



# Cisco ACI Virtual Machine Networking

---

This chapter contains the following sections:

- [Cisco ACI VM Networking Support for Virtual Machine Managers, on page 1](#)
- [Virtual Machine Manager Domain Main Components , on page 2](#)
- [Virtual Machine Manager Domains, on page 3](#)
- [VMM Domain VLAN Pool Association, on page 4](#)
- [VMM Domain EPG Association, on page 5](#)
- [About Trunk Port Group, on page 7](#)
- [Attachable Entity Profile, on page 8](#)
- [EPG Policy Resolution and Deployment Immediacy, on page 9](#)
- [Guidelines for Deleting VMM Domains, on page 11](#)
- [NetFlow with Virtual Machine Networking, on page 11](#)
- [Troubleshooting VMM Connectivity, on page 19](#)

## Cisco ACI VM Networking Support for Virtual Machine Managers

### Benefits of ACI VM Networking

Cisco ACI virtual machine (VM) networking supports hypervisors from multiple vendors. It provides the hypervisors programmable and automated access to high-performance scalable virtualized data center infrastructure.

Programmability and automation are critical features of scalable data center virtualization infrastructure. The Cisco ACI open REST API enables virtual machine integration with and orchestration of the policy model-based Cisco ACI fabric. Cisco ACI VM networking enables consistent enforcement of policies across both virtual and physical workloads managed by hypervisors from multiple vendors.

Attachable entity profiles easily enable VM mobility and placement of workloads anywhere in the Cisco ACI fabric. The Cisco Application Policy Infrastructure Controller (APIC) provides centralized troubleshooting, application health score, and virtualization monitoring. Cisco ACI multi-hypervisor VM automation reduces or eliminates manual configuration and manual errors. This enables virtualized data centers to support large numbers of VMs reliably and cost effectively.

## Supported Vendors

Cisco ACI supports virtual machine managers (VMMs) from the following products and vendors:

- Cisco Application Centric Infrastructure Virtual Edge  
For information, see the [Cisco ACI Virtual Edge documentation](#) on Cisco.com.
- Cisco Application Virtual Switch (AVS)  
For information, see the chapter "Cisco ACI with Cisco AVS" in the [Cisco ACI Virtualization Guide](#) and [Cisco AVS documentation](#) on Cisco.com.
- Cloud Foundry  
Cloud Foundry integration with Cisco ACI is supported beginning with Cisco APIC Release 3.1(2). For information, see the knowledge base article, [Cisco ACI and Cloud Foundry Integration](#) on Cisco.com.
- Kubernetes  
For information, see the knowledge base article, [Cisco ACI and Kubernetes Integration](#) on Cisco.com.
- Microsoft System Center Virtual Machine Manager (SCVMM)  
For information, see the chapters "Cisco ACI with Microsoft SCVMM" and "Cisco ACI with Microsoft Windows Azure Pack" in the [Cisco ACI Virtualization Guide](#).
- OpenShift  
For information, see the [OpenShift documentation](#) on Cisco.com.
- OpenStack  
For information, see the [OpenStack documentation](#) on Cisco.com.
- Red Hat Virtualization (RHV)  
For information, see the knowledge base article, [Cisco ACI and Red Hat Integration](#) on Cisco.com.
- VMware Virtual Distributed Switch (VDS)  
For information, see the chapter "Cisco "ACI with VMware VDS Integration" in the [Cisco ACI Virtualization Guide](#).

See the [Cisco ACI Virtualization Compatibility Matrix](#) for the most current list of verified interoperable products.

# Virtual Machine Manager Domain Main Components

ACI fabric virtual machine manager (VMM) domains enable an administrator to configure connectivity policies for virtual machine controllers. The essential components of an ACI VMM domain policy include the following:

- **Virtual Machine Manager Domain Profile**—Groups VM controllers with similar networking policy requirements. For example, VM controllers can share VLAN pools and application endpoint groups (EPGs). The APIC communicates with the controller to publish network configurations such as port groups that are then applied to the virtual workloads. The VMM domain profile includes the following essential components:

- **Credential**—Associates a valid VM controller user credential with an APIC VMM domain.
- **Controller**—Specifies how to connect to a VM controller that is part of a policy enforcement domain. For example, the controller specifies the connection to a VMware vCenter that is part a VMM domain.



