Cisco ACI Virtual Machine Networking

This chapter contains the following sections:

Cisco ACI VM Networking Support for Virtual Machine Managers

Benefits of ACI VM Networking

Cisco Application Centric Infrastructure (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 that are 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 Products and Vendors

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

- **Cisco Unified Computing System Manager (UCSM)**

  Integration of Cisco UCSM is supported beginning in Cisco Cisco APIC Release 4.1(1). For information, see the chapter "Cisco ACI with Cisco UCSM Integration in the *Cisco ACI Virtualization Guide, Release 4.1(1)*.

- **Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod)**

  Cisco ACI vPod is in general availability beginning in Cisco APIC Release 4.0(2). For information, see the Cisco ACI vPod 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 Found 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* on Cisco.com

- **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.

Mapping Cisco ACI and VMware Constructs

Cisco Application Centric Infrastructure (ACI) and VMware use different terms to describe the same constructs. This section provides a table for mapping Cisco ACI and VMware terminology; the information is relevant to VMware vSphere Distributed Switch (VDS).

Cisco ACI Terms	VMware Terms
Endpoint group (EPG)	Port group, portgroup

Cisco ACI Terms	VMware Terms
LACP Active	• Route based on IP hash (downlink port group) • LACP Enabled/Active (uplink port group)
LACP Passive	• Route based on IP hash (downlink port group) • LACP Enabled/Active (uplink port group)
MAC Pinning	• Route based on originating virtual port • LACP Disabled
MAC Pinning-Physical-NIC-Load	• Route based on physical NIC load • LACP Disabled
Static Channel - Mode ON	• Route based on IP Hash (downlink port group) • LACP Disabled
Virtual Machine Manager (VMM) domain	VDS
VM controller	vCenter (Datacenter)

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**—Specifes 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.

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 Cisco 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, and for encapsulation blocks contained in static pools to dynamic. Entries from blocks with static allocation type are excluded from dynamic allocation.

A VMM domain can associate with only one dynamic VLAN pool. Beginning with APIC release 6.0(4), a VMM domain for VMware vDS can associate a static VLAN pool instead of a dynamic VLAN pool. However, a static VLAN pool remains not supported with other types of VMM domains. By default, the assignment of

VLAN identifiers to EPGs that are associated with VMM domains is done dynamically by the Cisco APIC. While dynamic allocation is the default and preferred configuration, an administrator can statically assign a VLAN identifier to an endpoint group (EPG) instead. In that case, the identifiers used must be selected from encapsulation blocks with static allocation type in the VLAN pool.

When migrating from a physical workload to a VM (with an existing VLAN pool), it is recommended that the VMM domain reference the same VLAN pool used by the physical domain. In case of Layer 4 to Layer 7 devices, to support VMM integration, the VLAN pool referenced must be a dynamic pool. Dynamic pool is required for creating auto-creating EPGs when service graphs are deployed. You can use static pools for Layer 4 to Layer 7 devices which do not use the ACI service graph functionality.

The Cisco 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.

Note In dynamic VLAN pools, if a VLAN is disassociated from an EPG, it is automatically reassociated with the EPG in five minutes.

Note Dynamic VLAN association is not a part of configuration rollback, that is, in case an EPG or tenant was initially removed and then restored from the backup, a new VLAN is automatically allocated from the dynamic VLAN pools.

VMM Domain EPG Association

The Cisco Application Centric Infrastructure (ACI) fabric associates tenant application profile endpoint groups (EPGs) to virtual machine manager (VMM) domains, The Cisco ACI does so either automatically by an orchestration component such as Microsoft Azure, or by a Cisco Application Policy Infrastructure Controller (APIC) administrator creating such configurations. An EPG can span multiple VMM domains, and a VMM domain can contain multiple EPGs.

