in this section when configuring Layer 4 to Layer 7 service graphs for Cisco Application Centric Infrastructure Virtual Edge when it is part of Cisco ACI Virtual Pod (vPod).

- Layer 4 to Layer 7 services are supported using only Layer 3 policy-based redirect (PBR) is supported.

For information about PBR, see the chapter "Configuring Policy-based Redirect" in the [Cisco APIC Layer 4 to Layer 7 Services Deployment Guide](#).

- Do not deploy both service VMs of an HA pair behind the same Cisco ACI Virtual Edge.

Create a VM-host affinity rule. The rule ensures that both service VMs of an HA pair do not end up behind the same Cisco ACI Virtual Edge after deployment. The rule also ensures that each service VM of an HA pair runs on different hosts.

When creating VM-host affinity rule, for **Type**, choose **Virtual Machines to Hosts** and in DRS groups, choose **Must run on hosts in group**. For more information about creating a VM-host-affinity rule, refer to VMware documentation for the corresponding vSphere version.

- Do not manually associate non-service VMs to a service EPG. At any point on a single host, only one endpoint for each service EPG is supported.
- Do not tag service VM interfaces deployed on Cisco ACI Virtual Edge; Cisco ACI Virtual Edge does not support trunk port groups.
- Virtual MAC-based service VM deployment is supported for Cisco ACI Virtual Edge when it is deployed inside Cisco ACI Virtual Pod.

When creating a device policy for Citrix NetScaler Load Balancer in HA mode on the Cisco Application Policy Infrastructure Controller (APIC), enable Promiscuous Mode. The virtual MAC configuration on the service device requires that Promiscuous Mode be enabled.

- The supported modes of service VM deployment on Cisco ACI Virtual Edge are standalone, HA mode (active/standby), and 1-arm mode (in case of Citrix Netscaler LoadBalancer).
- Cisco ACI Virtual Edge supports VMware vMotion of service VMs.
- Only service-graph based deployments are supported on Cisco ACI Virtual Edge.
- Cisco ACI Virtual Edge does not support Route-Peering or Trunking Port.
- For Cisco ACI Virtual Edge used in the cloud, equal-cost multipath routing (ECMP) and health group features are not supported. This is the case although you may be able to configure those features using service graphs.
- You cannot migrate service VMs from Cisco ACI vPod to a physical pod or service VMs from a physical pod to a Cisco ACI vPod.
- Layer 4-Layer 7 PBR policy does not have pod awareness.
- When Citrix NetScaler service is deployed in HA mode, VMware vMotion is not supported between Cisco ACI vPod and the physical pod.
- If you use Citrix NetScaler Load Balancer in HA mode, non-vMAC switchover cannot be done, because listing multiple redirect entries in the redirect policy is not supported.
- VMware vMotion of ASAv devices is not supported after one switchover. VMware vMotion of ASAv devices is supported only when operational ACTIVE is same as configured ACTIVE device.

Qualified Service Devices

Service graph deployments for Cisco Application Centric Infrastructure (ACI) Virtual Edge are qualified for the following service devices:

- Cisco Adaptive Security Virtual Appliance (ASAv) firewall Version 9.9(1)



Note Before you deploy ASAv on the Cisco ACI Virtual Edge VMM domain, enable monitoring of `externalIf` and `internalIf`. To enable monitoring through the CLI, you can use the commands **`monitor-interface externalIf`** and **`monitor-interface internalIf`** on ASAv.

- INDUS: Revised at Pooja's request. Originally: "Citrix NetScaler VPX (Unmanaged mode) Version 11.0 build 70.16." ~ catortiz 8/16/19. Citrix NetScaler VPX Version 11.1 build 48.10nc.

Supported Deployments

The Cisco Application Centric Infrastructure (ACI) Virtual Edge supports the following deployments:

- ASAv in Routed Mode
- F5 Networks BIG-IP load balancer (Unmanaged mode), Cisco ACI Virtual Edge when not used with Cisco ACI Virtual Pod (vPod)

- One-arm mode
- Two-arm mode
- Citrix NetScaler VPX Version 11.1 build 48.10nc for Cisco ACI Virtual Edge when used with Cisco ACI vPod
- Standalone and HA mode (active/standby)
- One-arm and two-arm deployment modes

Bridge Domain Configuration for Cisco ASAV, Citrix NetScaler, or F5 BIG-IP ADC

When you configure the bridge domains for Cisco ASAv, Citrix NetScaler, or F5 BIG-IP ADC, configure the bridge domains as you do for a generic configuration, except as follows:

Configuration	Action
L2 Unknown Unicast	Choose Flood .
ARP Flooding check box	Check the check box.
Unicast Routing check box	This configuration depends on deployment. For example, put a check in the Unicast Routing check box if you want the Cisco ACI fabric to route the traffic. Also, when configuring the inside bridge domain, enable Unicast Routing if you plan to use endpoint attach.

References

For more information on configuring Bridge domains on Cisco ACI, see the [Cisco APIC Layer 2 Networking Configuration Guide](#).

For general information about bridge domain setting that is related to service graph design, see [Service Graph Design with Cisco Application Centric Infrastructure White Paper](#).



CHAPTER 15

Intrusion Detection System

- [IDS Overview, on page 69](#)
- [Guidelines and Limitations for IDS, on page 69](#)
- [IDS Check, on page 69](#)

IDS Overview

ACI Virtual Edge (AVE) provides IPv4 Intrusion Detection System (IDS) packet checks to increase security in the network by dropping packets that match specific criteria that are typically not required in most production networks. IDS packet checks are enabled by default and should be left enabled unless there is a specific reason to disable them.

Guidelines and Limitations for IDS

This section describes the guidelines and limitations for IDS:

- IDS can be disabled by logging into to AVE and run the `vemcmd set ids disable` command.