---

**Note** A single VMM domain can contain multiple instances of VM controllers, but they must be from the same vendor (for example, from VMware or from Microsoft).


---

- **EPG Association**—Endpoint groups regulate connectivity and visibility among the endpoints within the scope of the VMM domain policy. VMM domain EPGs behave as follows:
  - The APIC pushes these EPGs as port groups into the VM controller.
  - An EPG can span multiple VMM domains, and a VMM domain can contain multiple EPGs.
- **Attachable Entity Profile Association**—Associates a VMM domain with the physical network infrastructure. An attachable entity profile (AEP) is a network interface template that enables deploying VM controller policies on a large set of leaf switch ports. An AEP specifies which switches and ports are available, and how they are configured.
- **VLAN Pool Association**—A VLAN pool specifies the VLAN IDs or ranges used for VLAN encapsulation that the VMM domain consumes.

## Virtual Machine Manager Domains

An APIC VMM domain profile is a policy that defines a VMM domain. The VMM domain policy is created in APIC and pushed into the leaf switches.

Figure 1: ACI VMM Domain VM Controller Integration



VMM domains provide the following:

- A common layer in the ACI fabric that enables scalable fault-tolerant support for multiple VM controller platforms.
- VMM support for multiple tenants within the ACI fabric.

VMM domains contain VM controllers such as VMware vCenter or Microsoft SCVMM Manager and the credential(s) required for the ACI API to interact with the VM controller. A VMM domain enables VM mobility within the domain but not across domains. A single VMM domain can contain multiple instances of VM controllers but they must be the same kind. For example, a VMM domain can contain many VMware vCenters managing multiple controllers each running multiple VMs but it may not also contain SCVMM Managers. A VMM domain inventories controller elements (such as pNICs, vNICs, VM names, and so forth) and pushes policies into the controller(s), creating port groups, and other necessary elements. The ACI VMM domain listens for controller events such as VM mobility and responds accordingly.

## VMM Domain VLAN Pool Association

VLAN pools represent blocks of traffic VLAN identifiers. A VLAN pool is a shared resource and can be consumed by multiple domains such as VMM domains and Layer 4 to Layer 7 services.

Each pool has an allocation type (static or dynamic), defined at the time of its creation. The allocation type determines whether the identifiers contained in it will be used for automatic assignment by the APIC (dynamic) or set explicitly by the administrator (static). By default, all blocks contained within a VLAN pool have the same allocation type as the pool but users can change the allocation type for encapsulation blocks contained in dynamic pools to static. Doing so excludes them from dynamic allocation.

A VMM domain can associate with only one dynamic VLAN pool. By default, the assignment of VLAN identifiers to EPGs that are associated with VMM domains is done dynamically by the APIC. While dynamic allocation is the default and preferred configuration, an administrator can statically assign a VLAN identifier


to an EPG instead. In that case, the identifiers used must be selected from encapsulation blocks in the VLAN pool associated with the VMM domain, and their allocation type must be changed to static.

The APIC provisions VMM domain VLAN on leaf ports based on EPG events, either statically binding on leaf ports or based on VM events from controllers such as VMware vCenter or Microsoft SCVMM.

## VMM Domain EPG Association

The ACI fabric associates tenant application profile EPGs to VMM domains, either automatically by an orchestration component such as Microsoft Azure, or by an APIC administrator creating such configurations. An EPG can span multiple VMM domains and a VMM domain can contain multiple EPGs.

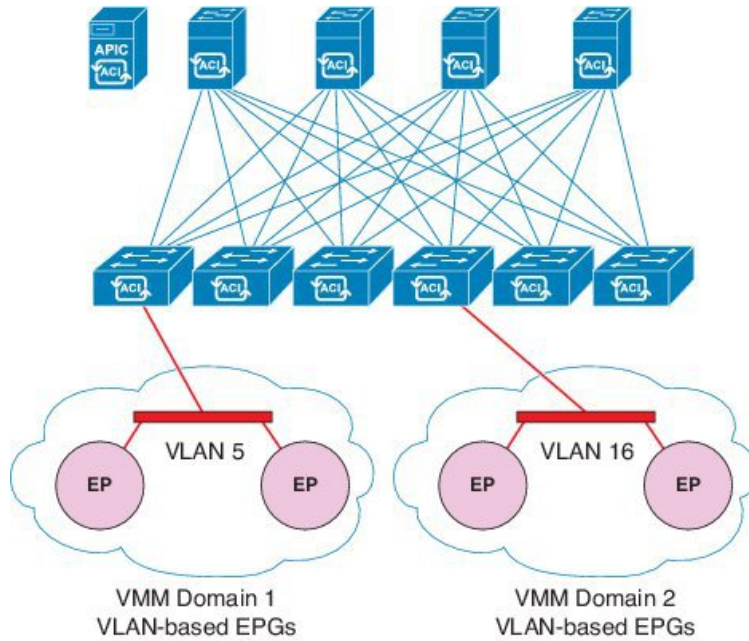
**Figure 2: VMM Domain EPG Association**



