



Cisco ACI vPod: Enabling Outside Communication with Devices in the Remote Location

New and Changed Information 2

[Cisco ACI vPod: Enabling Outside Communication with Devices in the Remote Location](#) 2

[Cisco ACI vPod: Extending the Cisco ACI Fabric](#) 3

[Local L3Out Configuration for Cisco ACI vPod](#) 6

New and Changed Information

The following table provides an overview of the significant changes to this guide up to this current release. The table does not provide an exhaustive list of all changes that are made to the guide or of the new features up to this release.

Table 1: New Features and Changed Behavior

Cisco APIC Release Version	Feature	Description	Where Documented
4.2(3)	A local Layer 3 outside network connection (L3Out) is available for a Cisco Application Centric Infrastructure Virtual Edge that is connected to a Cisco ACI Virtual Pod.	WAN traffic from Cisco ACI Virtual Edge attached to a Cisco ACI Virtual Pod no longer must go through an L3Out on a physical leaf in the on-premises data center. You can now configure the fabric to send traffic to leave a Cisco Cloud Service Router (CSR) using Cisco GOLF in the remote location.	<ul style="list-style-type: none">• The section Local L3Out Configuration for Cisco ACI vPod, on page 6 in this guide• Cisco ACI Virtual Edge Release Notes, Release 4.2(1)

Cisco ACI vPod: Enabling Outside Communication with Devices in the Remote Location

Beginning in Cisco Application Policy Infrastructure Controller (APIC) Release 4.2(1), you can improve the efficiency of a network that is extended to a remote site. You can do so by configuring a Layer 3 outside network connection (L3Out) through Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) in the remote site.

Before Cisco APIC Release 4.2(1), to connect to outside networks, Cisco ACI vPod had to use an L3Out on a physical leaf in the on-premises data center. However, you can now configure an L3Out directly through a Cisco Cloud Services Router (CSR) using the Cisco ACI GOLF feature.

The new direct L3Out configuration is simpler and more efficient than communicating over the interpod network (IPN) connection to the on-premises data center. Connecting through an L3Out and Cisco Cloud Services Router in the remote location provides greater bandwidth and lower latency than connecting through the IPN to the physical leaf.

You configure the L3Out for Cisco ACI vPod as you configure an L3Out for Cisco ACI GOLF in Cisco APIC. This guide provides the following information:

- Information about Cisco ACI vPod
 - [Cisco ACI vPod: Extending the Cisco ACI Fabric](#), on page 3
 - [Cisco ACI vPod Architecture](#), on page 4
 - [Cisco ACI vPod Licensing](#), on page 5
 - [Cisco ACI vPod in a Multipod Environment](#), on page 6
- Procedures for configuring L3Out

- [Prerequisites for Cisco ACI vPod L3Out Configuration, on page 6](#)
 - [Workflow for Cisco ACI vPod L3Out Configuration, on page 7](#)
 - [Create a Tenant, VRF, and Bridge Domain, on page 7](#)
 - [Creating EPGs, on page 8](#)
 - [Configuring an L3Out Using Cisco ACI GOLF, on page 10](#)
- [Links to more Cisco ACI vPod and Cisco ACI documentation](#)

