

Transform Your Business and Protect Your Cisco Nexus Investment While Adopting Cisco Application Centric Infrastructure

What You Will Learn

The new Cisco® Application Centric Infrastructure (ACI) fabric architecture brings a revolutionary policy-based approach to data center networks with the introduction of the Cisco Nexus® 9000 Series Switches and the Cisco Application Policy Infrastructure Controller (APIC). Enhancements introduced in the Cisco Nexus 9000 Series facilitate programmability of the fabric according to application-centric policies, as well as optimization of the overlay networks with native hardware support.

In addition to delivering the Cisco ACI policy model, Cisco provides interoperability and investment protection across the Cisco Nexus Family switching portfolio by focusing on two scenarios:

- Scenario 1: Enabling virtual or physical servers on existing Cisco Nexus networks to participate in the Cisco ACI fabric using Cisco APIC to provision workload (or endpoint group) policies and enable ACI forwarding mechanisms across both ACI (Nexus 9000-based) and existing Nexus fabrics (Nexus 3000-7000).
- Scenario 2: Use the Cisco Nexus 7000 Series Switches as data center interconnect (DCI) devices between the Cisco ACI fabric and an existing Cisco Nexus fabric (local or remote data center).

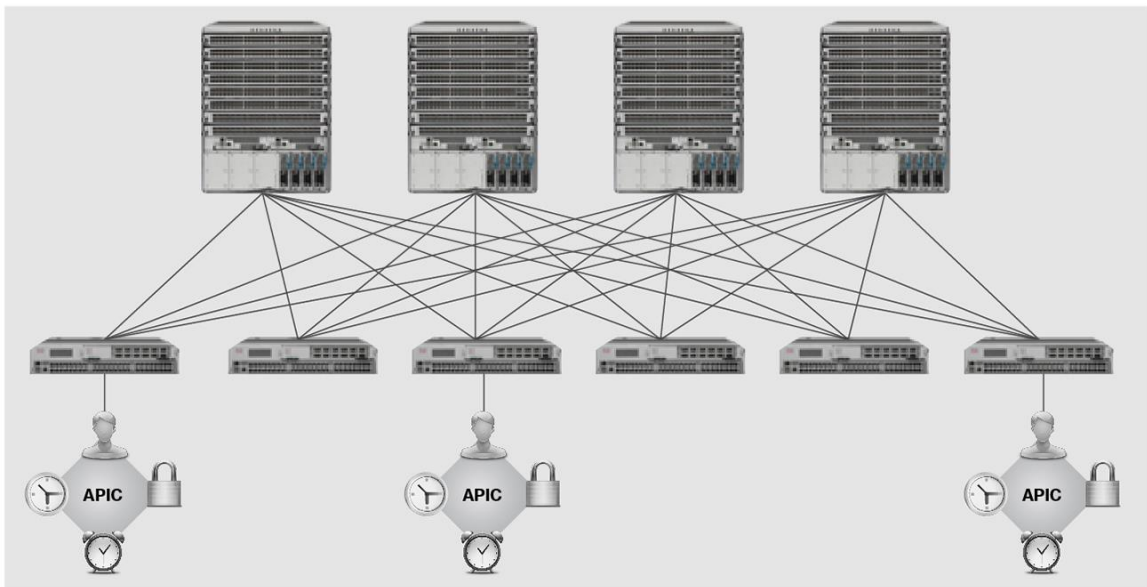
Scenario 1: The first scenario benefits customers that want applications on their current Cisco Nexus networks to participate in the Cisco ACI fabric policy. In this case, Cisco APIC manages policy for virtual or physical servers in the existing network. For virtual servers, they connect to an application-centric virtual switch (AVS) and enabled for Cisco ACI, in the existing Cisco Nexus network. AVS acts as a virtual leaf for the Cisco ACI spine-and-leaf fabric, and as an edge switch that is Cisco ACI aware, it can forward traffic according to Cisco ACI policy rules and apply Layer 4 through 7 services managed by Cisco ACI. For physical servers, a Cisco Nexus 9300 platform switch acts as an access-layer switch to the existing overlay network. Compared to other software-defined networking (SDN) overlay solutions, this solution provides a common infrastructure for physical and virtual workloads, along with a more advanced application-centric policy model.