In the illustration above, end points (EP) of the same color are part of the same end point group. For example, all the green EPs are in the same EPG even though they are in two different VMM domains.

Refer to the latest Verified Scalability Guide for Cisco ACI document for virtual network and VMM domain EPG capacity information.

Figure 3: VMM Domain EPG VLAN Consumption



**Note** Multiple VMM domains can connect to the same leaf switch if they do not have overlapping VLAN pools on the same port. Similarly, the same VLAN pools can be used across different domains if they do not use the same port of a leaf switch.

EPGs can use multiple VMM domains in the following ways:

- An EPG within a VMM domain is identified by using an encapsulation identifier that is either automatically managed by the APIC, or statically selected by the administrator. An example is a VLAN, a Virtual Network ID (VNID).
- An EPG can be mapped to multiple physical (for baremetal servers) or virtual domains. It can use different VLAN or VNID encapsulations in each domain.




**Note** By default, the APIC dynamically manages allocating a VLAN for an EPG. VMware DVS administrators have the option to configure a specific VLAN for an EPG. In that case, the VLAN is chosen from a static allocation block within the pool associated with the VMM domain.

Applications can be deployed across VMM domains.



Figure 4: Multiple VMM Domains and Scaling of EPGs in the Fabric



While live migration of VMs within a VMM domain is supported, live migration of VMs across VMM domains is not supported.

## About Trunk Port Group

A trunk port group is used to aggregate the traffic of EPGs. Currently, it is supported under a VMware domain only. The trunk port group's naming scheme does not follow an EPG's T|A|E format. The name can be any ASCII string, as a trunk port group is not tenant-aware.

The aggregation of EPGs under the same domain is based on a VLAN range, which is specified as encapsulation blocks contained in the trunk port group. Whenever a EPG's encapsulation is changed or a trunk port group's encapsulation block is changed, the aggregation will be re-evaluated to determine if the EPG should be aggregated. A trunk port group controls the deployment in leafs of network resources, such as VLANs, allocated to EPGs being aggregated, including both the base EPG and uSeg EPG. In the case of a uSeg EPG, the trunk port group's VLAN ranges need to include both the primary and secondary VLANs.



**Note** Cisco ACI does not support IP fragmentation. Therefore, when you configure Layer 3 Outside (L3Out) connections to external routers, or multipod connections through an Inter-Pod Network (IPN), it is critical that the MTU is set appropriately on both sides. On some platforms, such as ACI, Cisco NX-OS, and Cisco IOS, the configurable MTU value takes into account the IP headers (resulting in a max packet size to be set as 9216 bytes for ACI and 9000 for NX-OS and IOS). However, other platforms such as IOS-XR configure the MTU value exclusive of packet headers (resulting in a max packet size of 8986 bytes).

For the appropriate MTU values for each platform, see the relevant configuration guides.

Cisco highly recommends you test the MTU using CLI-based commands. For example, on the Cisco NX-OS CLI, use a command such as `ping 1.1.1.1 df-bit packet-size 9000 source-interface ethernet 1/1.`

**Caution**

If you install 1 Gigabit Ethernet (GE) or 10GE links between the leaf and spine switches in the fabric, there is risk of packets being dropped instead of forwarded, because of inadequate bandwidth. To avoid the risk, use 40GE or 100GE links between the leaf and spine switches.

**Note**

Multiple Spanning Tree (MST) is not supported on interfaces configured with the Per Port VLAN feature (configuring multiple EPGs on a leaf switch using the same VLAN ID with localPort scope).