Figure 1: VMM Domain EPG Association

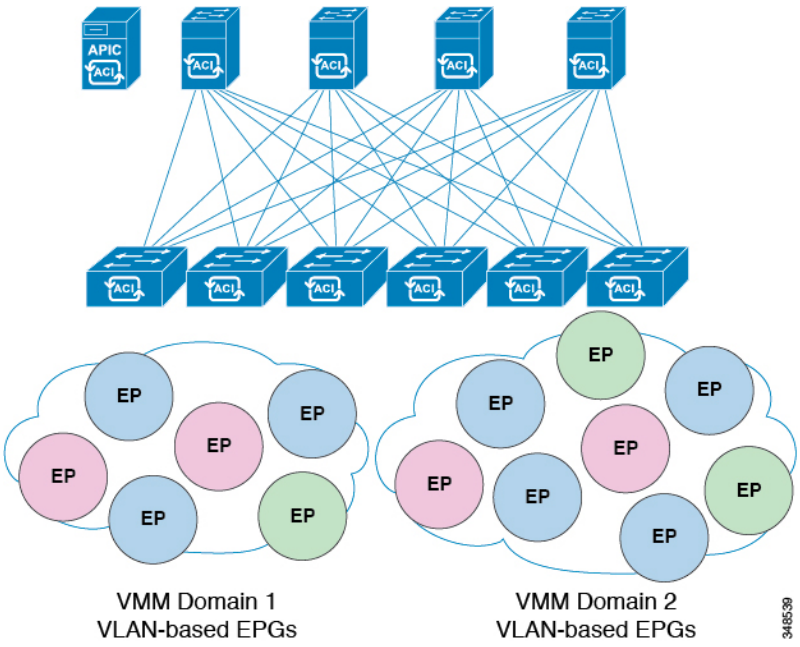

In the preceding illustration, end points (EPs) of the same color are part of the same EPG. For example, all the green EPs are in the same EPG although they are in two different VMM domains.

See the latest *Verified Scalability Guide for Cisco ACI* for virtual network and VMM domain EPG capacity information.

Figure 2: VMM Domain EPG VLAN Consumption

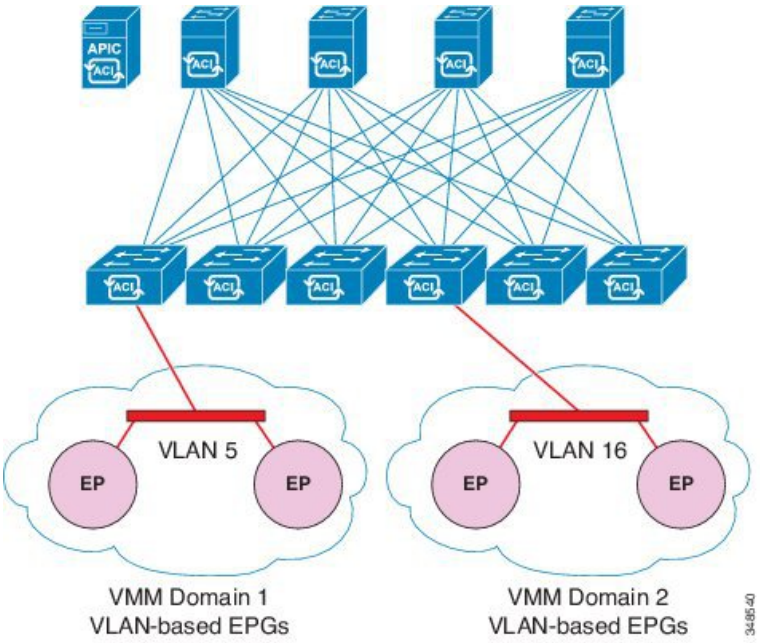

Note When multiple VMM domains with an overlapping VLAN ID range are connected to the same leaf switch, those domains should use the same VLAN pool. With the same VLAN pool, Cisco APIC can make sure to pick a different VLAN ID for each domain-to-EPG association. Otherwise, Cisco APIC might pick a VLAN ID that is already used on the switch for another domain-to-EPG association, which causes the VLAN deployment fail.