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

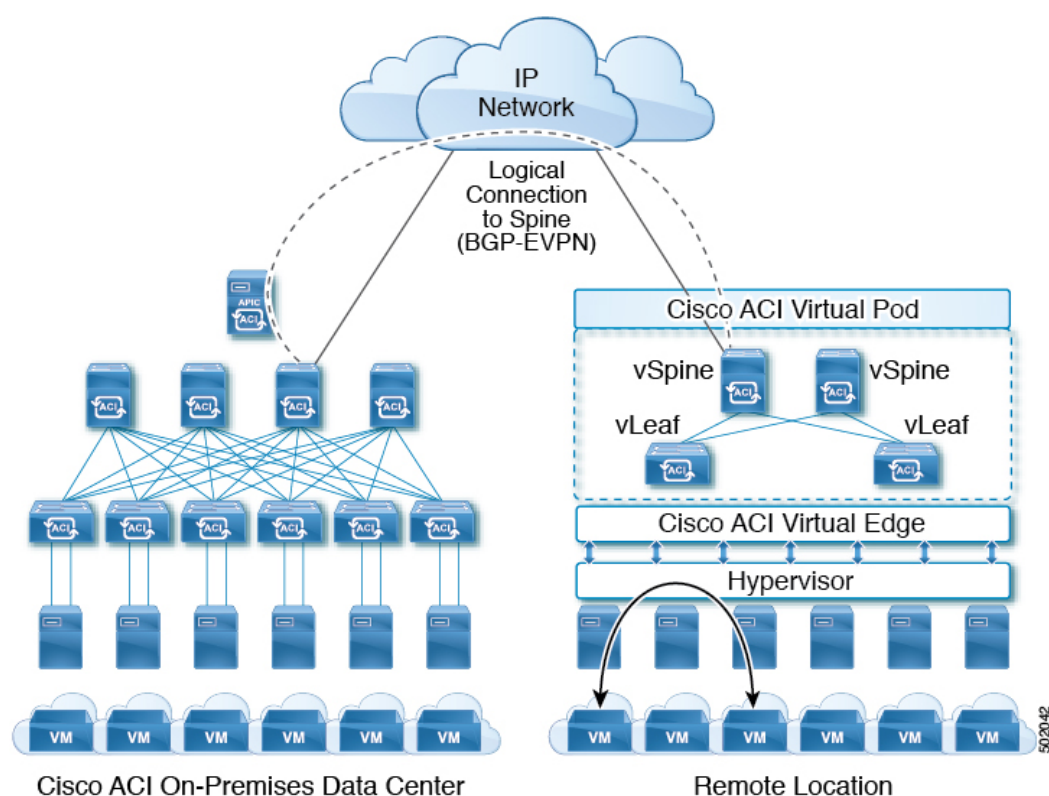
For general information, see the *Cisco ACI Virtual Pod Release Notes* on Cisco.com.

Cisco ACI vPod Architecture

Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) extends the multipod architecture into the virtual infrastructure. It consists of virtual spines (vSpines), virtual leafs (vLeafs), and one to 32 instances of Cisco ACI Virtual Edge. These elements are deployed on a VMware ESXi hypervisor in the remote location.

The vSpines and Cisco ACI Virtual Edge logically connect to the on-premises data center over an interpod network (IPN). The Cisco Application Policy Infrastructure Controller (APIC) cluster in the on-premises data center discovers, configures, and extends the policy on the Cisco ACI vPod.

Figure 1: Cisco ACI vPod in a Remote Location



Cisco ACI vPod includes the following elements:

- **Virtual spine (vSpine) and Virtual leaf (vLeaf):** The vSpine and vLeaf are separate virtual machines (VMs) within Cisco ACI vPod. They emulate the functionality of a physical spine and physical leaf. The vSpines and Cisco ACI Virtual Edge logically connect to the on-premises data center over an IPN.

The vSpine and vLeaf function as services inside the VMs in the Cisco ACI vPod, handling control plane management. They are contained in the Cisco ACI vPod installation package.


- **vSpine:** Establishes a Border Gateway Protocol (BGP) Ethernet VPN (EVPN) connection to the on-premises spine and the vPod control plane.

- **vLeaf:** Distributes the policies that are created on Cisco Application Policy Infrastructure Controller (APIC) to the Cisco ACI Virtual Edge.
- **Cisco ACI Virtual Edge:** Connects to spines and border leaves in the physical Cisco ACI fabric through Insieme VXLAN (iVXLAN) tunnels. It also uses iVXLAN for communication within the remote site. Cisco ACI Virtual Edge handles packet forwarding, policy enforcement, and all data plane management. The iVXLAN overlay technology in the data plane provides Layer 2 and Layer 3 connectivity services between endpoints, regardless of their location.

You deploy one Cisco ACI Virtual Edge on each VMware ESXi hypervisor host that participates in the virtual pod, creating a distributed data plane across the virtual pod. You can have up to 32 Cisco ACI Virtual Edge virtual leafs in a Cisco ACI vPod, one on each host.

Cisco ACI Virtual Edge uses the fast packet-processing framework Data Plane Development Kit (DPDK) for high-performance forwarding in the user space.

Figure 2: Cisco ACI vPod Components



A single Cisco ACI can contain multiple physical pods and multiple virtual pods. A single Cisco APIC manages the entire fabric. Tenant configurations—including virtual routing and forwarding (VRF), bridge domains, and endpoint groups (EPGs)—are available across all the pods.

Cisco ACI vPod Licensing

You obtain the following two licenses for Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) on Cisco.com:

- **Management Cluster:** You need a single license for each Cisco ACI vPod.