Scenario 2: The second scenario provides benefits to customers by allowing the Cisco ACI fabric to effectively connect to the Cisco Nexus fabric at remote data centers over the WAN, in addition to connecting to a local Cisco Nexus fabric. In a typical DCI scenario, a Cisco Nexus 7000 Series Switch (or Cisco Nexus 7700 platform switch) would interconnect two or more remote data centers across a WAN. This same capability is supported between the Cisco ACI fabric and a Cisco Nexus fabric (local or remote). In either case, the Cisco Nexus 7000 Series or 7700 platform switch is functioning as a peripheral router to the Cisco ACI fabric, using OpFlex as an interface to automate provisioning and exchange tenant information.

Overview of Cisco ACI Architecture and Integration Requirements for Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series Switches

Cisco ACI (Figure 1), Cisco's vision for the next-generation programmable and automated data center, was introduced in November 2013. Cisco ACI is based on some of the foundational principles of SDN, primarily to increase programmability of the fabric architecture for greater automation and on-demand workload deployment and optimization. In addition, Cisco ACI includes important enhancements that allow dramatic improvements over prior SDN approaches.

Figure 1. The Cisco ACI Fabric Is a Spine-and-Leaf Architecture of Cisco Nexus 9000 Series Switches Supported by Multiple Redundant Cisco APIC Controllers for Scalability and Fault Tolerance



Cisco ACI introduced a revolutionary application-centric language and programmability model that uniquely aligns with business and application objectives, simplifying policy implementations and better aligning IT with business needs. Cisco ACI also extends the network-specific approaches of SDN to the entire data center infrastructure, incorporating servers, storage, application services, and security under one policy model, unifying siloed IT teams and providing the capability to automate all network aspects of application deployment and the application lifecycle.

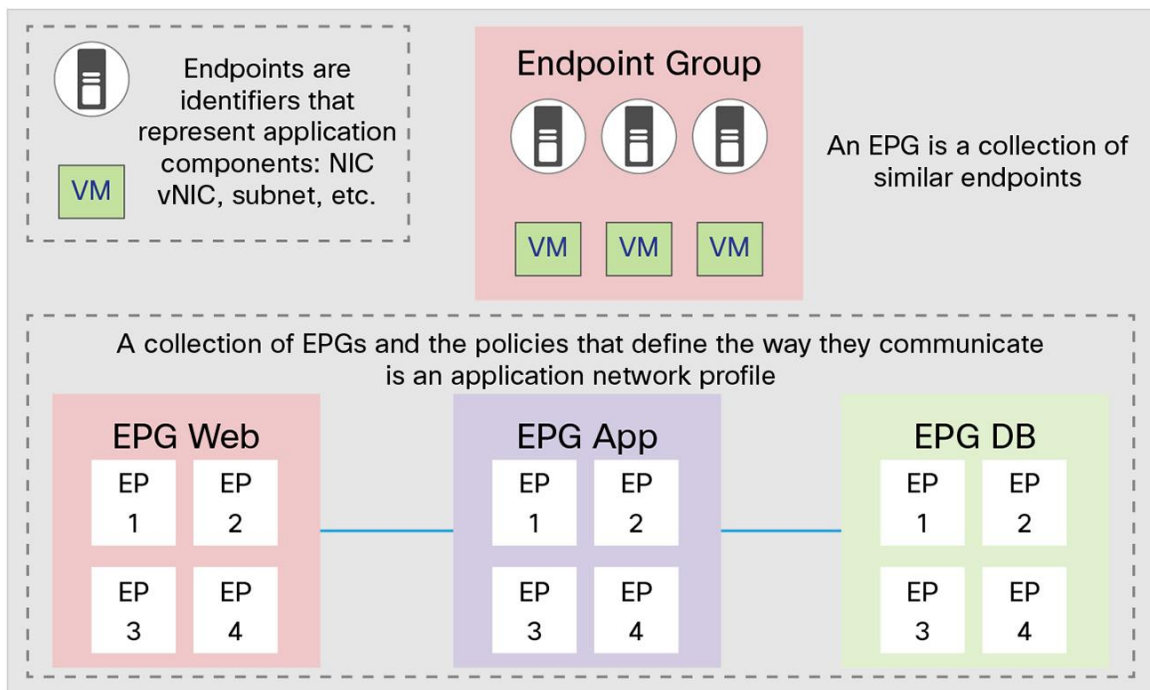
To achieve this revolutionary new fabric architecture, critical features were built into the Cisco Nexus 9000 Series Switches, including hardware-optimized overlays that simplify the network topology and segment tenant and application networks. These Virtual Extensible LAN (VXLAN) overlays are optimized through a combination of merchant and proprietary application-specific integrated circuits (ASICs) in the Cisco Nexus 9000 Series.

