

Integrate Cisco Application Centric Infrastructure with Existing Networks

What You Will Learn

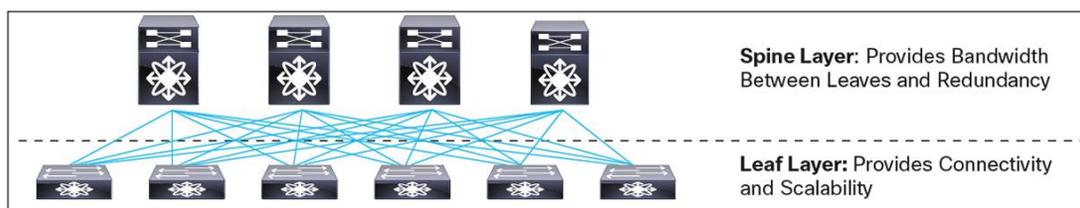
Cisco Application Centric Infrastructure (ACI) offers a revolutionary way of deploying, managing, and applying policy to networks. Unlike other software-defined networking (SDN) solutions, which focus solely on traffic forwarding or virtualization of physical devices, Cisco ACI approaches data center networking starting with policy and business requirements. Data center networks are deployed to support the delivery of applications; therefore, Cisco ACI starts there.

Cisco ACI fully automates the deployment of applications across the entire data center, including processes to meet security and user-experience requirements, using the business-level language of applications. To do this, Cisco ACI uses a combination of best-in-class hardware and software, rather than assuming one can exist without the other. This document discusses how to insert a Cisco ACI fabric into an existing data center environment and use it to automate the deployment of policy while integrating with existing systems.

Cisco ACI Overview

Cisco ACI fabrics use the Cisco Nexus® 9000 Series Switches as the core of the transport system. The Cisco Nexus 9000 Series was designed from the foundation to meet the rapidly changing requirements of data center networks, while enabling the advanced capabilities of Cisco ACI. At the physical layer, the Cisco ACI fabric consists of a leaf-and-spine design, or Clos network. This design is well suited to the east-west traffic patterns of modern data centers, moving traffic between application tiers or components. Figure 1 shows a typical spine-and-leaf design.

Figure 1. Typical Spine - and-Leaf Data Center Network Design



As shown in Figure 1, with this design each leaf connects to each spine, and no connections are created between pairs of leafs or pairs of spines. Leaf switches are used for all connectivity outside the fabric, including servers, service devices, and other networks such as intranets and the Internet.

With this architecture a spine switch provides cross-sectional bandwidth between leaf switches, plus additional redundancy. Bandwidth is determined by the number of spines and number of links to each spine. Redundancy is dictated by the amount of bandwidth lost in the event of a spine failure. In the topology in the figure, a single spine failure would reduce overall bandwidth and paths by 25 percent because of the use of four spines.

Above this physical layer Cisco ACI uses a controller, the Cisco Application Policy Infrastructure Controller (APIC), to manage the data center network and its policy centrally. Cisco APIC not only provides central management and automation, but also a policy model that maps application requirements directly onto the network as a cohesive system for application delivery. For more information about Cisco APIC, see <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-fabric/white-paper-c11-730021.html>.

This model provides full automation for the deployment and management of applications end to end, including Layer 4 through 7 policy, providing a single policy design, deployment, and monitoring point for applications. For more information about Cisco ACI and Layer 4 through 7 services, see <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/aci-fabric-controller/white-paper-c11-729998.html>.

Existing Network Infrastructure

Most existing networks are not built with a spine-and-leaf design, but may consist of various disparate devices typically configured in a three-tier architecture using a core-layer, aggregation-layer, and access-layer topology (Figure 2).