The management cluster is licensed as a pair of redundant virtual spine (vSpine) and virtual leaf (vLeaf) virtual machines (VMs). The cluster provides Cisco ACI management and control plane functions. It also programs the individual Cisco Application Centric Infrastructure Virtual Edge within the Cisco ACI vPod.

The single Cisco ACI vPod management cluster license includes the redundant vSpines and vLeafs, a total of four VMs.



Note VMware ESXi 6.0 or later is required for the management cluster.

- **Cisco ACI Virtual Edge:** You need one license for each instance of Cisco ACI Virtual Edge; you can have up to 32 instances on each Cisco ACI vPod.

Cisco ACI Virtual Edge is licensed in vPod mode, when it is part of Cisco ACI vPod. It is licensed as a VM that operates on each workload server. It provides data plane and policy enforcement functions. (Cisco ACI Virtual Edge is not licensed when installed on a physical leaf—when it functions outside of Cisco ACI vPod deployments.)



Note VMware ESXi 6.0 or later is required for the workload.

Cisco ACI vPod in a Multipod Environment

Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) can be used with multipod Cisco ACI fabrics. You must configure an interpod network (IPN) connection to a physical pod in the on-premises data center before you can add the Cisco ACI vPod.

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.

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

- *Cisco Application Centric Infrastructure Fundamentals*
- *Cisco Layer 3 Network Configuration Guide*
- *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*

See the *Cisco ACI Virtual Edge Release Notes* for information about Cisco Application Centric Infrastructure Virtual Edge for features not supported with multipod.

Local L3Out Configuration for Cisco ACI vPod

This section provides the information that you need—including prerequisites and procedures—to enable local Layer 3 outside network configuration (L3out) for Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod).

Prerequisites for Cisco ACI vPod L3Out Configuration

You must have performed the following tasks before you can improve network efficiency by configuring a Layer 3 outside (L3Out) connection for Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod).

- Have installed a Cisco Application Centric Infrastructure (ACI) fabric, including Cisco Application Policy Infrastructure Controller (APIC).

See the *Cisco APIC Getting Started Guide* and the *Cisco APIC Basic Configuration Guide* on [Cisco.com](https://www.cisco.com) for details.

- Installed Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) in a remote location.

See the *Cisco ACI Virtual Pod Installation Guide* and *Cisco ACI Virtual Pod Getting Started Guide* on [Cisco.com](https://www.cisco.com) for details.

- Configured the external router connectivity with Cisco ACI vPod to communicate with virtual spines for local L3Out configuration.

Cisco ACI GOLF is supported on CSR1000V, ASR1000, ASR9000 and Nexus 7000.

- Have read and followed the guidelines for configuring tenant Cisco ACI GOLF.

See the chapter "Routed Connectivity to External Networks" in the [Cisco APIC Layer 3 Networking Configuration Guide](#).

- Have read the chapter "Cisco ACI GOLF" in the [Cisco APIC Layer 3 Networking Configuration Guide](#).

The GOLF feature and its configuration are nearly identical to those of Cisco ACI vPod L3Out.

Workflow for Cisco ACI vPod L3Out Configuration

This section provides a high-level description of the tasks that are required to configure a Layer 3 outside network connection (L3Out) through the Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod).

- Fulfill all the prerequisites, which include meeting all hardware and software requirements.

See the section [Prerequisites for Cisco ACI vPod L3Out Configuration, on page 6](#) in this guide.

- Create a tenant, virtual routing and forwarding (VRF) instance for the tenant, and a bridge domain on a virtual leaf in the remote site.

See the section [Create a Tenant, VRF, and Bridge Domain, on page 7](#) in this guide.

- Create an application profile.

See the section [Creating an Application Profile Using the GUI, on page 8](#) in this guide.

- Create endpoint groups (EPGs) that will consume the L3Out.

See the section [Creating EPGs, on page 8](#) in this guide.

- Configure the L3Out for the created VRF.

See the section [Configuring an L3Out Using Cisco ACI GOLF, on page 10](#) in this guide.

Create a Tenant, VRF, and Bridge Domain

Complete the procedure in this section to create a tenant, virtual routing and forwarding (VRF) instance, and a bridge domain.

