

Cisco Application Centric Infrastructure and VMware Integration

What You Will Learn

The need for quick deployment of applications is challenging IT infrastructure in new ways. Infrastructure must become application aware and more agile to support dynamic application instantiation and removal. Cisco® Application Centric Infrastructure (ACI) is a highly flexible, open, programmable environment that can be transparently integrated into VMware virtual environments. Cisco ACI integration with VMware helps deliver simplicity without compromising infrastructure scale, performance, responsiveness, security, or end-to-end visibility.

Cisco ACI simplifies and accelerates the entire application deployment lifecycle for the next-generation data center and cloud deployments. Cisco ACI takes a systems-based approach and provides tight integration between physical infrastructure ready to support Cisco ACI and VMware virtual elements. Cisco ACI enables customers to use a common policy-based operating model across their physical, virtual, and cloud-based environments with penalty-free overlays.

The Cisco Application Policy Infrastructure Controller (APIC) is the main architectural component of Cisco ACI integration with the VMware virtual environment. It provides a unified point of automation and management for Cisco ACI fabric, policy enforcement, and health monitoring for both physical and virtual environments.

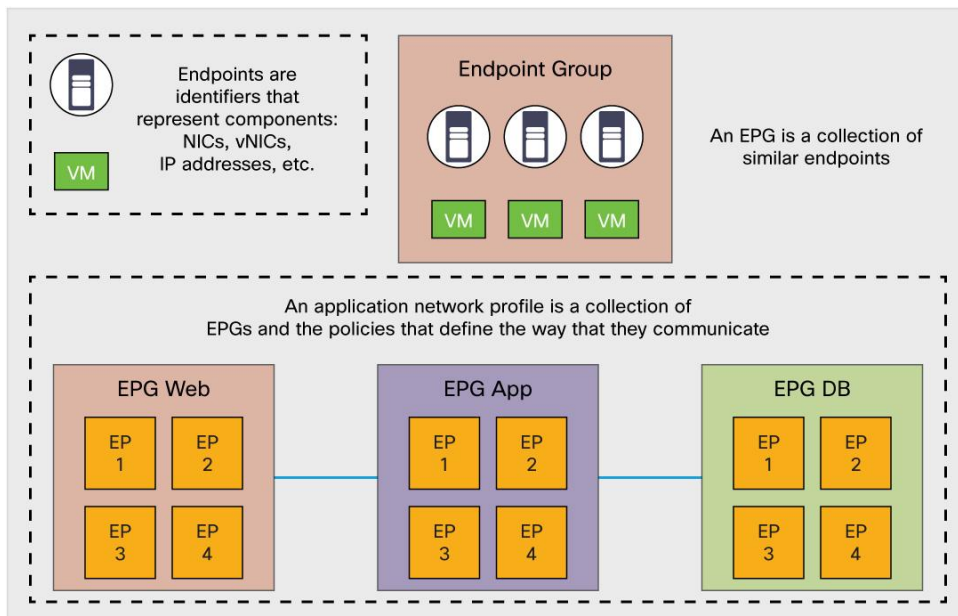
Fundamentals of Cisco ACI

Cisco ACI abstracts policy and connectivity from the underlying fundamental network constructs of VLANs, IP addresses, access control lists (ACLs), and quality-of-service (QoS) policies. It then can be used to describe application network connectivity more abstractly, in terms of endpoint groups, providers and consumers, service-level agreements (SLAs), etc., to make this connectivity relevant for the end user of the fabric.

Cisco ACI provides a secure multitenant solution, allowing the network infrastructure administration and data flows to be segregated. Tenants can be divided into customers, business units, groups, etc. Tenants also can be subdivided into Layer 3 constructs, known as Virtual Routing and Forwarding (VRF) instances. Contexts provide an additional way to separate the organizational and forwarding requirements of a tenant. Within each VRF instance, a bridge domain is created. A bridge domain is a Layer 2 namespace in which the various subnets are defined. All the subnets and default gateways are assigned within the bridge domain. By using separate forwarding instances, IP addressing can be duplicated in separate contexts for multi-tenancy.

Within the context, a series of objects, called endpoints and endpoint groups (EPGs), define an application (Figure 1). Endpoints are devices (physical or virtual) that connect to the fabric and use it to interface with the network infrastructure. These endpoints can be computing, storage, and network services and security devices that attach to the fabric. At first customer shipment (FCS), Cisco ACI will support endpoints classified as network interface cards (NICs) or virtual NICs (vNIC) and their associated VLANs or Virtual Extensible LANs (VXLANs). In the future, endpoints will be extended to include IP addresses, MAC addresses, Domain Name System (DNS) names, virtual machine attributes, IEEE 802.1x identity, and other common attributes.