Note This is not persistent upon AVE reboot.

There is no knob in the APIC GUI to turn this feature on or off.

IDS Check

The following packet validations are done by IDS:

- consistency checks between frame length, IHL, Total length (no strict)
 - IHL >= 5
 - payload length + 8*frag_offset <= 64K
 - If DF==1 then must have frag_offset == 0
 - Invalid packet padding
 - disallow SA == 255.255.255.255
 - disallow SA == 127.x.x.x
 - disallow DA == 127.x.x.x

```
- disallow IPSA = IPDA
- disallow DA = 0.0.0.0
- disallow SA of class D
- disallow SA of class E
- disallow DA of class E
```




APPENDIX **A**

Performing Tasks Using the NX-OS Style CLI

- [Mixed-Mode Encapsulation, on page 71](#)
- [Port Channel and Virtual Port Channel Configuration, on page 73](#)
- [Enhanced LACP Policy Support, on page 75](#)
- [SPAN Features, on page 76](#)
- [BPDU Features, on page 77](#)
- [IGMP Querier and Snooping, on page 78](#)
- [Intra-EPG Isolation Enforcement for Cisco ACI Virtual Edge, on page 79](#)
- [Distributed Firewall, on page 80](#)

Mixed-Mode Encapsulation

Check or Change the VMM Domain Encapsulation Mode Using the NX-OS CLI

You can use the NX-OS CLI to check or change the encapsulation mode of a VMM domain.

**Note**

If EPGs are associated to the VMM domain, you cannot change its switching mode. If you want the domain to use a different switching mode, delete and re-create it. However, you can change the switching mode of the VMM domain if no EPGs are associated to it.

Step 1 Check the VMM domain encapsulation mode.

Example:

```
apic1(config-vmware-ave)# show run
# Command: show running-config vmware-domain mininet1 configure-ave
# Time: Tue Nov 21 07:07:58 2017
vmware-domain mininet1
  configure-ave
    switching mode vlan
    multicast-address 230.1.2.3
  exit
exit
apic1(config-vmware-ave)#
```

Step 2 Change the VMM domain encapsulation mode.

Example:

```
apic1# configure
apic1(config)# vmware-domain mininet
apic1(config-vmware)# configure-ave
apic1(config-vmware-ave)# switching mode ?
vlan          VLAN/SW Mode
vxlan         VXLAN/SW Mode
vxlan-ns      VXLAN/HW Mode
```

Override the VMM Domain Encapsulation Mode for an EPG Using the NX-OS Style CLI

After you create an EPG and associate it with a VMM domain, you can change the encapsulation mode of the EPG so it differs from or is the same of the VMM domain encapsulation mode.

Before you begin

You must already have created an EPG and have associated it with a VMM domain.

Specify the encapsulation mode for an EPG:

Example:

```
apic1(config)# tenant <tenant name>
apic1(config-tenant)# application <application name>
apic1(config-tenant-app)# epg <epg name>conf
apic1(config-tenant-app-epg)# vmware-domain member <vmm domain name>
apic1(config-tenant-app-epg-domain)# encap-mode auto | vlan | vxlan
apic1(config-tenant-app-epg-domain)# switching-mode AVE
```

You can choose one of the following encapsulation modes:

- **Auto**—This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
- **VLAN**—This overrides the domain's VXLAN configuration, and the EPG will use VLAN encapsulation. However, a fault will be triggered for the EPG if a VLAN pool is not configured on the domain.
- **VXLAN**—This overrides the domain's VLAN configuration, and the EPG will use VXLAN encapsulation. However, a fault will be triggered for the EPG if a multicast pool is not configured on the domain.

Port Channel and Virtual Port Channel Configuration

Configure Port Channel Mode Using the NX-OS Style CLI

Configure port channel mode.

Example:

```
apic1# conf t
apic1(config)# vmware-domain mininet
apic1(config-vmware)# configure-ave
apic1(config-vmware-ave)# channel-mode ?
  active      Set channeling mode to ACTIVE
  mac-pinning Set channeling mode to MAC-PINNING
  on          Set channeling mode to ON (static)
  passive     Set channeling mode to PASSIVE
apic1(config-vmware-ave)# channel-mode <mode>
```

Configure a Port Channel Using the NX-OS Style CLI

Create a port channel.

Example:

```
apic1# config
apic1(config)# template port-channel cli-pc1
apic1(config-if)# channel-mode active
apic1(config-if)# vlan-domain member cli-vdom1

apic1(config-if)# show running-config
# Command: show running-config interface port-channel cli-pc1
# Time: Thu Oct 1 10:38:30 2015
interface port-channel cli-pc1
  vlan-domain member cli-vdom1
  channel-mode active
exit
```

Configure a VPC Using the NX-OS Style CLI

Configuring a Virtual Port Channel (VPC) using the NX-OS style CLI consists of two tasks. Your first configure a VPC domain and then configure the VPC on the switch interfaces.

Configure a VPC Domain Using the NX-OS Style CLI

Configure a VPC domain.

Example:

```

apic1# config
apic1(config)# vpc domain explicit 10 leaf 101 102

apic1(config-vpc)# show running-config
# Command: show running-config vpc domain explicit 10 leaf 101 102
# Time: Thu Oct 1 10:39:26 2015
vpc domain explicit 10 leaf 101 102
exit

```

Configure a VPC on Switch Interfaces Using NX-OS Style CLI

Configure a VPC on switch interfaces

Example:

```

apic1# config
apic1(config)# leaf 101 - 102
apic1(config-leaf)# interface ethernet 1/3
apic1(config-leaf-if)# channel-group cli-pc1 vpc

apic1(config-leaf-if)# show running-config
# Command: show running-config leaf 101 - 102 interface ethernet 1/3
# Time: Thu Oct 1 10:41:15 2015
leaf 101
  interface ethernet 1/3
    channel-group cli-pc1 vpc
  exit
exit
leaf 102
  interface ethernet 1/3
    channel-group cli-pc1 vpc
  exit
exit

```

Configure Endpoint Retention Using the NX-OS Style CLI

Before you begin

You must have created a vCenter domain.