Procedure

- Step 1** On the menu bar, choose **Tenants > Add Tenant**.
- Step 2** In the **Create Tenant** dialog box, perform the following tasks:
- a) In the **Name** field, enter a name.
 - b) Click **Submit** to complete the tenant configuration.
- Step 3** In the **Navigation** pane, expand **Tenant-name > Networking**, and in the **Work** pane, drag the **VRF** icon to the canvas to open the **Create VRF** dialog box, and perform the following tasks:
- a) In the **Name** field, enter a name.
 - b) Click **Submit** to complete the VRF configuration.
- Step 4** In the **Networking** pane, drag the **BD** icon to the canvas while connecting it to the **VRF** icon. In the **Create Bridge Domain** dialog box that displays, perform the following tasks:
- a) In the **Name** field, enter a name.

- b) Click the **L3 Configurations** tab.
 - c) Expand **Subnets** to open the **Create Subnet** dialog box, enter the subnet mask in the **Gateway IP** field.
 - d) In the **Scope** area, check the **Advertised Externally** check box to ensure that the bridge domain can communicate with the Layer 3 outside network connection (L3Out).
 - e) Click **Submit** to complete bridge domain configuration.
-

Creating an Application Profile Using the GUI

Procedure

- Step 1** On the menu bar, choose **TENANTS**. In the **Navigation** pane, expand the tenant, right-click **Application Profiles**, and click **Create Application Profile**.
 - Step 2** In the **Create Application Profile** dialog box, in the **Name** field, add the application profile name (OnlineStore).
-

Creating EPGs

After you create an application profile, you create endpoint groups (EPGs). You can deploy EPGs on specific nodes or ports, or you can associate it to a specific virtual machine manager (VMM) domain.

Deploying an EPG on a Specific Node or Port

You can create an endpoint group (EPG) on a specific node or a specific port on a node.

Before you begin

You must have created the tenant where you deploy the EPG.

Procedure

- Step 1** Log in to the Cisco Application Policy Infrastructure Controller (APIC).
- Step 2** Choose **Tenants > tenant**.
- Step 3** In the left navigation pane, expand *tenant*, **Application Profiles**, and the *application profile*.
- Step 4** Right-click **Application EPGs** and choose **Create Application EPG**.
- Step 5** In the **Create Application EPG STEP 1 > Identity** dialog box, complete the following steps:
 - a) In the **Name** field, enter a name for the EPG.
 - b) From the **Bridge Domain** drop-down list, choose a bridge domain.
 - c) Check the **Statically Link with Leaves/Paths** check box.

This check box allows you to specify on which port you want to deploy the EPG.
 - d) Click **Next**.
 - e) From the **Path** drop-down list, choose the static path to the destination EPG.
- Step 6** In the **Create Application EPG STEP 2 > Leaves/Paths** dialog box, from the **Physical Domain** drop-down list, choose a physical domain.

Step 7

Complete one of the following sets of steps:

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

Step 8

Click **Update** and click **Finish**.

Step 9

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

Step 10

Complete one of the following actions:

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

Creating an EPG and Associating It with a VMM

The port the EPG uses must belong to one of the VM Managers (VMM) or physical domains associated with the EPG.

Procedure**Step 1**

On the menu bar, choose **Tenants** and the tenant where you want to create an EPG.

Step 2

In the navigation pane, expand the folder for the tenant, the **Application Profiles** folder, and the folder for the application profile.

Step 3

Right-click the **Application EPG** folder, and in the **Create Application EPG** dialog box, perform the following actions:

- a) In the **Name** field, add the EPG name (db).
- b) In the **Bridge Domain** field, choose the bridge domain from the drop-down list (bd1).
- c) Check the **Associate to VM Domain Profiles** check box. Click **Next**.
- d) In the **STEP 2 > Domains** area, expand **Associate VM Domain Profiles** and from the drop-down list, choose the desired VMM domain.
- e) From the **Deployment Immediacy** drop-down list, accept the default or choose when policies are deployed from Cisco APIC to the physical leaf switch.
- f) From the **Resolution Immediacy** drop-down list, choose when policies are deployed from the physical leaf switch to the virtual leaf.

If you have Cisco AVS, choose **Immediate** or **On Demand**; if you have Cisco ACI Virtual Edge or VMware VDS, choose **Immediate**, **On Demand**, or **Pre-provision**.

- g) (Optional) In the **Delimiter** field, enter one of the following symbols: |, ~, !, @, ^, +, or =.

If you do not enter a symbol, the system uses the default | delimiter in the VMware portgroup name.