Figure 1. Endpoints and Endpoint Groups



An EPG is a collection of endpoints with the same types of attributes and identical network behavior (connectivity, security, QoS requirements, etc.).

Here are some examples of EPGs:

- EPG defined by traditional network VLANs; all endpoints connected to a given VLAN are placed in an EPG
- EPG defined by VXLANs; same as for VLANs except using VXLAN
- EPG defined by security zone
- EPG defined by application tier (web, application, or database)
- EPG mapped to a VMware ESXi port group


The policies used to describe the communication, services insertion, and QoS and SLAs embedded in EPGs are referred to as contracts. A contract is a set of policy requirements that describe how EPGs can communicate with each other and with the outside world. The Cisco ACI contract defines a filter, which includes a Layer 4 inbound/outbound ports and an associated action that dictates whether the traffic is permitted, denied, logged, marked, redirected, or copied. The default behavior of Cisco ACI security follows an allowed list model, which denies all traffic, this is a configurable option. In the default mode administrators must explicitly allow communication between EPGs.

For each tenant, EPGs and policies are summarized in an application network profiles (ANP). These ANPs are the logical representation of the application infrastructure requirements. When the application is ready to be provisioned, Cisco APIC can push the ANP and provision the entire stateless Cisco ACI fabric and L4-7 service devices instantaneously.

Cisco ACI Integration with VMware vCenter

A main benefit of Cisco ACI is the capability to be hypervisor independent. Cisco ACI is hypervisor independent to the tenant traffic, which can be tagged with IEEE 802.1q (VLAN), VXLAN, or Network Virtualization Using Generic Routing Encapsulation (NVGRE). Traffic forwarding is not limited to or constrained within the encapsulation type or encapsulation overlay network (Figure 2).

Figure 2. Hypervisor-Independent Traffic in Cisco ACI



Cisco ACI uses an extended VXLAN encapsulation frame format in the fabric. All tenant traffic in the fabric is tagged at the first-hop leaf ingress port with an extended VXLAN header, which identifies the policy attributes of the application endpoint in the fabric. The VXLAN header carries the VXLAN network identifier (VNID) along with the EPG policy. As a result, the policy attributes are carried in every packet. As workloads move within the virtual environment, the policies attached to the workloads are enforced transparently and consistently within the infrastructure. When the packet leaves the fabric, the VXLAN header is deencapsulated, and the packet is encapsulated with any tag of the tenant's choice: VLAN, VXLAN, or NVGRE based on the destination EPGs encapsulation.

In Cisco ACI, virtual machine management (VMM) refers to a hypervisor management system that has control over virtual machines. A Cisco ACI fabric can have multiple VMM domains across hypervisors from different vendors. The VMM domain for VMware is VMware vSphere vCenter. Each VMM domain has a local significant association with a VLAN or VXLAN and contains its own mobility domain, which means that mobility is restricted to that VMM domain. Across the various VMM domains, the VLAN or VXLAN namespace can be reused. At FCS, the VLAN ID has local significance for the leaf node. In the future, the VLAN ID will have local port significance. As a best practice, do not overlap VLAN namespaces for a given leaf node. Designate a few leaf nodes and put them in a VMM domain so they can own the VLAN namespace. Using VMM domains, a customer can scale beyond the usual restriction of 4000 VLANs. In a large-scale cloud environment, 4000 VLANs or identifiers are not enough to uniquely identify each tenant, segment or network. Each VMM domain can scale up to 4000 VLANs. As a result, organizations can span and scale horizontally using VMM domains and achieve up to 64,000 logical network segments. The only restriction, as mentioned earlier, is that VM mobility is limited to the VMM domain.

The location and reachability of VMware virtual endpoints can be learned in several ways:

- Control-plane learning
 - Out-of-band (OOB) handshake: The OOB control protocol is used to communicate between the VMM domain and Cisco APIC. VMware vCenter knows where the hosts are because it places the virtual host. That information is used to understand where the host resides.
 - Link Layer Discovery Protocol (LLDP) and Cisco Discovery Protocol: These protocols are used to resolve the virtual host ID to the attached port on a leaf node.
- Data-plane learning
 - Distributed switch learning: Distributed switch learning is used to distribute and configure policy in hardware.