The Cisco ACI infrastructure controller, Cisco APIC, maintains the application policies and device templates that configure, provision, and manage fabric nodes, including application services, security devices, and the virtual infrastructure to support automated application deployment and optimization. Network nodes and devices managed under the Cisco ACI fabric must support a southbound communication mechanism from Cisco APIC, which in many cases will be the OpFlex protocol. OpFlex is specific to the Cisco ACI application policy model and is being implemented across multivendor network devices, application delivery controllers (ADNs), firewalls, virtual switches, and more as the Cisco ACI ecosystem evolves. OpFlex is also being implemented across other SDN

controllers that will support the Cisco ACI policy model, such as OpenDaylight, through open source contributions. OpFlex has been contributed to the IETF as a standard southbound controller protocol, and reference implementations are being distributed through leading open source vendors.

This Cisco ACI architecture is designed to apply policy rules to a fundamental application element, an endpoint, or a class of similar endpoints, called an endpoint group (EPG). Endpoints are essentially workloads, either physical or virtual, with multiple EPGs making up the various tiers of a multitier application, for example. Application networks are defined as connected EPGs with policy-based contracts or requirements between each pair of endpoints, such as security rules and quality of service (QoS) policy (Figure 2). The business-relevant policy language applied to application EPGs, including definitions of application networks, with the configuration and control of network devices, is what makes the Cisco ACI fabric application centric.

Figure 2. Endpoint Groups



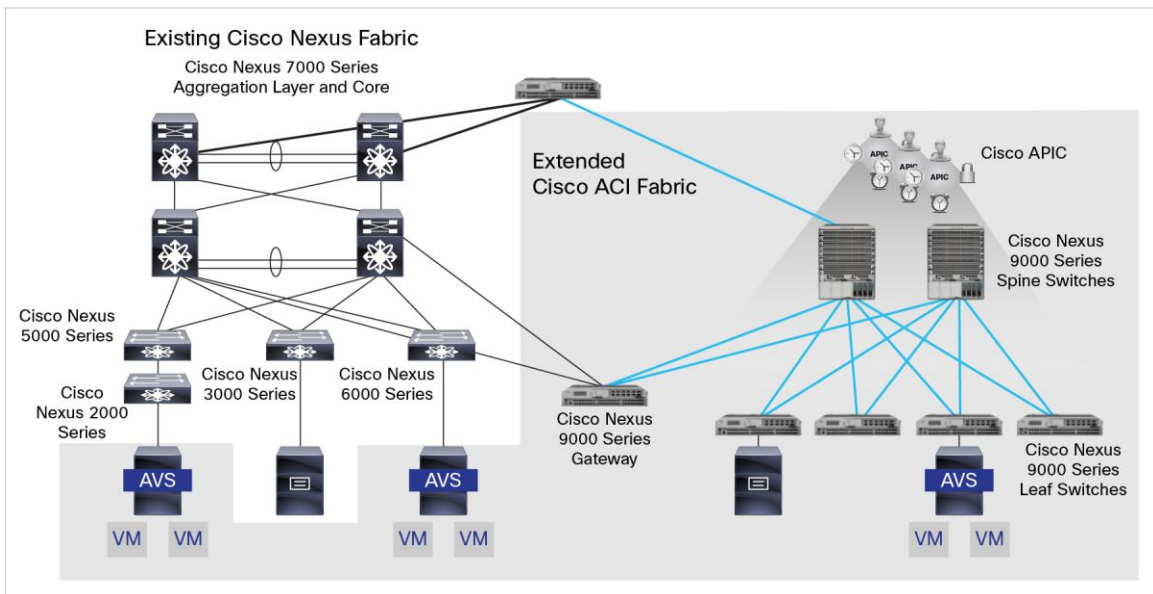
As organizations adopt and migrate to the Cisco ACI architecture, a critical requirement is that they be able to use their existing data center switching investments in Cisco Nexus hardware that predate Cisco ACI, specifically the Cisco Nexus 2000, 3000, 5000, 6000 and 7000 Series Switches. In general, they want to manage their data center applications as EPGs managed under Cisco APIC, but connected through networks other than Cisco Nexus 9000 Series networks. Additionally, as customers build out Cisco ACI fabrics, they want to preserve their investment in existing Cisco models by using them in their new Cisco ACI network, under a common orchestration platform and a common policy wherever possible.