Step 1 In the CLI, enter configuration mode:

Example:

```

apic1# configure
apic1(config)#

```

Step 2 Configure a retention time for detached endpoints:

You can choose a delay of between 0 and 600 seconds. The default is 0.

Example:

```
apicl(config)# vmware-domain <domainName>

apicl(config-vmware)# ep-retention-time <value>
```

Enhanced LACP Policy Support

Create LAGs for DVS Uplink Port Groups Using the NX-OS Style CLI

Improve distributed virtual switch (DVS) uplink port group load balancing by putting the port groups into link aggregation groups (LAGs) and associating them with specific load-balancing algorithms. You can perform this task using the NX-OS style CLI.

Before you begin

You must have created a VMware vCenter virtual machine manager (VMM) domain for VMware VDS or Cisco Application Centric Infrastructure (ACI) Virtual Edge.

Create or delete an enhanced LACP policy.

Example:

```
apicl(config-vmware)# enhancedlacp LAG name
apicl(config-vmware-enhancedlacp)# lbmode loadbalancing mode
apicl(config-vmware-enhancedlacp)# mode mode
apicl(config-vmware-enhancedlacp)# numlinks max number of uplinks
apicl(config-vmware)# no enhancedlacp LAG name to delete
```

What to do next

If you are using VMware VDS, associate endpoint groups (EPGs) to the domain with the enhanced LACP policy. If you are using Cisco Application Centric Infrastructure (ACI) Virtual Edge, associate internally created inside and outside port groups with the enhanced LACP policy, then associate EPGs to the domain with the policy.

Associate Internal Port Groups to VMware vCenter Domains with Enhanced LACP Policies Using the NX-OS Style CLI

Associate Cisco Application Centric Infrastructure (ACI) Virtual Edge internally created inside and outside port groups with a VMware vCenter domain with an enhanced LACP policy. You can perform this task using the NX-OS style CLI.

Before you begin

You must have created link aggregation groups (LAGs) for the distributed virtual switch (DVS) uplink port groups and associated a load-balancing algorithm to the LAGs.

Associate (or deassociate) internal endpoint groups (EPGs) to the VMM domain with the enhanced LACP policy.

Example:

```
apic1(config-vmware)# lag-policy name of the policy to associate
apic1(config-vmware)# no lag-policy name of the policy to deassociate
```

What to do next

Associate EPGs with the VMware vCenter domain containing the enhanced LACP policy.

Associate Application EPGs to VMware vCenter Domains with Enhanced LACP Policies Using the NX-OS Style CLI

Associate application endpoint groups (EPGs) with the VMware vCenter domain with LAGs and a load-balancing algorithm. You can perform this task using NX-OS style CLI. You can also deassociate application EPGs from the domain.

Before you begin

You must have created link aggregation groups (LAGs) for distributed virtual switch (DVS) uplink port groups and associated a load-balancing algorithm to the LAGs.

Step 1 Associate an application EPG with the domain or deassociate it from the domain.

Example:

```
apic1(config-tenant-app-epg-domain)# lag-policy name of the LAG policy to associate
apic1(config-tenant-app-epg-domain)# no lag-policy name of the LAG policy to deassociate
```

Step 2 Repeat Step 1 for other application EPGs in the tenant as desired.

SPAN Features

Configure SPAN Using the NX-OS CLI

Step 1 Configure SPAN.

Example:

```

apicl(config)# monitor virtual session cli-vspan1
apicl(config-monitor-virtual)# source tenant cli-esx1 application cli-esx1 epg cli-vspan1 mac
<00:50:56:BA:BE:0F>
apicl(config-monitor-virtual-source)# direction both
apicl(config-monitor-virtual-source)# exit
apicl(config-monitor-virtual)# destination tenant cli-esx1 application cli-vspan1 epg cli-esx1b mac
<00:50:56:BA:F0:E0>

apicl(config)# vmware-domain cli-esx
apicl(config-vmware)# configure-ave
apicl(config-vmware-ave)# monitor virtual session cli-vspan1

```

Step 2 Verify the configuration.

Example:

```

apicl(config-monitor-virtual)# show running-config
# Command: show running-config monitor virtual session cli-vspan1
# Time: Thu Oct 8 11:20:09 2015
monitor virtual session cli-vspan1
  source tenant cli-esx1 application cli-esx1 epg cli-esx1 mac 00:50:56:BA:BE:0F
  exit
  destination tenant cli-esx1 application cli-esx1 epg cli-esx1b mac 00:50:56:BA:F0:E0
  exit

```

BPDU Features

Configure BPDU Features Using the NX-OS Style CLI

Step 1 Enter the vmware-domain mode.

Example:

```

apicl# configure
apicl(config)# vmware-domain domain name
AVE-Vlan AVE2-VXLAN Test Test2

```

Step 2 Create a spanning-tree interface policy.

Example:

```

apicl(config-vmware)# configure-ave
apicl(config-vmware-ave)# spanning-tree
  bpdu-filter bpdu-guard
apicl(config-vmware-ave)# spanning-tree
  bpdu-filter Configure BPDU filter override on AVE uplink ports
  bpdu-guard Configure BPDU guard override on AVE uplink ports

```

Step 3 Disable or enable BPDU filter.