**Note**

If you are using Cisco ACI Multi-Site with this Cisco APIC cluster/fabric, look for a cloud icon on the object names in the navigation bar. This indicates that the information is derived from Multi-Site. It is recommended to only make changes from the Multi-Site GUI. Please review the Multi-Site documentation before making changes here.

**Note**

For a Cisco APIC REST API query of event records, the APIC system limits the response to a maximum of 500,000 event records. If the response is more than 500,000 events, it returns an error. Use filters to refine your queries. For more information, see [Composing Query Filter Expressions](#).

For more information, see

- [Creating a Trunk Port Group Using the GUI](#)
- [Creating a Trunk Port Group Using the NX-OS Style CLI](#)
- [Creating a Trunk Port Group Using the REST API](#)

## Attachable Entity Profile

The 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 (for example, the Cisco UCS fabric interconnect), or Layer 3 routers (for example Cisco Nexus 7000 Series switches). These attachment points can be physical ports, FEX ports, port channels, or a virtual port channel (vPC) on leaf switches.





**Note** When creating a VPC domain between two leaf switches, both switches must be in the same switch generation, one of the following:

- Generation 1 - Cisco Nexus N9K switches without “EX” or “FX” on the end of the switch name; for example, N9K-9312TX
- Generation 2 – Cisco Nexus N9K switches with “EX” or “FX” on the end of the switch model name; for example, N9K-93108TC-EX

Switches such as these two are not compatible VPC peers. Instead, use switches of the same generation.

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, such as Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), or Link Aggregation Control Protocol (LACP).

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.

The following AEP requirements and dependencies must be accounted for in various configuration scenarios, including network connectivity, VMM domains, and multipod configuration:

- The AEP defines the range of allowed VLANs but it does not provision them. No traffic flows unless an EPG is deployed on the port. Without defining a VLAN pool in an AEP, a VLAN is not enabled on the leaf port even if an EPG is provisioned.
- A particular VLAN is provisioned or enabled on the leaf port that is based on EPG events either statically binding on a leaf port or based on VM events from external controllers such as VMware vCenter or Microsoft Azure Service Center Virtual Machine Manager (SCVMM).
- Attached entity profiles can be associated directly with application EPGs, which deploy the associated application EPGs to all those ports associated with the attached entity profile. The AEP has a configurable generic function (infraGeneric), which contains a relation to an EPG (infraRsFuncToEpg) that is deployed on all interfaces that are part of the selectors that are associated with the attachable entity profile.

A virtual machine manager (VMM) domain automatically derives physical interface policies from the interface policy groups of an AEP.

An override policy at the AEP can be used to specify a different physical interface policy for a VMM domain. This policy is useful in scenarios where a VM controller is connected to the leaf switch through an intermediate Layer 2 node, and a different policy is desired at the leaf switch and VM controller physical ports. For example, you can configure LACP between a leaf switch and a Layer 2 node. At the same time, you can disable LACP between the VM controller and the Layer 2 switch by disabling LACP under the AEP override policy.

## EPG Policy Resolution and Deployment Immediacy

Whenever an EPG associates to a VMM domain, the administrator can choose the resolution and deployment preferences to specify when a policy should be pushed into leaf switches.

### Resolution Immediacy

- **Pre-provision**—Specifies that a policy (for example, VLAN, VXLAN binding, contracts, or filters) is downloaded to a leaf switch even before a VM controller is attached to the virtual switch (for example, VMware VDS). This pre-provisions the configuration on the switch.

This helps the situation where management traffic for hypervisors/VM controllers are also using the virtual switch associated to APIC VMM domain (VMM switch).

Deploying a VMM policy such as VLAN on ACI leaf switch requires APIC to collect CDP/LLDP information from both hypervisors via VM controller and ACI leaf switch. However if VM Controller is supposed to use the same VMM policy (VMM switch) to communicate with its hypervisors or even APIC, the CDP/LLDP information for hypervisors can never be collected because the policy required for VM controller/hypervisor management traffic is not deployed yet.

When using pre-provision immediacy, policy is downloaded to ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host connected to the VMM switch.

- **Immediate**—Specifies that EPG policies (including contracts and filters) are downloaded to the associated leaf switch software upon ESXi host attachment to a DVS. LLDP or OpFlex permissions are used to resolve the VM controller to leaf node attachments.

The policy will be downloaded to leaf when you add host to the VMM switch. CDP/LLDP neighborship from host to leaf is required.