The challenges in achieving these objectives include lack of Cisco APIC integration in the existing Cisco Nexus models and limited VXLAN overlay support in those devices, required for Cisco ACI fabric topologies. Subsequent sections discuss how to overcome these challenges.

Scenario 1: Enabling EPG Policies for Applications Attached to Existing Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series Networks

Most customers deploying Cisco ACI fabrics today have existing Cisco Nexus networks, usually three-tier (edge, aggregation, and core), based on some combination of Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series devices. These existing networks support a combination of physical and virtual workloads, with the virtual workloads attached to fabric extenders or virtual switches, usually the Cisco Nexus 1000V virtual switch. In addition to building out new pods based on Cisco ACI, many of these customers want to use the existing Cisco Nexus architecture to apply Cisco ACI policies to EPGs attached to the Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series fabric, with little recabling, reconfiguration, etc. (Figure 3).

Figure 3. Extending Cisco APIC Policy Control Across Servers and Endpoints in a Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series Network



To meet this requirement, Cisco extends Cisco ACI policy support to application workloads across the entire existing network through a Cisco Nexus 1000V AVS enabled for Cisco ACI. In this scenario, all data traffic flows through the infrastructure that existed before Cisco ACI was added; the traffic flows to a Cisco Nexus 9000 Series network switch managed by Cisco APIC (shown as the Cisco Nexus 9000 Series gateway in Figure 3). The Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series network can be any number of layers (such as a two- or three-tier network), with AVS switches running on the servers to connect to the EPG workloads. Initially, the tunnel through the Cisco Nexus network in place before Cisco ACI was installed must be a Layer 2 connection to the Cisco Nexus 9000 Series portion of the network, but future enhancements will allow routed networks in the overlay to the Cisco ACI fabric (target: Q4CY14).