Example:

```

apicl(config-vmware-ave)# spanning-tree bpdu-filter
  default disable enable
apicl(config-vmware-ave)# spanning-tree bpdu-filter
  default Remove BPDU filter/guard override policy
  disable Disable BPDU filter

```

```
enable    Enable BPDU filter
```

Step 4 Disable or enable BPDU guard.

```
apic1(config-vmware-ave)# spanning-tree bpdu-guard
                        default  disable  enable
```

IGMP Querier and Snooping

Configure IGMP Querier Using the NX-OS Style CLI

Configure IGMP querier.

Example:

```
apic1# configure
apic1(config)# tenant t1
apic1(config-tenant)# interface bridge-domain bd1
apic1(config-tenant-interface)# ip address <192.168.1.1/24> snooping-querier
<CR>
multi-site  Set the address as multi-site address
scope       Scope of the address among ['public', 'private']
secondary   Set the address as secondary address
```

Configure IGMP Snooping to Take Effect Immediately Using the NX-OS Style CLI

Configure IGMP snooping to take effect immediately.

Example:

```
apic1# configure
apic1(config)# tenant t1
apic1(config-tenant)# interface bridge-domain bd1
apic1(config-tenant-interface)# ip igmp snooping querier
```

Configure IGMP Snooping to Take Effect Later Using the NX-OS Style CLI

Configure IGMP snooping to take effect later.

Example:


```
apicl# configure
apicl(config)# tenant t1
apicl(config-tenant)# template ip igmp snooping policy <foo_igmp>
apicl(config-tenant-template-ip-igmp-snooping)# ip igmp snooping querier
```

Intra-EPG Isolation Enforcement for Cisco ACI Virtual Edge

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).

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 41](#) and [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 42](#) in this guide.

Distributed Firewall

Configure a Stateful Policy for Distributed Firewall Using the NX-OS Style CLI

Configure a stateful policy in the Cisco APIC.

Example:

```
apic1(config)# tenant Tenant1
apic1(config-tenant)# access-list TCP-511 apic1
apic1(config-tenant-acl)# match icmp
apic1(config-tenant-acl)# match raw TCP-511 dFromPort 443 dToPort 443 etherT ip prot 6 stateful yes

apic1(config-tenant-acl)# match raw tcp etherT ip prot 6 sFromPort 443 sToPort 443 stateful yes
apic1(config-tenant-acl)# match raw tcp-22out dFromPort 22 dToPort 22 etherT ip prot 6 stateful yes
apic1(config-tenant-acl)# match raw tcp-all etherT ip prot 6 stateful yes
apic1(config-tenant-acl)# match raw tcp22-from etherT ip prot 6 sFromPort 22 sToPort 22 stateful yes
apic1(config-tenant-acl)# exit apic1(config-tenant)# contract TCP511
apic1(config-tenant-contract)# subject TCP-ICMP
apic1(config-tenant-contract-subj)# access-group TCP-511 both
apic1(config-tenant-contract-subj)# access-group arp both
apic1(config-tenant-contract-subj)#
```

What to do next

Create a Distributed Firewall policy.

Enable Distributed Firewall or Change Its Mode Using the NX-OS Style CLI

You can use the NX-OS style CLI to enable Distributed Firewall or change its mode.

Enable Distributed Firewall or change its mode.

Example:

```
apic1# configure
apic1(config)# vmware-domain Direct-AVE2-VXLAN
apic1(config-vmware)# configure-ave
apic1(config-vmware-ave)# firewall mode < any of below 3>
disabled Disabled mode
enabled Enabled mode
learning Learning mode
```

Configure Parameters for Distributed Firewall Flow Information Using the NX-OS Style CLI

Step 1 Configure the parameters for the syslog server or servers.

Example:

```
apicl# configure
apicl(config)# logging server-group group name
apicl(config-logging)# server IP address severity severity level facility facility name port
1-65535 mgmtEpg MgmtEpg
```

You can repeat the last command for additional syslog servers; you can configure up to three syslog servers.

Step 2 Configure the parameters for the syslog source.

Example:

```
apicl# configure
apicl(config)# vmware-domain Direct-AVE
apicl(config-vmware)# configure-ave
apicl(config-vmware-ave)# firewall mode enabled
apicl(config-vmware-ave)# firewall-logging server-group group name action-type permit,
deny severity severity polling-interval 60-86400
```

Note You must enter the **firewall mode enabled** command before you enter the **firewall-logging** command.

Note For the **firewall-logging** command, you can enter either **permit** or **deny**. You can also enter both, separated by a comma.



APPENDIX **B**

Performing Tasks Using REST API

- [Mixed-Mode Encapsulation, on page 83](#)
- [Port Channel and Virtual Port Channel Configuration, on page 84](#)
- [Enhanced LACP Policy Support, on page 88](#)
- [SPAN Features, on page 90](#)
- [BPDU Features, on page 93](#)
- [IGMP Querier and Snooping, on page 94](#)
- [Intra-EPG Isolation Enforcement, on page 95](#)
- [Distributed Firewall, on page 96](#)

Mixed-Mode Encapsulation

Check or Change the VMM Domain Encapsulation Mode Using the REST API

You can use the REST API to discover and change the encapsulation mode of a VMM domain.

**Note**

If EPGs are associated to the VMM domain, you cannot change its switching mode. If you want the domain to use a different switching mode, delete and re-create it. However, you can change the switching mode of the VMM domain if no EPGs are associated to it.

Discover and change the encapsulation mode of a VMM domain.

Example:

```
<vmmProvP vendor="VMware">
  <vmmDomP name="mininet" enableAVE="true" enfPref="sw" mcastAddr="225.1.1.1" encapMode="vxlan"
  prefEncapMode="vxlan">
</vmmProvP>
```

Override the VMM Domain Encapsulation Mode for an EPG Using the REST API

Override the VMM domain encapsulation mode for an EPG.