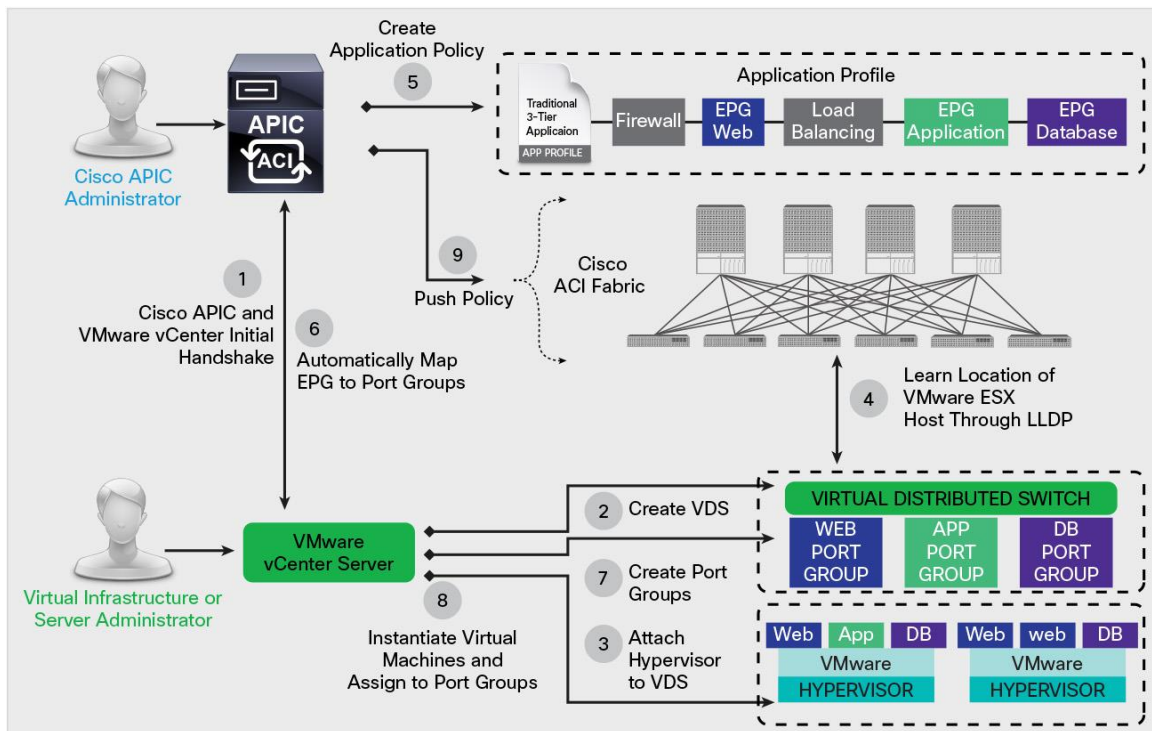
When a virtual endpoint is discovered, the policy is pushed to and configured on the leaf nodes based on the resolution immediacy policy and instrumentation immediacy policy, respectively. Both policies have immediate and on-demand (default) options that are defined when the VMM domain is associated on Cisco APIC. The on-demand option conserves resources and uses the reserved space in the policy content addressable memory (CAM) when necessary.

- Resolution immediacy policy: When the immediate option is used, all policies (VLAN, VXLAN, and NVGRE) bindings, contracts, and filters are immediately pushed to the leaf node when the hypervisor physical NIC (pNIC) is attached. With the on-demand option, policies are pushed to the leaf node when the pNIC and vNIC are attached to the port group (EPG).
- Deployment immediacy policy: The option is defined when the policies are configured in the hardware. If the immediate option is selected, the policies are programmed in the policy CAM after reachability information is received by Cisco APIC. The on-demand option configures policies in the hardware policy CAM only when reachability is learned through the data path.

Integrating Cisco ACI with VMware

Cisco APIC integrates with the VMware vCenter instances to transparently extend the Cisco ACI policy framework to VMware vSphere workloads. Cisco APIC uses ANPs to represent the Cisco ACI policy. Cisco APIC creates a virtual distributed switch (VDS) in VMware vCenter to create the virtual network. From that point onward, Cisco APIC manages all application infrastructure components. The network administrator creates EPGs and pushes them to VMware vCenter as port groups on the DVS. Server administrators can then associate the virtual machines and provision them accordingly (Figure 3).

Figure 3. Integrating Cisco ACI with VMware



For More Information

- Cisco ACI: <http://www.cisco.com/go/aci>
- VMware: <http://www.vmware.com>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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.

Recycling symbol: 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. (1110R)