When multiple VMM domains with an overlapping VLAN ID range are connected to the same leaf switch and those domains use the same VLAN pool, you can have multiple VMM domains associated with the same EPG. However, each domain-to-EPG association deploys a different VLAN ID, respectively, even though the VLANs are for the same EPG and potentially are on the same port. If using VLAN IDs in this manner is suboptimal to your requirements, you can use the same VMM domain with multiple VMM controllers instead of having multiple VMM domains.

EPGs can use multiple VMM domains in the following ways:

- An EPG within a VMM domain is identified by using an encapsulation identifier. Cisco APIC can manage the identifier automatically, or the administrator can statically select it. 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 Cisco APIC dynamically manages the allocatiion of 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 that is associated with the VMM domain.

Applications can be deployed across VMM domains.

Figure 3: Multiple VMM Domains and Scaling of EPGs in the Fabric

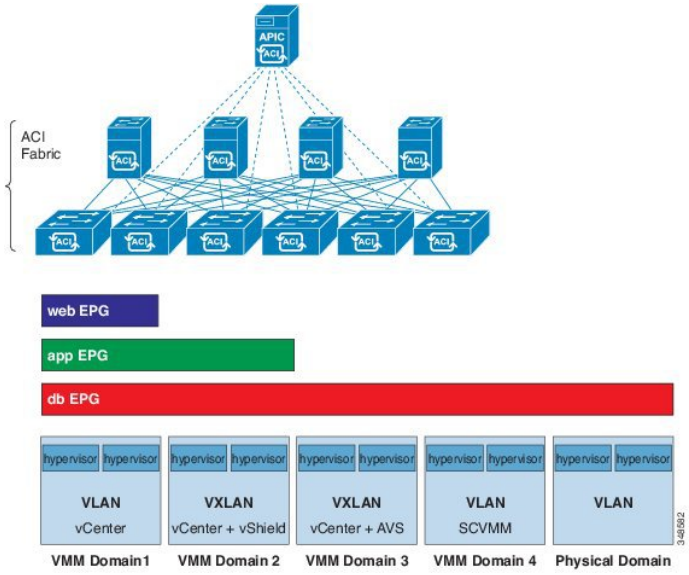

While live migration of VMs within a VMM domain is supported, live migration of VMs across VMM domains is not supported.

Note When you change the VRF on a bridge domain that is linked to an EPG with an associated VMM domain, the port-group is deleted and then added back on vCenter. This results in the EPG being undeployed from the VMM domain. This is expected behavior.

About Trunk Port Group

You use a trunk port group to aggregate the traffic of endpoint groups (EPGs) for VMware virtual machine manager (VMM) domains. Unlike regular port groups, which are configured under the Tenants tab in the Cisco Application Policy Infrastructure Controller (APIC) GUI, trunk port groups are configured under the VM Networking tab. Regular port groups follow an EPG's *T*|*A*|*E* name format.

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 is be re-evaluated to determine if the EGP should be aggregated.

A trunk port group controls the leaf deployment of network resources, such as VLANs, that allocated to the EPGs being aggregated. The EPGs include both base EPG and microsegmented (uSeg) EPGs. 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 Multi-Pod connections through an Inter-Pod Network (IPN), it is recommended that the interface MTU is set appropriately on both ends of a link. On some platforms, such as Cisco ACI, Cisco NX-OS, and Cisco IOS, the configurable MTU value does not take into account the Ethernet headers (matching IP MTU, and excluding the 14-18 Ethernet header size), while other platforms, such as IOS-XR, include the Ethernet header in the configured MTU value. A configured value of 9000 results in a max IP packet size of 9000 bytes in Cisco ACI, Cisco NX-OS, and Cisco IOS, but results in a max IP packet size of 8986 bytes for an IOS-XR untagged interface.

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

We highly recommend that 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 Cisco 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* in the *Cisco APIC REST API Configuration Guide*.