Figure 2. Typical Three-Tier Network Architecture

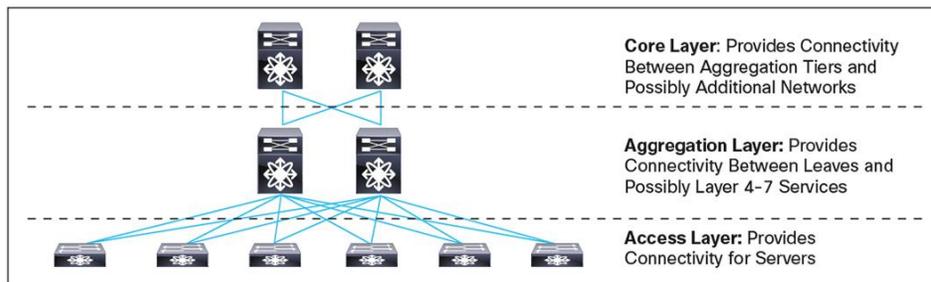


Figure 2 shows a typical three-tier data center network design. In this design, a single pair of switches is used at the aggregation layer and at the core layer, to provide redundancy for failure events. No more than two switches or routers are used at these tiers because of traditional Spanning Tree protocol constraints, which cause redundant links to be blocked, therefore negating the benefits of adding more devices.

In this model, the leaf switches are responsible for server connectivity and are then redundantly connected upstream to the aggregation layer. The aggregation provides connectivity between leaf switches and is typically also the point at which Layer 4 through 7 services are inserted. These services can consist of firewalls, load balancers, etc. Additionally, the aggregation layer is often the Layer 3 or routed boundary, or in some cases the core may provide this boundary.

This Layer 3 boundary design again must accommodate traditional Spanning Tree Protocol constraints and the need for Layer 2 adjacency for some server workloads. In addition, in this design the aggregation tier is the policy boundary for data center traffic. VLANs are typically created with one Layer 3 subnet within them. Broadcast traffic is allowed freely between devices within that subnet or VLAN. Policy (security, quality of service [QoS], services, etc.) is then applied only when traffic is sent to the default gateway to be forwarded between VLANs.

Investment Protection

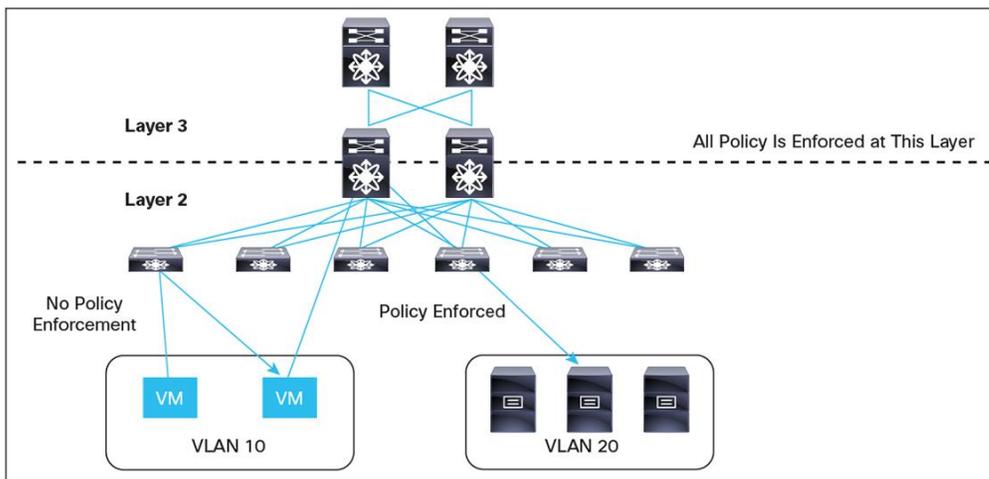
These topologies have been fairly standard for years. Therefore, customers have a large investment in the networking equipment that is in place. Other than in a completely new, greenfield environment, implementation of an all-new leaf-and-spine design will not be an option. Additionally, major changes to the existing physical or logical topology are typically not welcome because they can induce risk.

Because of these requirements, the design of Cisco ACI fabrics must include both compatibility with existing data center networks and the capability to easily integrate with those networks. The Cisco ACI fabric must be able to be inserted transparently into existing infrastructure while providing the same advantages of policy automation, linear scalability, and application mobility and visibility.

Cisco ACI is designed to provide integration with any existing network in any topology. Layers 2 and 3 can be extended into Cisco ACI, as well as Layer 3 data center overlay technologies such as Virtual Extensible LAN (VXLAN). Beyond these, the topology design and integration points must be carefully considered.