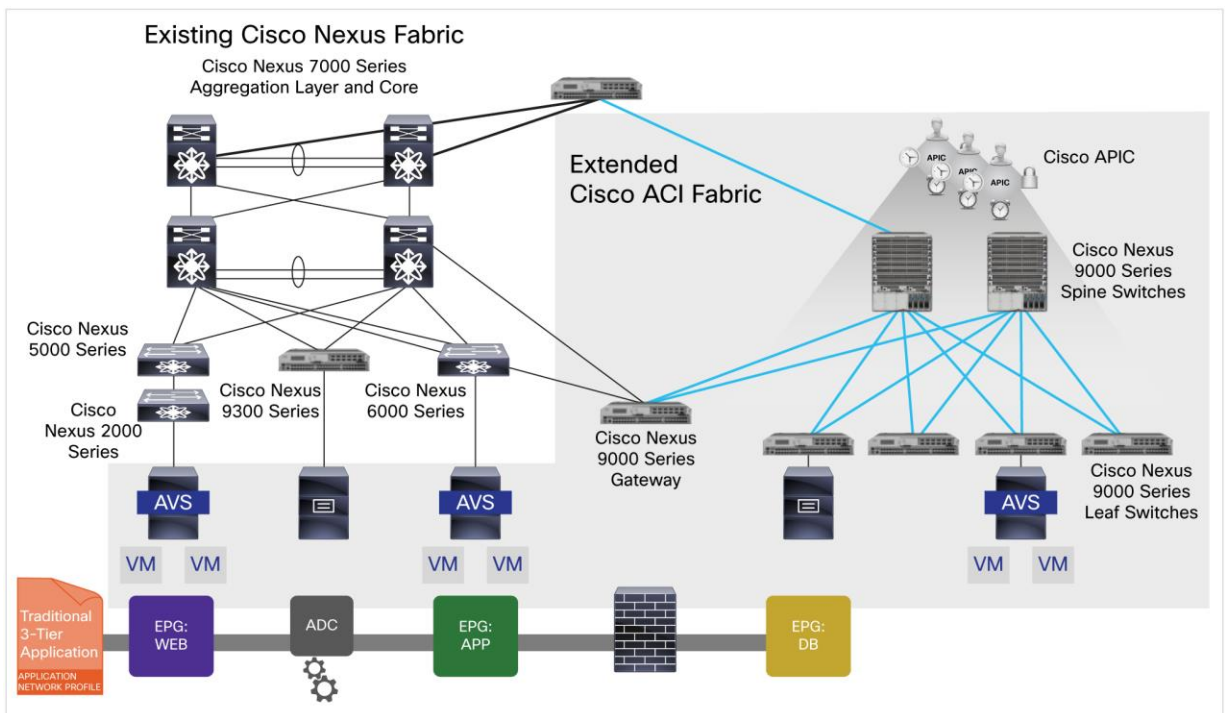
The Cisco ACI Directory/Proxy Service shown at the top of Figure 3 is required to map VM addresses to endpoint IP addresses. As workloads migrate throughout the fabric, this address mapping which tracks to location of all workloads is required to avoid multicast in the VXLAN fabric.

EPG-to-EPG workload connectivity within this network on separate servers could be handled normally, for example, through a VXLAN tunnel between servers, without reaching the Cisco Nexus 9000 Series gateway to communicate with workloads on the Cisco ACI network. In this deployment scenario, physical workloads and virtual workloads are handled consistently, with physical servers sending traffic directly to the Cisco Nexus 9000 Series gateway. As an alternative, the Cisco Nexus 9000 Series gateway can be inserted as a top-of-rack (leaf) switch for physical workloads in the existing topology and communicate with the Cisco Nexus 9000 Series and Cisco ACI spine through the existing Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series fabric as well (shown in Figure 4 with the addition of the Cisco Nexus 9300 platform on the left and the incorporation of the physical server into the Cisco ACI fabric).

All workloads can then be managed and provisioned as an EPG, with consistent policies across all servers and segments of the network, under a single policy controller and orchestration model. Cisco APIC can also support cloud orchestration software, such as OpenStack and Cisco UCS® Director, through northbound interfaces on the controller.

Figure 4 shows the three-tiers of a web application spread across the existing LAN environment and the new Cisco ACI pod. The web and application tiers are segregated across pods, using VLANs and subnets, and the database tier sits on the new Cisco ACI fabric. Policy enforcement will need to be provided for those connections, through additional Layer 4 through 7 application and security services. These services can be located in either the preexisting network or the new Cisco ACI fabric.

Figure 4. A Three-Tier Web Application Can Be Deployed Across Existing and New Server Pods, Supported by a Mix of Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series Switches and Cisco Nexus 9000 Series Devices, with Cisco APIC Policies Applied to All EPGs and Layer 4 Through 7 Services; Although Not Required in this Application Scenario, the Physical Server on the Left Is Also Now Incorporated into the ACI Policy Model with the Insertion of the Cisco Nexus 9300 Platform Switch in the Access Layer of the Network



The Cisco ACI policy model provides the capability to group EPG objects based on the service they deliver and the policy and connectivity required. This feature alleviates constraints inflicted by today's networks, which require forwarding constructs such as VLANs and subnets to be used for this purpose. In this model, groups of objects are constructed, and connectivity and policy enforcement is built in. The connection policy defined and managed in the three-tier web application in Figure 4 would include the ADC and firewall policies required between each tier of the application.