- h) If you have Cisco ACI Virtual Edge or Cisco AVS, from the **Encap Mode** drop-down list, choose an encapsulation mode.

You can choose one of the following encapsulation modes:

- **VXLAN**: This overrides the domain's VLAN configuration, and the EPG uses VXLAN encapsulation. However, a fault is 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.

- i) If you have Cisco ACI Virtual Edge, from the **Switching Mode** drop-down list, choose **native** or **AVE**.

If you choose **native**, the EPG is switched through the VMware VDS; if you choose **AVE**, the EPG is switched through the Cisco ACI Virtual Edge. The default is **native**.

- j) Click **Update** and then click **Finish**.

Step 4 In the **Create Application Profile** dialog box, create two more EPGs. Create the three EPGs—db, app, and web—in the same bridge domain and data center.

Configuring an L3Out Using Cisco ACI GOLF


This section contains information about Cisco Application Centric Infrastructure (ACI) GOLF and how to use it to configure a local Layer 3 outside network connection (L3Out).

For general information about Cisco ACI GOLF, see the chapter "Cisco ACI GOLF" in the [Cisco APIC Layer 3 Networking Configuration Guide](#).

Cisco ACI Golf and Cisco ACI vPod

The Cisco Application Centric Infrastructure GOLF feature (also known as Layer 3 Ethernet VPN (EVPN) Services for Fabric WAN) enables much more efficient and scalable Cisco ACI fabric WAN connectivity. It uses the Border Gateway Protocol (BGP) EVPN protocol over IP for WAN routers that are connected to spine switches.

Figure 3: Cisco ACI GOLF Topology



All tenant WAN connections use a single session on the spine switches where the WAN routers are connected. This aggregation of tenant BGP sessions towards the Data Center Interconnect Gateway (DCIG) improves control plane scale by reducing the number of tenant BGP sessions and the amount of configuration required for all of them. The network is extended out using Layer 3 configured on the virtual spine switches. Transit routing with shared services using GOLF is not supported.

A Layer 3 external outside network (L3extOut) for GOLF connectivity for a spine switch is specified under the *infra* tenant, and includes the following:

- LNodeP (l3extInstP is not required within the L3Out in the *infra* tenant.)
- A provider label for the L3extOut for GOLF in the *infra* tenant.
- BGP protocol policies

All regular tenants use the above-defined connectivity. The L3extOut defined in regular tenants requires the following:

- An l3extInstP (EPG) with subnets and contracts. The scope of the subnet is used to control import/export route control and security policies. The bridge domain subnet must be set to advertise externally and it must be in the same VRF as the application EPG and the GOLF L3Out EPG.
- Communication between the application EPG and the GOLF L3Out EPG is governed by explicit contracts (not Contract Preferred Groups).
- An l3extConsLbl consumer label that must be matched with the same provider label of an L3Out for GOLF in the *infra* tenant. Label matching enables application EPGs in other tenants to consume the LNodeP external L3Out EPG.
- The BGP EVPN session in the matching provider L3extOut in the *infra* tenant advertises the tenant routes defined in this L3Out.

Guidelines and Limitations

Observe the following GOLF guidelines and limitations:



Note These guidelines are applicable to both physical pods and Cisco ACI Virtual Pod (vPod).

- GOLF routers must advertise at least one route to Cisco ACI in order to accept traffic. No tunnel is created between leaf switches or Cisco ACI Virtual Edge and the external routers until Cisco ACI receives a route from the external routers.
- All Cisco Nexus 9000 Series ACI-mode switches and all of the Cisco Nexus 9500 platform ACI-mode switch line cards and fabric modules support GOLF. With Cisco Application Policy Infrastructure Controller (APIC) , release 3.1(x) and later, this includes the N9K-C9364C switch.
- Cisco ACI vPod that correspond to Cisco APIC Release 4.2(3) and later support L3Out over GOLF.
- At this time, only a single GOLF provider policy can be deployed on spine switch interfaces for the whole fabric.
- Up to Cisco APIC release 2.0(2), GOLF is not supported with multipod. In release 2.0 (2) the two features are supported in the same fabric only over Cisco Nexus N9000K switches without “EX” on the end of the switch name; for example, N9K-93120TX. Since the 2.1(1) release, the two features can be deployed together over all the switches used in the multipod and EVPN topologies.
- When configuring GOLF on a spine switch, wait for the control plane to converge before configuring GOLF on another spine switch.
- In Cisco ACI vPod, at least one virtual spine from each Cisco ACI vPod must establish EVPN BGP peering with external GOLF router(s).