Because Cisco ACI focuses primarily on the design, automation, and enforcement of policy, the aggregation tier is the most logical insertion point for the Cisco ACI fabric. As stated earlier the aggregation tier is already responsible for policy enforcement and typically acts as the Layer 3 boundary. Therefore, traffic is already being moved to that tier for that purpose. Figure 3 provides a detailed view of the existing traffic pattern that will need to be integrated.

Figure 3. Three-Tier Traffic Patterns



Cisco ACI provides three methods of integrating with existing network infrastructure, these methods are:

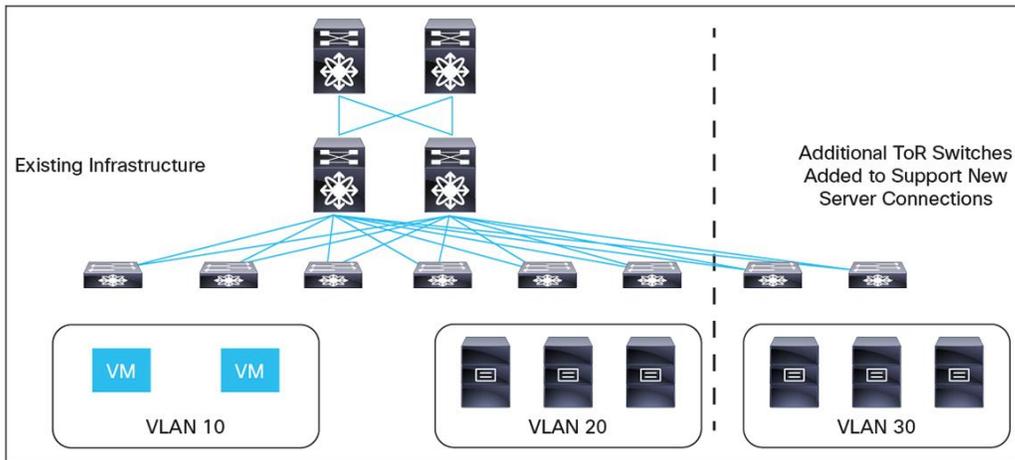
1. Cisco ACI Fabric as an Additional Data Center Pod
2. Cisco ACI Fabric as a Data Center Policy Engine
3. Cisco ACI Fabric Extended to Non-Directly Attached Virtual and Physical Leaf Switches

This white paper will cover options 1 and 2. Option 3 will be covered in a separate paper.

Method-1: Cisco ACI Fabric as an Additional Data Center Pod

This method utilizes a new pod build out to insert the ACI fabric. In this method, existing servers and services will not be modified or changed. ACI will be inserted as an aggregation tier for a new pod build out. This will act the same as attaching a new aggregation tier to an existing core for the purpose of a new pod. Figure 4 shows the traditional insertion of a new pod.

Figure 4. Traditional Method of Adding a New Data Center Pod



As you can see in Figure 4 the traditional method to add a pod is to attach a pair of new Aggregation switches to the existing core. New access switches are then connected to this aggregation tier to support new server racks. This method allows for the addition of new servers with the additional stability of separating out the aggregation layer services between pods.

Using this same methodology a new pod can be added to the existing network using an ACI Fabric. Rather than attaching two Aggregation switches and several Access switches a small ACI Spine/Leaf Fabric can be added. This is shown in Figure 5.