- **On Demand**—Specifies that a policy (for example, VLAN, VXLAN bindings, contracts, or filters) is pushed to the leaf node only when an ESXi host is attached to a DVS and a VM is placed in the port group (EPG).

The policy will be downloaded to leaf when host is added to VMM switch and virtual machine needs to be placed into port group (EPG). CDP/LLDP neighborship from host to leaf is required.

With both immediate and on demand, if host and leaf lose LLDP/CDP neighborship the policies are removed.

### Deployment Immediacy

Once the policies are downloaded to the leaf software, deployment immediacy can specify when the policy is pushed into the hardware policy content-addressable memory (CAM).

- **Immediate**—Specifies that the policy is programmed in the hardware policy CAM as soon as the policy is downloaded in the leaf software.
- **On demand**—Specifies that the policy is programmed in the hardware policy CAM only when the first packet is received through the data path. This process helps to optimize the hardware space.



#### Note

When you use on demand deployment immediacy with MAC-pinned VPCs, the EPG contracts are not pushed to the leaf ternary content-addressable memory (TCAM) until the first endpoint is learned in the EPG on each leaf. This can cause uneven TCAM utilization across VPC peers. (Normally, the contract would be pushed to both peers.)

# Guidelines for Deleting VMM Domains

Follow the sequence below to assure that the APIC request to delete a VMM domain automatically triggers the associated VM controller (for example VMware vCenter or Microsoft SCVMM) to complete the process normally, and that no orphan EPGs are stranded in the ACI fabric.

1. The VM administrator must detach all the VMs from the port groups (in the case of VMware vCenter) or VM networks (in the case of SCVMM), created by the APIC.

In the case of Cisco AVS, the VM admin also needs to delete vmk interfaces associated with the Cisco AVS.

2. The ACI administrator deletes the VMM domain in the APIC. The APIC triggers deletion of VMware VDS or Cisco AVS or SCVMM logical switch and associated objects.



---

**Note** The VM administrator should not delete the virtual switch or associated objects (such as port groups or VM networks); allow the APIC to trigger the virtual switch deletion upon completion of step 2 above. EPGs could be orphaned in the APIC if the VM administrator deletes the virtual switch from the VM controller before the VMM domain is deleted in the APIC.

---

If this sequence is not followed, the VM controller does delete the virtual switch associated with the APIC VMM domain. In this scenario, the VM administrator must manually remove the VM and vtep associations from the VM controller, then delete the virtual switch(es) previously associated with the APIC VMM domain.

## NetFlow with Virtual Machine Networking

### About NetFlow with Virtual Machine Networking

The NetFlow technology provides the metering base for a key set of applications, including network traffic accounting, usage-based network billing, network planning, as well as denial of services monitoring, network monitoring, outbound marketing, and data mining for both service providers and enterprise customers. Cisco provides a set of NetFlow applications to collect NetFlow export data, perform data volume reduction, perform post-processing, and provide end-user applications with easy access to NetFlow data. If you have enabled NetFlow monitoring of the traffic flowing through your datacenters, this feature enables you to perform the same level of monitoring of the traffic flowing through the Cisco Application Centric Infrastructure (Cisco ACI) fabric.

Instead of hardware directly exporting the records to a collector, the records are processed in the supervisor engine and are exported to standard NetFlow collectors in the required format.

For more information about NetFlow, see the *Cisco APIC and NetFlow* knowledge base article.

## About NetFlow Exporter Policies with Virtual Machine Networking

A virtual machine manager exporter policy (netflowVmmExporterPol) describes information about the data collected for a flow that is sent to the reporting server or NetFlow collector. A NetFlow collector is an external entity that supports the standard NetFlow protocol and accepts packets marked with valid NetFlow headers.

An exporter policy has the following properties:

- **VmmExporterPol.dstAddr**—This mandatory property specifies the IPv4 or IPv6 address of the NetFlow collector that accepts the NetFlow flow packets. This must be in the host format (that is, "/32" or "/128"). An IPv6 address is supported in vSphere Distributed Switch (vDS) version 6.0 and later.
- **VmmExporterPol.dstPort**—This mandatory property specifies the port on which the NetFlow collector application is listening on, which enables the collector to accept incoming connections.
- **VmmExporterPol.srcAddr**—This optional property specifies the IPv4 address that is used as the source address in the exported NetFlow flow packets.