We recommend that you peer both virtual spines from each Cisco ACI vPod with external GOLF router(s).

- A spine switch can be added to multiple provider GOLF outside networks (GOLF L3Outs), but the provider labels have to be different for each GOLF L3Out. Also, in this case, the OSPF Area has to be different on each of the `L3extOuts` and use different loopback addresses.
- The BGP EVPN session in the matching provider `L3Out` in the `infra` tenant advertises the tenant routes defined in this `L3extOut`.
- When deploying three GOLF L3Outs, if only 1 has a provider/consumer label for GOLF, and 0/0 export aggregation, APIC will export all routes. This is the same as existing `L3extOut` on leaf switches for tenants.
- If there is direct peering between a spine switch and a data center interconnect (DCI) router, the transit routes from leaf switches to the ASR have the next hop as the PTEP of the leaf switch. In this case, define a static route on the ASR for the TEP range of that ACI pod. Also, if the DCI is dual-homed to the same pod, then the precedence (administrative distance) of the static route should be the same as the route received through the other link.
- The default `bgpPeerPfxPol` policy restricts routes to 20, 000. For ACI WAN Interconnect peers, increase this as needed.
- In a deployment scenario where there are two `L3extOuts` on one spine switch, and one of them has the provider label `prov1` and peers with the DCI 1, the second `L3extOut` peers with DCI 2 with provider label `prov2`. If the tenant VRF has a consumer label pointing to any 1 of the provider labels (either `prov1` or `prov2`), the tenant route will be sent out both DCI 1 and DCI 2.
- When aggregating GOLF OpFlex VRFs, the leaking of routes cannot occur in the ACI fabric or on the GOLF device between the GOLF OpFlex VRF and any other VRF in the system. An external device (not the GOLF router) must be used for the VRF leaking.
- A Cisco Cloud Service Router (CSR), or GOLF router, is not controlled by Cisco APIC. Before Cisco APIC Release 4.2(3), physical spine used to push only the VRF configurations to the GOLF router through the OpFlex channel. Because the OpFlex channel is not supported on the virtual spine (vSpine), a network administrator should manually configure the VRF configurations in the GOLF router.



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 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 maximum IP packet size of 9000 bytes in Cisco ACI, Cisco NX-OS, and Cisco IOS, but results in a maximum 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**.

Configure Cisco ACI GOLF for Cisco ACI vPod Using the GUI

The following steps describe how to configure infra GOLF services that any tenant network can consume.

Procedure

Step 1 On the menu bar, click **Tenants**, then click **infra** to select the infra tenant.

Step 2 In the **Navigation** pane, expand the **Networking** option and perform the following actions:

- Right-click **L3Outs** and click **Create L3Out** to open the **Create L3Out** wizard.
- In the **Identity** window, enter the necessary information in the **Name** and **L3 Domain** fields.

The **VRF** field is automatically populated because you configured a VRF earlier.

- In the **Use For:** field, select **Golf**.

The **Provider Label** and **Route Target** fields appear.

- In the **Provider Label** field, enter a provider label (for example, `golf`).
- In the **Route Target** field, choose whether to use automatic or explicit policy-governed BGP route target filtering policy:

- **Automatic:** Implements automatic BGP route-target filtering on VRFs associated with this routed outside configuration.
- **Explicit:** Implements route-target filtering through use of explicitly configured BGP route-target policies on VRFs associated with this routed outside configuration.

Note Explicit route target policies are configured in the **BGP Route Target Profiles** table on the **BGP Page** of the **Create VRF Wizard**. If you select the **Automatic** option in the **Route Target** field, configuring explicit route target policies in the **Create VRF Wizard** might cause BGP routing disruptions.

- Leave the remaining fields (BGP selected, and so on) as-is, and click **Next**.

Step 3 Enter the necessary information in the **Nodes and Interfaces** window of the **Create L3Out** wizard.

- Leave the **Use Defaults** check box checked.
- In the **Node ID** drop-down list, choose a spine switch node ID.
- In the **Router ID** field, enter the router ID.
- (Optional) You can configure another IP address for a loopback address, if necessary.

If you leave the **Loopback Address** field empty, the loopback address will be the same as in the neighboring **Router ID** field. Enter a different IP address for a loopback address if you don't want to use the router ID.