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 endpoint group (EPG) associates to a virtual machine manager (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 vSphere Distributed Switch (VDS). This pre-provisions the configuration on the switch.

  This helps the situation where management traffic for hypervisors/VM controllers is also using the virtual switch associated to the Cisco Application Policy Infrastructure Controller (APIC) VMM domain (VMM switch).

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

  When using pre-provision immediacy, policy is downloaded to Cisco ACI leaf switch regardless of CDP/LLDP neighborship. Even without a hypervisor host that is 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 the leaf when host is added to the VMM switch. The VM needs to be placed into a 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.

Note In OpFlex-based VMM domains, an OpFlex agent on the hypervisor reports a VM/EP virtual network interface card (vNIC) attachment to an EPG to the leaf OpFlex process. When using On Demand Resolution Immediacy, the EPG VLAN/VXLAN is programmed on **all** leaf port channel ports, virtual port channel ports, or both when the following are true:

- Hypervisors are connected to leafs on port channel or virtual port channel attached directly or through blade switches.
- A VM or instance vNIC is attached to an EPG.
- Hypervisors are attached as part of the EPG or VMM domain.

Opflex-based VMM domains are Microsoft Security Center Virtual Machine Manager (SCVMM) and HyperV, and Cisco Application Virtual Switch (AVS).

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-addressble 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.)

Dynamic EPG deployment

The current behavior (prior to Cisco APIC release 6.2(1)) of EPG association with a VMM domain using the "immediate" resolution mode results in the deployment of EPGs and their associated VLANs on all the leaf interfaces connected to the DVS ports. Especially, in case of single DVS spanning multiple clusters, the "immediate" resolution mode deploys EPGs to all DVS-connected leaf interfaces, irrespective of whether traffic for a given EPG is expected to pass through those ports. This may result in unnecessary resource consumption on some nodes and ports.

Beginning with APIC release 6.2(1), dynamic cluster-aware EPG deployment is supported. The dynamic EPG deployment through VMM integration with VMware vCenter, can now be aware of VMware ESXi cluster inside the VMware vCenter. This allows APIC to deploy the EPGs only to the ACI leaf switch interfaces connected to VMware ESXi hosts part of the VMware ESXi cluster selected by the user.

Cluster-aware deployment is only supported on Cisco VMware vCenter with APIC. Cluster creation must be done using the VMware vCenter. Those cluster names can then be used for EPG associations in Cisco APIC. If cluster names are not specified, regular *Immediate* deploy behavior applies. If cluster names specified in this association do not match clusters defined in the Datacenter, EPG deployment would be equivalent to *On-Demand* resolution mode.

Benefits of cluster-aware EPG deployment:

- EPGs are deployed only to leaf interfaces connected to hosts within the relevant cluster, thus optimizing resource utilization
- Improved scalability
- Reduces overhead in cases where a single-DVS is used across multiple clusters

Associate the cluster group

To associate the cluster group on which the EPGs are deployed, on the **Add VMM Domain Association** screen, click the (+) and select the cluster (fetched from VMware vCenter).

Note This is applicable only when you select the **Immediate** option for the **Resolution Immediacy** field.

Guidelines for Deleting VMM Domains

Follow the sequence below to assure that the Cisco Application Policy Infrastructure Controller (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 Cisco Application Centric Infrastructure (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 Cisco APIC.
2. The Cisco ACI administrator deletes the VMM domain in the Cisco APIC. The Cisco APIC triggers deletion of VMware VDS 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 Cisco APIC to trigger the virtual switch deletion upon completion of step 2 above. EPGs could be orphaned in the Cisco APIC if the VM administrator deletes the virtual switch from the VM controller before the VMM domain is deleted in the Cisco APIC.

If this sequence is not followed, the VM controller does delete the virtual switch associated with the Cisco 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 Cisco 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.

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 **Policies** > **Interface** > **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; however, in the **NetFlow** area, choose **Enable**.

Step 6 Click **Submit**.

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

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

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.