## NetFlow Support with VMware vSphere Distributed Switch

The VMware vSphere Distributed Switch (VDS) supports NetFlow with the following caveats:

- The external collector must be reachable through the ESX. ESX does not support virtual routing and forwardings (VRFs).
- A port group can enable or disable NetFlow.
- VDS does not support flow-level filtering.

Configure the following VDS parameters in VMware vCenter:

- Collector IP address and port. IPv6 is supported on VDS version 6.0 or later. These are mandatory.
- Source IP address. This is optional.
- Active flow timeout, idle flow timeout, and sampling rate. These are optional.

## NetFlow Support with Cisco Application Virtual Switch

The Cisco Application Virtual Switch (AVS) supports NetFlow with the following caveats:

- The external collector must be reachable through the ESX. ESX does not support virtual routing and forwardings (VRFs).
- A port group can enable or disable NetFlow and specify the direction of the traffic to be collected.
- Cisco AVS does not support flow-level filtering.

## Configuring a NetFlow Exporter Policy for VM Networking Using the GUI

The following procedure configures a NetFlow exporter policy for VM networking.

### Procedure

---

- Step 1** On the menu bar, choose **Fabric > Access Policies**.
  - Step 2** In the **Navigation** pane, expand **Interface Policies > Policies > NetFlow**.
  - Step 3** Right-click **NetFlow Exporters for VM Networking** and choose **Create NetFlow Exporter for VM Networking**.
  - Step 4** In the **Create NetFlow Exporter for VM Networking** dialog box, fill in the fields as required.
  - Step 5** Click **Submit**.
- 

## Consuming a NetFlow Exporter Policy Under a VMM Domain Using the GUI

The following procedure consumes a NetFlow exporter policy under a VMM domain using the GUI.

### Procedure

---

- Step 1** On the menu bar, choose **Virtual Networking > Inventory**.
  - Step 2** In the **Navigation** pane, expand the **VMM Domains** folder, right-click **VMware**, and choose **Create vCenter Domain**.
  - Step 3** In the **Create vCenter Domain** dialog box, fill in the fields as required, except as specified:
    - a) In the **NetFlow Exporter Policy** drop-down list, choose the desired exporter policy or create a new one.
    - b) In the **Active Flow Timeout** field, enter the desired active flow timeout, in seconds.

The **Active Flow Timeout** parameter specifies the delay that NetFlow waits after the active flow is initiated, after which NetFlow sends the collected data. The range is from 60 to 3600. The default value is 60.
    - c) In the **Idle Flow Timeout** field, enter the desired idle flow timeout, in seconds.

The **Idle Flow Timeout** parameter specifies the delay that NetFlow waits after the idle flow is initiated, after which NetFlow sends the collected data. The range is from 10 to 300. The default value is 15.
    - d) (VDS only) In the **Sampling Rate** field, enter the desired sampling rate.

The **Sampling Rate** parameter specifies how many packets that NetFlow will drop after every collected packet. If you specify a value of 0, then NetFlow does not drop any packets. The range is from 0 to 1000. The default value is 0.
  - Step 4** Click **Submit**.
- 

## Enabling NetFlow on an Endpoint Group to VMM Domain Association Using the GUI

The following procedure enables NetFlow on an endpoint group to VMM domain association.

**Before you begin**

You must have configured the following:

- An application profile
- An application endpoint group

**Procedure**

- 
- Step 1** On the menu bar, choose **Tenants > All Tenants**.
- Step 2** In the **Work** pane, double-click the tenant's name.
- Step 3** In the left navigation pane, expand *tenant\_name* > **Application Profiles** > *application\_profile\_name* > **Application EPGs** > *application\_EPG\_name*
- Step 4** Right-click **Domains (VMs and Bare-Metals)** and choose **Add VMM Domain Association**.
- Step 5** In the **Add VMM Domain Association** dialog box, fill in the fields as required, except as specified below:
- In the **NetFlow** are, choose **Enable**.
  - (Cisco AVS only) Click **NetFlow Direction** and choose **ingress**, **egress**, or **both** for the flows that need to be monitored and collected.
- Step 6** Click **Submit**.
- 

## Configuring a NetFlow Exporter Policy for Virtual Machine Networking Using the NX-OS-Style CLI

The following example procedure uses the NX-OS-style CLI to configure a NetFlow exporter policy for virtual machine networking.