- e) Leave the **External Control Peering** field checked.
- f) Enter necessary additional information in the **Nodes and Interfaces** window.

The fields that are shown in this window vary, depending on the options that you select in the **Layer 3** and **Layer 2** areas.

- g) When you have entered the remaining additional information in the **Nodes and Interfaces** window, click **Next**.

Step 4

In the **Create Infra BGP Peer Connectivity Profile** window of the **Create L3Out** wizard, enter the necessary information.

- a) In the **BGP Loopback Policies** and **BGP Interface Policies** areas, enter the following information:

- **Peer Address:** Enter the peer IP address

- **EBGP Multihop TTL:** Enter the connection Time To Live (TTL). The range is 1–255 hops; if zero, no TTL is specified. The default is zero.

- **Source Interface Type:** Choose **L3Out Loopback as Source** or **Routable Loopback as Source**.

Choose **Routable Loopback as Source** if both of the following are true: You will deploy on-premises GOLF (GOLF on physical spines) with routable IP pools for L3Out BGP peering with GOLF routers, and you will use the same GOLF routers for BGP peering with the Cisco Application Centric Infrastructure (ACI) Virtual Pod (vPod) virtual spines (vSpines). You also must choose **Routable Loopback as Source** for the Cisco ACI vPod.

If you choose **Routable Loopback as Source** an additional field, **Data Plane Address** appears. This address will be configured on GOLF router's nve interface. The GOLF router also will use it to send WAN traffic towards the Cisco ACI fabric (physical pod and Cisco ACI vPod). The physical spines will use the same address to send WAN traffic towards the GOLF router and by Cisco ACI Virtual Edge to send traffic towards the GOLF router.

- **Data Plane Address:** Enter the data plane address.

By specifying the **Data Plane Address** in the Cisco Application Policy Infrastructure Controller (APIC) GUI, you tell the Cisco ACI that this is the address in data path already configured on the GOLF router to which we need to send WAN traffic.

- **Remote ASN:** Enter a number that uniquely identifies the neighbor autonomous system. The Autonomous System Number can be in 4 byte as plain format 1–4294967295.

If you need to use BGP passwords, enter the information in the **Password** field.

Note Cisco ACI does not support asdot or asdot+ format autonomous system numbers.

- b) In the **OSPF** area, choose the default OSPF policy, a previously created OSPF policy, or **Create OSPF Interface Policy**.
- c) Click **Next**.

The **External EPG** window appears.

Step 5

Enter the necessary information in the **External EPG** window of the **Create L3Out** wizard.

- a) In the **Name** field, enter a name for the external network.
- b) In the **Provided Contract** field, enter the name of a provided contract.
- c) In the **Consumed Contract** field, enter the name of a consumed contract.
- d) In the **Allow All Subnet** field, uncheck if you don't want to advertise all the transit routes out of this L3Out connection.

The Subnets area appears if you uncheck this box. Specify the desired subnets and controls as described in the following steps.

- e) Click **Finish** to complete the necessary configurations in the **Create L3Out** wizard.

Step 6 In the **Navigation** pane for any tenant, expand the *tenant_name* > **Networking** and perform the following actions:

- a) Right-click **L3Outs** and click **Create L3Out** to open the wizard.
- b) In the **Create L3Out** dialog box, enter the necessary information in the **Name**, **VRF**, and **L3 Domain** fields.
- c) Check the box next to the **Use for GOLF** field.
- d) In the **Label** field, choose **Consumer**.
- e) Assign a **Consumer Label**.

Use the same name as you did when you created the provider label in Step 2d for the infra tenant.

- f) Click **Next**.

Step 7 In the **External EPG** dialog box, complete the following steps:

- a) In the **Name** field, enter a name for the external EPG.
- b) From the **Provided Contract** drop-down list, choose or create a provided contract.
- c) From the **Consumed Contract** drop-down list, choose or create a consumed contract.
- d) In the **Default EPG for all external networks**, leave the check box checked if you want to use the EPG for all external networks.

What to do next

Configure EVPN type 3 and type 5 route leaking from Cisco Application Centric Infrastructure (ACI). Follow the procedure in the section ""Distributing BGP EVPN Type-2 Host Routes to a DCIG" in the "Cisco ACI GOLF" chapter of the *Cisco APIC Layer 3 Networking Configuration Guide*.

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: www.cisco.com/go/trademarks. 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)

© 2019 Cisco Systems, Inc. All rights reserved.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

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