Layer 4 through 7 application services and security devices can be automatically provisioned and configured to support EPG workloads as in any Cisco ACI fabric. Virtual services, such as a virtual ADC or firewall, would likely be integrated through AVS. Physical application services can also be deployed and provisioned on the application network through VLAN stitching, and in the future through the proposed Network Services Header (NSH) standard based on Cisco vPath technology.

To further improve investment protection of existing ADC and firewall appliances, existing devices can be used, whether they are deployed on the existing or new Cisco Nexus 9000 Series fabric. Cisco ACI supports an extensive ecosystem of Layer 4 through 7 and security partners, including A10 Networks, Citrix, F5 Networks, Embrane, and Cisco Adaptive Security Appliances (ASA).

Continuing with the preceding example, Layer 4 through 7 rules for the existing tiers would already be in place for the web and application tiers on those devices, because they were in place on the existing infrastructure. Therefore, simply redirecting traffic to those devices is sufficient. For the new database tier on the Cisco ACI fabric, new rules may need to be configured. New rules for the database tier can be automatically pushed to the given Layer 4 through 7 devices through the use of Cisco ACI device packages from Cisco or the device vendor, as well as through custom scripts. Device packages are template definitions of functions for each device within Cisco APIC, allowing it to configure the desired policies and capabilities of each service for each EPG. This approach allows a single centralized point of policy management and deployment.

Because policy configuration is the most complex and tedious change in a network as applications are added, removed, expanded, or moved, automation of policy changes is critical. Existing VLANs, subnets, etc. have already been configured and do not require frequent modification. It is the individual device policies applied to each application that change frequently. Thus, device packages and policy definitions within Cisco APIC are capable of fully automating policy changes, without the need to reconfigure the existing network topology or equipment.

In the scenario discussed here, the existing network is essentially a transit network, with overlays running between the AVS or Cisco Nexus 9000 Series leaf nodes and the policy enforcement points in the Cisco ACI fabric. Cisco customers can use an existing network with little change, other than connection of new Cisco Nexus 9000 Series devices as needed to support new servers and workloads, thus protecting the investment in the existing Cisco Nexus network. AVS switches, where required, are easy and cost effective to deploy or upgrade to a version that supports Cisco ACI.

As a result, customers can transform their businesses and greatly accelerate their IT processes by aligning their business and application policies with their data center infrastructure, while still relying on whatever existing network they were using before.

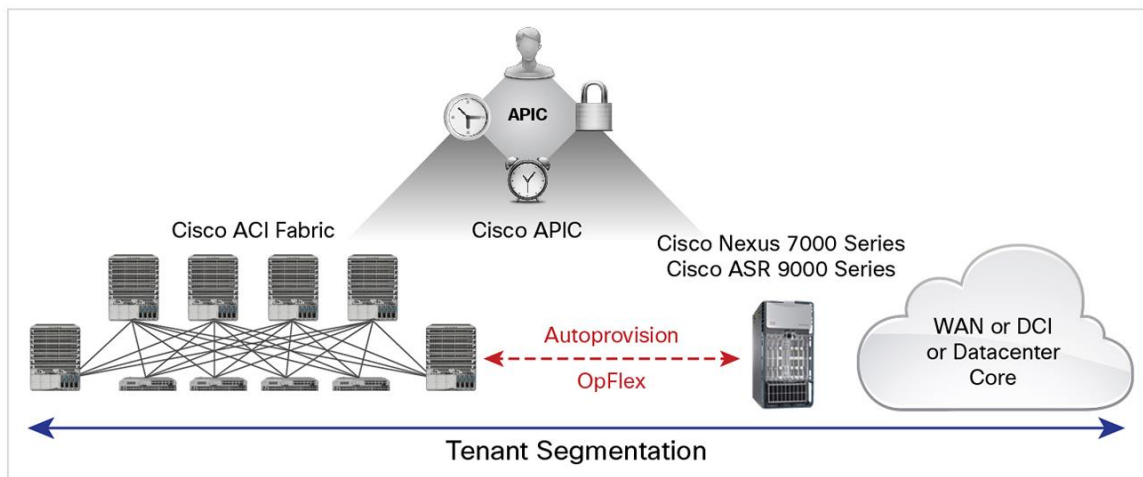
Scenario 2: Integrating Cisco Nexus 7000 Series Switches or Cisco ASR 9000 Series Routers as DCI Routers or Gateways Through OpFlex