**Procedure**

- 
- Step 1** Enter the configuration mode.
- Example:**
- ```
apicl# config
```
- Step 2** Configure the exporter policy.
- Example:**
- ```
apicl(config)# flow vm-exporter vmExporter1 destination address 2.2.2.2 transport udp 1234
apicl(config-flow-vm-exporter)# source address 4.4.4.4
apicl(config-flow-vm-exporter)# exit
apicl(config)# exit
```
-

## Consuming a NetFlow Exporter Policy Under a VMM Domain Using the NX-OS-Style CLI for VMware VDS

The following procedure uses the NX-OS-style CLI to consume a NetFlow exporter policy under a VMM domain.

### Procedure

---

**Step 1** Enter the configuration mode.

**Example:**

```
apic1# config
```

**Step 2** Consume the NetFlow exporter policy.

**Example:**

```
apic1(config)# vmware-domain mininet
apic1(config-vmware)# configure-dvs
apic1(config-vmware-dvs)# flow exporter vmExporter1
apic1(config-vmware-dvs-flow-exporter)# active-flow-timeout 62
apic1(config-vmware-dvs-flow-exporter)# idle-flow-timeout 16
apic1(config-vmware-dvs-flow-exporter)# sampling-rate 1
apic1(config-vmware-dvs-flow-exporter)# exit
apic1(config-vmware-dvs)# exit
apic1(config-vmware)# exit
apic1(config)# exit
```

---

## Consuming a NetFlow Exporter Policy Under a VMM Domain Using the NX-OS-Style CLI for Cisco AVS

The following procedure uses the NX-OS-style CLI to consume a NetFlow exporter policy under a VMM domain.

### Procedure

---

**Step 1** Enter the configuration mode.

**Example:**

```
apic1# config
```

**Step 2** Consume the NetFlow exporter policy.

**Example:**

```
apic1(config)# vmware-domain mininet
apic1(config-vmware)# configure-avs
apic1(config-vmware-dvs)# flow exporter vmExporter1
apic1(config-vmware-dvs-flow-exporter)# active-flow-timeout 62
apic1(config-vmware-dvs-flow-exporter)# idle-flow-timeout 16
apic1(config-vmware-dvs-flow-exporter)# exit
```



```

apicl(config-vmware-dvs)# exit
apicl(config-vmware)# exit
apicl(config)# exit

```

---

## Enabling or Disabling NetFlow on an Endpoint Group Using the NX-OS-Style CLI for VMware VDS

The following procedure enables or disables NetFlow on an endpoint group using the NX-OS-style CLI.

### Procedure

---

**Step 1** Enable NetFlow:

#### Example:

```

apicl# config
apicl(config)# tenant tn1
apicl(config-tenant)# application appl
apicl(config-tenant-app)# epg epg1
apicl(config-tenant-app-epg)# vmware-domain member mininet
apicl(config-tenant-app-epg-domain)# flow monitor enable
apicl(config-tenant-app-epg-domain)# exit
apicl(config-tenant-app-epg)# exit
apicl(config-tenant-app)# exit
apicl(config-tenant)# exit
apicl(config)# exit

```

**Step 2** (Optional) If you no longer want to use NetFlow, disable the feature:

#### Example:

```

apicl(config-tenant-app-epg-domain)# no flow monitor enable

```

---

## Enabling or Disabling NetFlow on an Endpoint Group Using the NX-OS-Style CLI for Cisco AVS

The following procedure enables or disables NetFlow on an endpoint group using the NX-OS-style CLI.

### Procedure

---

**Step 1** Enable NetFlow:

#### Example:

```

apicl# config
apicl(config)# tenant tn1
apicl(config-tenant)# application appl
apicl(config-tenant-app)# epg epg1
apicl(config-tenant-app-epg)# vmware-domain member mininet
apicl(config-tenant-app-epg-domain)# flow monitor enable

```

```

apic1(config-tenant-app-epg-domain)#flow direction {ingress | egress | both}
apic1(config-tenant-app-epg-domain)# exit
apic1(config-tenant-app-epg)# exit
apic1(config-tenant-app)# exit
apic1(config-tenant)# exit
apic1(config)# exit