Example:

```
<polUni>
<fvTenant name="coke">
<fvAp name="sap">
<fvAEPg name="web1">
<fvRsDomAtt resImedcy="immediate"
tDn="uni/vmmp-VMware/dom-mininet"
switchingMode="AVE" encapMode="vxlan"/>
</fvAEPg>
</fvAp>
</fvTenant>
</polUni>
```

For **encapMode=**, you can enter one of the following:

- **auto**—This causes the EPG to use the same encapsulation mode as the VMM domain. This is the default configuration.
- **vlan**—This overrides the domain's VXLAN configuration, and the EPG will use VLAN encapsulation. However, a fault will be triggered for the EPG if a VLAN pool is not configured on the domain.
- **vxlan**—This overrides the domain's VLAN configuration, and the EPG will use VXLAN encapsulation. However, a fault will be triggered for the EPG if a multicast pool is not configured on the domain.

Port Channel and Virtual Port Channel Configuration

Configure an LACP Port Channel Policy Using the REST API

Step 1 Create a node profile that specifies the leaf IDs that the access port profiles are associated with.

Example:

```
<infraInfra dn="uni/infra">
  <infraNodeP name="bLeaf">
    <infraLeafS name="leafs" type="range">
      <infraNodeBlk name="nblk" from="_17" to="_17">
        </infraNodeBlk>
      </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-shipping1"/>
      <infraRsAccPortP tDn="uni/infra/accportprof-shipping2"/>
    </infraNodeP>
```

Step 2 Create an access port profile that specifies the port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping1">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="19" toPort="20"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag1" />
  </infraHPortS>
</infraAccPortP>
```

Step 3 Create an access port profile that specifies a second port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping2">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="21" toPort="22"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag2" />
  </infraHPortS>
</infraAccPortP>
```

Step 4 Create an access bundle group that points to the port channel interface policy.

Example:

```
<infraFuncP>
  <infraAccBndlGrp name="accountingLag1" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp1' />
    <infraRsAttEntP tDn="uni/infra/attentp-default" />
  </infraAccBndlGrp>
  <infraAccBndlGrp name="accountingLag2" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp2' />
    <infraRsAttEntP tDn="uni/infra/attentp-default" />
  </infraAccBndlGrp>
</infraFuncP>
```

Step 5 Create a port channel interface policy.

Example:

```
</infraFuncP>
<lacpLagPol name='accountingLacp1' ctrl='15' descr='accounting' maxLinks='14' minLinks='1'
mode='active' />
<lacpLagPol name='accountingLacp2' ctrl='15' descr='accounting' maxLinks='14' minLinks='1'
mode='active' />
```

You can set the mode to 'passive' instead of 'active'.

Step 6 Associate the VMM domain to the attachable entity profile.

Example:

```
<infraAttEntityP name="default"> <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
</infraAttEntityP>

</infraInfra>
```

Configure a MAC Pinning Port Channel Policy Using the REST API

Step 1 Create a node profile that specifies the leaf IDs that the access port profiles are associated with.

Example:

```
<infraInfra dn="uni/infra">
  <infraNodeP name="bLeaf">
    <infraLeafS name="leafs" type="range">
      <infraNodeBlk name="nblk" from_="17" to_="17">
        </infraNodeBlk>
      </infraLeafS>
    <infraRsAccPortP tDn="uni/infra/accportprof-shipping1"/>
    <infraRsAccPortP tDn="uni/infra/accportprof-shipping2"/>
  </infraNodeP>
```

Step 2 Create an access port profile that specifies the port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping1">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="19" toPort="20"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag1" />
  </infraHPortS>
</infraAccPortP>
```

Step 3 Create an access port profile that specifies a second port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping2">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="21" toPort="22"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag2" />
  </infraHPortS>
</infraAccPortP>
```

Step 4 Create an access bundle group that points to the port channel interface policy.

Example:

```
<infraFuncP>
  <infraAccBndlGrp name="accountingLag1" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp1'/>
    <infraRsAttEntP tDn="uni/infra/attentp-default"/>
  </infraAccBndlGrp>
  <infraAccBndlGrp name="accountingLag2" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp2'/>
    <infraRsAttEntP tDn="uni/infra/attentp-default"/>
  </infraAccBndlGrp>
</infraFuncP>
```

Step 5 Create a port channel interface policy.

Example:

```
<lacpLagPol name='accountingLacp1' ctrl='15' descr='accounting' maxLinks='14' minLinks='1'
mode='mac-pin' />
<lacpLagPol name='accountingLacp2' ctrl='15' descr='accounting' maxLinks='14' minLinks='1'
mode='mac-pin' />
```

Step 6 Associate the VMM domain to the attachable entity profile.

Example:

```
<infraAttEntityP name="default"> <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
</infraAttEntityP>

</infraInfra>
```


Configure a Static Port Channel Policy Using the REST API

Step 1 Create a node profile that specifies the leaf IDs that the access port profiles are associated with.

Example:

```
<infraInfra dn="uni/infra">
  <infraNodeP name="bLeaf">
    <infraLeafS name="leafs" type="range">
      <infraNodeBlk name="nblk" from_"17" to_"17">
        </infraNodeBlk>
      </infraLeafS>
    <infraRsAccPortP tDn="uni/infra/accportprof-shipping1"/>
    <infraRsAccPortP tDn="uni/infra/accportprof-shipping2"/>
  </infraNodeP>
```

Step 2 Create an access port profile that specifies the port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping1">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="19" toPort="20"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag1" />
  </infraHPortS>
</infraAccPortP>
```

Step 3 Create an access port profile that specifies a second port included in the access bundle group.

Example:

```
<infraAccPortP name="shipping2">
  <infraHPortS name="pselc" type="range">
    <infraPortBlk name="blk" fromCard="1" toCard="1" fromPort="21" toPort="22"/>
    <infraRsAccBaseGrp tDn="uni/infra/funcprof/accbundle-accountingLag2" />
  </infraHPortS>
</infraAccPortP>
```

Step 4 Create an access bundle group that points to the port channel interface policy.

Example:

```
<infraFuncP>
  <infraAccBndlGrp name="accountingLag1" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp1'/>
    <infraRsAttEntP tDn="uni/infra/attentp-default"/>
  </infraAccBndlGrp>
  <infraAccBndlGrp name="accountingLag2" lagT='link'>
    <infraRsLacpPol tnLacpLagPolName='accountingLacp2'/>
    <infraRsAttEntP tDn="uni/infra/attentp-default"/>
  </infraAccBndlGrp>
</infraFuncP>
```

Step 5 Create a port channel interface policy.

Example:

```
<lacpLagPol name='accountingLacp1' ctrl='15' descr='accounting' maxLinks='14' minLinks='1' mode='off'
/>
<lacpLagPol name='accountingLacp2' ctrl='15' descr='accounting' maxLinks='14' minLinks='1' mode='off'
/>
```

Step 6 Associate the VMM domain to the attachable entity profile.

Example:

```
<infraAttEntityP name="default"> <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
</infraAttEntityP>

</infraInfra>
```

Enhanced LACP Policy Support

Create LAGs for DVS Uplink Port Groups Using REST API

Improve distributed virtual switch (DVS) uplink port group load balancing by putting the port groups into link aggregation groups (LAGs) and associating them with specific load-balancing algorithms. You can perform this task using REST API.

Before you begin

You must have created a VMware vCenter virtual machine manager (VMM) domain for VMware VDS or Cisco Application Centric Infrastructure (ACI) Virtual Edge.

Step 1 Create the the LAG and associate it with a load-balancing algorithm.

Example:

```
<polUni>
<vmwProvP vendor="VMware">
  <vmwDomP name="mininetlacpavs">
    <vmwVSwitchPolicyCont>
      <lacpEnhancedLagPol name="lag2" mode="passive" lbmode="vlan" numLinks="4">
      </lacpEnhancedLagPol>
    </vmwVSwitchPolicyCont>
  </vmwDomP>
</vmwProvP>
</polUni>
```

Step 2 Repeat the step to create other LAGs for the DVS.

What to do next

If you are using VMware VDS, associate endpoint groups (EPGs) to the domain with the enhanced LACP policy. If you are using Cisco Application Centric Infrastructure (ACI) Virtual Edge, associate internally created inside and outside port groups with the enhanced LACP policy, then associate EPGs to the domain with the policy.

Associate Internal Port Groups to VMware vCenter Domains with Enhanced LACP Policies Using REST API

Associate Cisco Application Centric Infrastructure (ACI) Virtual Edge internally created inside and outside port groups with a VMware vCenter domain with an enhanced LACP policy. You can perform this task using REST API.

Before you begin

You must have created link aggregation groups (LAGs) for distributed virtual switch (DVS) uplink portgroups and associated a load-balancing algorithm to the LAGs.

Example:

```
<vmmProvP vendor="VMware">
  <vmmDomP name="mininetlacpavs" enfPref="sw" mcastAddr="225.1.1.1" prefEncapMode="vlan"
enableAVE="true">
    <vmmRsPrefEnhancedLagPol
tDn="uni/vmmp-VMware/dom-mininetlacpavs/vswitchpolcont/enlacplagp-lag2"/>
  </vmmDomP>
</vmmProvP>
```

What to do next

Associate endpoint groups (EPGs) with the VMware vCenter domain containing the enhanced LACP policy.

Associate Application EPGs to VMware vCenter Domains with Enhanced LACP Policies Using REST API

Associate application endpoint groups (EPGs) with the VMware vCenter domain with LAGs and a load-balancing algorithm. You can perform this task using REST API. You can also deassociate application EPGs from the domain.

Before you begin

You must have created link aggregation groups (LAGs) for distributed virtual switch (DVS) uplink port groups and associated a load-balancing algorithm to the LAGs.

Step 1 Associate an EPG to a VMware vCenter domain with LAGs associated to a load-balancing algorithm.

Example:

```
<polUni>
  <fvTenant
    dn="uni/tn-coke"
    name="coke">
    <fvCtx name="cokectx"/>
    <fvAp
      dn="uni/tn-coke/ap-sap"
      name="sap">
      <fvAEPg
```

```

    dn="uni/tn-coke/ap-sap/epg-web3"
    name="web3" >
      <fvRsBd tnFvBDName="cokeBD2" />
      <fvRsDomAtt resImedcy="immediate" switchingMode="native"
        tDn="uni/vmmp-VMware/dom-mininetlacpavs">
        <fvAEPgLagPolAtt >
          <fvRsVmmVSwitchEnhancedLagPol
tDn="uni/vmmp-VMware/dom-mininetlacpavs/vswitchpolcont/enlacplagp-lag2"/>
        </fvAEPgLagPolAtt>
      </fvRsDomAtt>
    </fvAEPg>
  </fvAp>
</fvTenant>
</polUni>

```

Step 2 Repeat Step 1 for other application EPGs in the tenant, as desired.

SPAN Features

Configure Local SPAN with a CEP Source Using the REST API

Configure local SPAN with a CEP source.

Example:

```

<polUni>
  <infraInfra>
    <spanVSrcGrp name="srcgrp2">
      <spanVSrc name="src1" dir="both" >
        <spanRsSrcToVPort tDn="uni/tn-t0/ap-a0/epg-g3/cep-00:50:56:B3:24:E1"/>