A more specialized deployment scenario uses existing Cisco Nexus 7000 Series Switches as DCI routers between a Cisco ACI fabric and either a remote or local WAN or existing Cisco Nexus pod. In this capacity, the Cisco Nexus 7000 Series provides accelerated, highly secure data replication, server clustering, and workload mobility between data centers. The Cisco Nexus 7000 Series uniquely provides many of the DCI protocols required for this connectivity between data centers, including Multiprotocol Label Switching (MPLS) VPNs, Overlay Transport Virtualization (OTV), and Location/ID Separation Protocol (LISP).

Cisco Nexus 7000 Series Switches support the OpFlex protocol for integration with Cisco APIC specifically for the purpose of DCI integration as the gateway to remote data centers. The Cisco Nexus 7000 Series DCI node is incorporated into the Cisco ACI fabric as a peripheral device, outside the spine-and-leaf topology.

In this configuration (shown in Figure 5), the DCI router can be thought of as having two sides: a Cisco ACI fabric-facing side that peers with a border leaf over OpFlex, and a WAN-facing side that can be configured manually or through a separate network management system or controller. Cisco ASR 9000 Series Aggregation Services Routers and Cisco Nexus 7000 Series Switches both support OpFlex in this DCI scenario.

Figure 5. The Cisco Nexus 7000 Series Switch or Cisco ASR 9000 Series Router Can Be Integrated as a DCI Gateway Node into the Cisco ACI Fabric and Communicate with the Cisco APIC Through the OpFlex Protocol



Additional Cisco ACI awareness is built into the Cisco Nexus 7000 Series and Cisco ASR 9000 Series through OpFlex on top of the traditional DCI capabilities. The DCI router communicates directly with the border leaf to exchange tenant attributes to provision the Cisco ACI interfaces. The goal of this integration is to focus on automation of frequently changing per-tenant policies and configurations. Today, the Cisco ACI information model does not support a complete end-to-end WAN configuration. For example, parameters that may be exchanged between a border leaf and DCI device include a tenant ID, Virtual Routing and Forwarding (VRF) ID, Autonomous System Number (ASN) ID, and DCI IP address.

Because the hand-off between the border leaf and DCI will handle a very large number of tenants, it is important that the control plane and data plane be optimized for such a task. Therefore, Cisco intends to introduce optimizations for Border Gateway Protocol (BGP) that will allow a single BGP adjacency to carry the information for all tenants that exist in the fabric. These BGP scale enhancements would be coupled with a VXLAN control plane capable of extending the segmentation policies for a very large number of tenants.

Conclusion

The revolutionary approach and capabilities of Cisco ACI have accelerated change and innovation across the data center network. Although Cisco ACI can accelerate ROI and is designed to increase operation efficiency and reduce application delivery costs, organizations are still looking for ways to use their existing investments in Cisco Nexus hardware as they migrate to the new Cisco ACI architecture.

The design of the Cisco ACI fabric allows incorporation of existing Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series networks under a common cloud orchestration platform. It also allows management of application workloads attached to the network that existed before Cisco ACI was added; these workloads are managed as Cisco APIC endpoint groups. This capability preserves existing investment in Cisco Nexus 2000, 3000, 5000, 6000, and 7000 Series Switches, which can be used in the new fabric model.

Customers now can choose to continue to invest in Cisco Nexus hardware that preceded Cisco ACI, using current processes and orchestration models, or to start investing immediately in Cisco ACI fabric appliances. Either approach will help ensure optimal utilization of existing investments now and in the future.

For More Information

<http://cisco.com/go/aci>



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.

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)