```

**Step 2** (Optional) If you no longer want to use NetFlow, disable the feature:

**Example:**

```

apic1(config-tenant-app-epg-domain)# no flow monitor enable

```

## Configuring a NetFlow Exporter Policy for VM Networking Using the REST API

The following example XML shows how to configure a NetFlow exporter policy for VM networking using the REST API:

```

<polUni>
  <infraInfra>
    <netFlowVmmExporterPol name="vmExporter1" dstAddr="2.2.2.2" dstPort="1234"
srcAddr="4.4.4.4"/>
  </infraInfra>
</polUni>

```

## Consuming a NetFlow Exporter Policy Under a VMM Domain Using the REST API for VMware VDS

The following example XML shows how to consume a NetFlow exporter policy under a VMM domain using the REST API:

```

<polUni>
  <vmmProvP vendor="VMware">
    <vmmDomP name="mininet">
      <vmmVSwitchPolicyCont>
        <vmmRsVswitchExporterPol tDn="uni/infra/vmmexporterpol-vmExporter1"
activeFlowTimeOut="62" idleFlowTimeOut="16" samplingRate="1"/>
      </vmmVSwitchPolicyCont>
    </vmmDomP>
  </vmmProvP>
</polUni>

```

## Consuming a NetFlow Exporter Policy Under a VMM Domain Using the REST API for Cisco AVS

### Procedure

To consume a NetFlow exporter policy under a VMM domain, send a POST message like the following example:

**Example:**

```
<polUni>
  <vmmProvP vendor="VMware">
    <vmmDomP name="mininet">
      <vmmVSwitchPolicyCont>
        <vmmRsVswitchExporterPol tDn="uni/infra/vmmexporterpol-vmExporter1"
activeFlowTimeOut="62" idleFlowTimeOut="16"/>
      </vmmVSwitchPolicyCont>
    </vmmDomP>
  </vmmProvP>
</polUni>
```

## Enabling NetFlow on an Endpoint Group for VMM Domain Association for VMware VDS

The following example XML shows how to enable NetFlow on an endpoint group for VMM domain association using the REST APIs:

```
<polUni>
  <fvTenant name="t1">
    <fvAp name="a1">
      <fvAEPg name="EPG1">
        <fvRsDomAtt tDn="uni/vmmp-VMware/dom-mininet" netflowPref="enabled" />
      </fvAEPg>
    </fvAp>
  </fvTenant>
</polUni>
```

## Enabling NetFlow on an Endpoint Group for VMM Domain Association for Cisco AVS

**Procedure**

Enable NetFlow on an endpoint group for VMM domain association by sending a POST message similar to the following example:

**Example:**

**Note** The example specifies NetFlow direction as "ingress." Alternatively, you can choose "egress" or "both."

```
<polUni>
  <fvTenant name="t1">
    <fvAp name="a1">
      <fvAEPg name="EPG1">
        <fvRsDomAtt tDn="uni/vmmp-VMware/dom-mininet" netflowPref="enabled"
netflowDir="ingress"/>
      </fvAEPg>
    </fvAp>
  </fvTenant>
</polUni>
```

# Troubleshooting VMM Connectivity

The following procedure resolves VMM connectivity issues:

## Procedure

---

- Step 1** Trigger inventory resync on the Application Policy Infrastructure Controller (APIC).  
For more information about how to trigger an inventory resync on APIC, see the following knowledge base article:  
[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b\\_KB\\_VMM\\_OnDemand\\_Inventory\\_in\\_APIC.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_VMM_OnDemand_Inventory_in_APIC.html)
- Step 2** If step 1 does not fix the issue, for the impacted EPGs, set the resolution immediacy to use preprovisioning in the VMM domain.  
"Pre-Provision" removes the need for neighbor adjacencies or OpFlex permissions and subsequently the dynamic nature of VMM Domain VLAN Programming. For more information about Resolution Immediacy types, see the following EPG Policy Resolution and Deployment Immediacy section:  
[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b\\_ACI-Fundamentals/b\\_ACI-Fundamentals\\_chapter\\_01011.html#concept\\_EF87ADDAD4EF47BDA741EC6EFDAECBBD](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html#concept_EF87ADDAD4EF47BDA741EC6EFDAECBBD)
- Step 3** If steps 1 and 2 do not fix the issue and you see the issue on all of the VMs, then delete the VM controller policy and readd the policy.
- Note** Deleting the controller policy impacts traffic for all VMs that are on that controller.
-