      </spanVSrc>
      <spanSpanLbl name="destgrp1">
        </spanSpanLbl>
    </spanVSrcGrp>
    <infraFuncP>
      <infraAccBndlGrp name="test-lvspan">
        <infraRsSpanVSrcGrp tnSpanVSrcGrpName="srcgrp1"/>
        <infraRsSpanVDestGrp tnSpanVDestGrpName="destgrp1"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test-lvspan"/>
      </infraAccBndlGrp>
    </infraFuncP>
    <spanVDestGrp
      name="destgrp2">
      <spanVDest name="dest1">
        <spanRsDestToVPort
tDn="uni/tn-t0/ap-a0/Promiscuous-EPG/cep-00:50:56:B3:5F:AA"/>
      </spanVDest>
    </spanVDestGrp>
    <infraAttEntityP name="test-lvspan">
      <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
    </infraAttEntityP>
  </infraInfra>
</polUni>

```

Configure Local SPAN with an EPG Source Using the REST API

Configure local SPAN with an EPG source.

Example:

```
<polUni>
  <infraInfra>
    <spanVSrcGrp
      name="srcgrp2" adminSt="start">
        <spanVSrc name="src2" dir="both">
          <spanRsSrcToEpg tDn="uni/tn-t0/ap-a0/epg-g11"/>
        </spanVSrc>
        <spanSpanLbl name="destgrp1">
          </spanSpanLbl>
        </spanVSrcGrp>
      <infraFuncP>
        <infraAccBndlGrp name="test-lvspan">
          <infraRsSpanVSrcGrp tnSpanVSrcGrpName="srcgrp2"/>
          <infraRsSpanVDestGrp tnSpanVDestGrpName="destgrp1"/>
        </infraAccBndlGrp>
      </infraFuncP>
    <spanVDestGrp
      name="destgrp2">
        <spanVDest name="dest1">
          <spanRsDestToVPort
            tDn="uni/tn-t0/ap-a0/Promiscuous-EPG/cep-00:50:56:B3:5F:AA"/>
          </spanVDest>
        </spanVDestGrp>
        <infraAttEntityP name="test-lvspan">
          <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
        </infraAttEntityP>
      </infraInfra>
    </polUni>
```

Configure ERSPAN with a CEP Source Using the REST API

Configure ERSPAN with a CEP source.

Example:

```
<polUni>
  <infraInfra>
    <spanVSrcGrp name="srcgrp2">
      <spanVSrc name="src1" dir="both" >
        <spanRsSrcToVPort tDn="uni/tn-t0/ap-a0/epg-g3/cep-00:50:56:B3:24:E1"/>
      </spanVSrc>
      <spanSpanLbl name="destgrp1">
        </spanSpanLbl>
      </spanVSrcGrp>
    <infraFuncP>
      <infraAccBndlGrp name="test-lvspan">
        <infraRsSpanVSrcGrp tnSpanVSrcGrpName="srcgrp1"/>
        <infraRsSpanVDestGrp tnSpanVDestGrpName="destgrp1"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test-lvspan"/>
      </infraAccBndlGrp>
```

```

    </infraFuncP>
    <spanVDestGrp
      name="destgrp1">
    <spanVDest name="dest1">
      <spanVEpgSummary name="summ1" dstIp="10.30.13.195" ttl="50" mtu="1500" dscp="2"/>
    </spanVDest>
    </spanVDestGrp>
    <infraAttEntityP name="test-lvspan">
      <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
    </infraAttEntityP>
  </infraInfra>
</polUni>

```

Configure ERSPAN with a Static Endpoint Using the REST API

Configure ERSPAN with a static CEP source.

Example:

```

<polUni>
  <fvTenant name="infra">
    <fvAp name="access">
      <fvAEPg name="default">
        <fvStCEp name="erspan-dest "
          type="tep"
          mac="00:50:56:B3:42:9C"
          ip="10.0.0.50"
          encap="vlan-4093">
          <fvRsStCEpToPathEp tDn="topology/pod-1/paths-110/pathep-[macpin-1]"/>
        </fvStCEp>
      </fvAEPg>
    </fvAp>
  </fvTenant>
</polUni>

```

Configure ERSPAN with an EPG Source Using the REST API

Configure ERSPAN with an EPG source.

Example:

```

<polUni>
  <infraInfra>
    <spanVSrcGrp
      name="srcgrp2" adminSt="start">
    <spanVSrc name="src2" dir="both">
      <spanRsSrcToEpg tDn="uni/tn-t0/ap-a0/epg-g11"/>
    </spanVSrc>
    <spanSpanLbl name="destgrp1">
    </spanSpanLbl>
  </spanVSrcGrp>
  <infraFuncP>
    <infraAccBndlGrp name="test-lvspan">
      <infraRsSpanVSrcGrp tnSpanVSrcGrpName="srcgrp2"/>
    </infraFuncP>
  </infraInfra>
</polUni>

```

```

        <infraRsSpanVDestGrp tnSpanVDestGrpName="destgrp1"/>
      </infraAccBndlGrp>
    </infraFuncP>
    <spanVDestGrp
      name="destgrp1">
      <spanVDest name="dest1">
        <spanVEpgSummary name="summ1" dstIp="10.30.13.195" ttl="50" mtu="1500" dscp="2"/>
      </spanVDest>
    </spanVDestGrp>
    <infraAttEntityP name="test-lvspan">
      <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
    </infraAttEntityP>
  </infraInfra>
</polUni>

```

BPDU Features

Configure BPDU Features Using the REST API

Step 1 Configure BPDU Guard.

Example:

```

<polUni>
  <infraInfra>
    <stpIfPol name="testStp5" ctrl="bpdu-guard"/>
    <infraFuncP>
      <infraAccBndlGrp name="test51">
        <infraRsStpIfPol tnStpIfPolName="testStp5"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test-bpdu"/>
      </infraAccBndlGrp>
    </infraFuncP>
  </infraInfra>
</polUni>

<vmmpProvP vendor="VMware">
  <vmmpDomP name="mininet">
    <vmmpVSwitchPolicyCont>
      <vmmpRsVswitchOverrideStpPol tDn="uni/infra/ifPol-testStp5"/>
    </vmmpVSwitchPolicyCont>
  </vmmpDomP>
</vmmpProvP>

```

Step 2 Configure BPDU filtering.

Example:

```

<polUni>
  <infraInfra>
    <stpIfPol name="testStp5" ctrl="bpdu-filter"/>
    <infraFuncP>
      <infraAccBndlGrp name="test51">
        <infraRsStpIfPol tnStpIfPolName="testStp5"/>
        <infraRsAttEntP tDn="uni/infra/attentp-test-bpdu"/>
      </infraAccBndlGrp>
    </infraFuncP>
  </infraInfra>
</polUni>

```

```

        </infraFuncP>
    </infraInfra>
</polUni>

<vmmProvP vendor="VMware">
    <vmmDomP name="mininet">
        <vmmVSwitchPolicyCont>
            <vmmRsVswitchOverrideStpPol tDn="uni/infra/ifPol-testStp5"/>
        </vmmVSwitchPolicyCont>
    </vmmDomP>
</vmmProvP>

```

IGMP Querier and Snooping

Enable IGMP Querier on the Bridge Domain Subnet Using the REST API

Enable IGMP querier on the bridge domain subnet.

Example:

```

<fvTenant name="ms10">
  <fvCtx name="msv10"/>

  <fvBD name="msb10">
    <fvSubnet ctrl="querier" descr="" ip="1.1.9.1/24" name="" nameAlias=""
preferred="no" scope="private" virtual="no"/>
    <fvRsCtx tnFvCtxName="msv10"/>
  </fvBD>
</fvTenant>

```

Configure an IGMP Snooping Policy Using the REST API

Create an IGMP snooping policy and apply it to the bridge domain.

Example:

```

<igmpSnoopPol name="igmp_snp_bd_21"
  adminSt="enabled"
  ctrl="fast-leave,querier"
  lastMbrIntvl="1"
  queryIntvl="125"
  rspIntvl="10"
  startQueryCnt="2"
  startQueryIntvl="31"
/>
<fvCtx name="msv10"/>

<fvBD name="msb10">
  <fvRsCtx tnFvCtxName="msv10"/>

```



```
<!-- Bind IGMP snooping to a BD -->
<fvRsIgmpsn tnIgmpSnoopPolName="igmp_snp_bd_21"/>
</fvBD></fvTenant>
```

Intra-EPG Isolation Enforcement

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

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).

Step 1 Send this HTTP POST message to deploy the application using the XML API.

Example:

```
POST
https://10.197.139.36/api/mo/uni/tn-Tenant2.xml
```

Step 2 For a VMM deployment, include the XML structure in the following example in the body of the POST message.

Example:

```
<fvTenant name="Tenant2" >
  <fvAp name="AP-1">
    <fvAEPg name="EPG-61" pcEnfPref="enforced">
      <!-- pcEnfPref="enforced" ENABLES ISOLATION-->
      <!-- pcEnfPref="unenforced" DISABLES ISOLATION-->
      <fvRsBd tnFvBDName="BD-61" />
      <fvRsDomAtt switchingMode="AVE" encapMode="vxlan" resImedcy="immediate"
tDn="uni/vmmp-VMware/dom-D-AVE-SITE-1-XXIII" >
        </fvRsDomAtt>
      </fvAEPg>
    </fvAp>
  </fvTenant>
```

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 41](#) and [View Statistics for Isolated Endpoints on Cisco ACI Virtual Edge Under the Tenants Tab, on page 42](#) in this guide.

Distributed Firewall

Configure a Stateful Policy for Distributed Firewall Using the REST API

Configure a stateful policy in the Cisco APIC.

Step 1 Log in to the Cisco APIC.

Step 2 Post the policy to `https://APIC-ip-address/api/node/mo/.xml`.

Example:

```
<polUni>
  <infraInfra>

    <nwsFwPol name="fwpol1" mode="enabled"/>      (enabled, disabled, learning)

    <infraFuncP>
      <infraAccBndlGrp name="fw-bundle">
        <infraRsFwPol tnNwsFwPolName="fwpol1"/>
        <infraRsAttEntP tDn="uni/infra/attentp-testfw2"/>
      </infraAccBndlGrp>
    </infraFuncP>

    <infraAttEntityP name="testfw2">
      <infraRsDomP tDn="uni/vmmp-VMware/dom-mininet"/>
    </infraAttEntityP>

  </infraInfra>
</polUni>
```

What to do next

Create a Distributed Firewall policy.

Configure Parameters for Distributed Firewall Flow Information Using the REST API

Step 1 Send an HTTP POST message to deploy the application using the XML API.

Example:

POST `https://10.197.139.36/api/node/mo/uni/fabric/slgroup-Syslog-Servers.xml`

Step 2 Configure the parameters for the syslog server or servers.

Example:

```
<syslogGroup descr="" dn="uni/fabric/slgroup-Syslog-Servers" format="aci" name="Syslog-Servers"
nameAlias="">
  <syslogRemoteDest adminState="enabled" descr="" format="aci" forwardingFacility="local7"
```

```
host="10.197.139.216" name="10.197.139.216" nameAlias="" port="1514" severity="debugging">
  <fileRsARemoteHostToEpg tDn="uni/tn-mgmt/mgmt-default/oob-default"/>
</syslogRemoteDest>
<syslogProf adminState="enabled" descr="" name="syslog" nameAlias=""/>
<syslogFile adminState="disabled" descr="" format="aci" name="" nameAlias="" severity="information"/>

<syslogConsole adminState="disabled" descr="" format="aci" name="" nameAlias="" severity="alerts"/>
</syslogGroup>
```