Figure 5. Adding a New Data Center Pod Using Cisco ACI

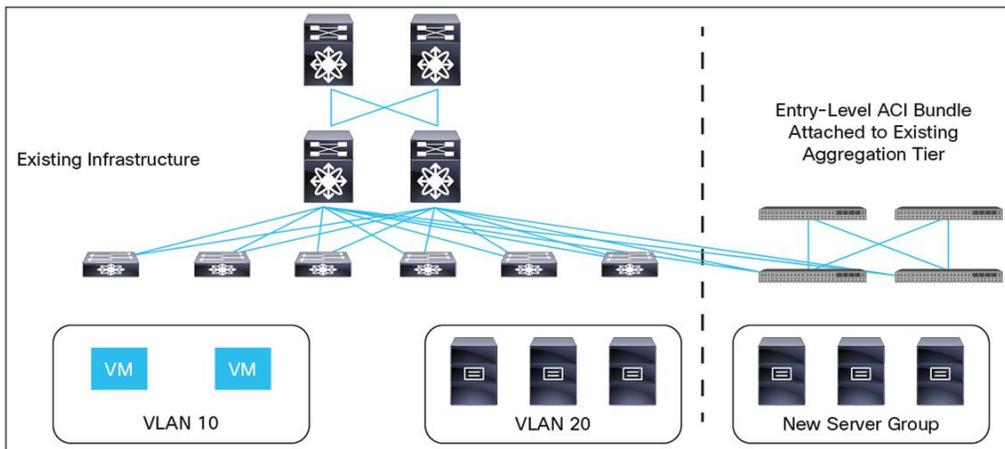
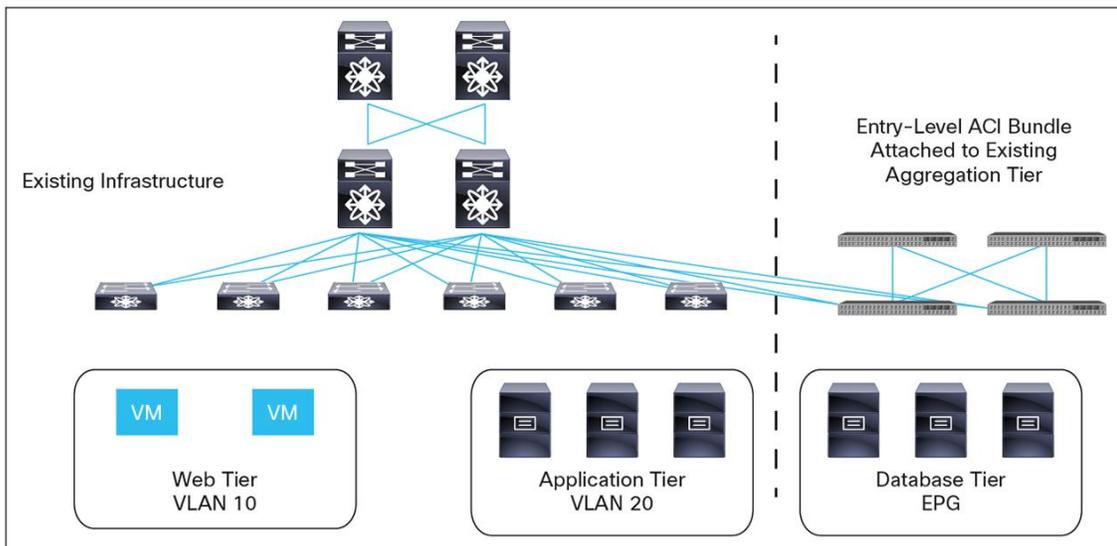


Figure 5 shows that this methodology works in a very similar fashion to the traditional pod addition shown in Figure 4. The key difference is that with a Spine/Leaf topology everything will connect to the Leaf switches, thus the existing Core is shown connected to the ACI leaf switches, and not to the Spine.

With the physical topology in place, the logical topology will need to be built. Cisco ACI is designed to integrate seamlessly with existing network infrastructure using standard protocols. Connections to outside networks are supported using OSPF, BGP, VxLAN and VLANs. The connection from the ACI leaf switches to the Core switches can be made using any of these. Figure 6 shows this.

Figure 6. Logical Connectivity to the New ACI Pod



Using Cisco ACI as a new data center pod provides investment protection for existing infrastructure while allowing growth into the benefits of the ACI Fabric. This methodology requires no topology, connectivity, or policy changes for existing workloads while providing a platform for ACI for new applications and services.

Cisco ACI Fabric as a Data Center Policy Engine Physical Topology

Aggregation tiers are used to provide cross connectivity between switches at the access tier and to enforce policy. Access-tier switches are added on an as-needed basis as the server environment expands, commonly using top-of-rack (ToR) switches when a new rack of servers is added. This traditional expansion-and-purchase cycle provides the optimal place to integrate Cisco ACI into an existing data center network environment. Figure 7 shows this expansion model.

Figure 7. Expanding the Access Tier

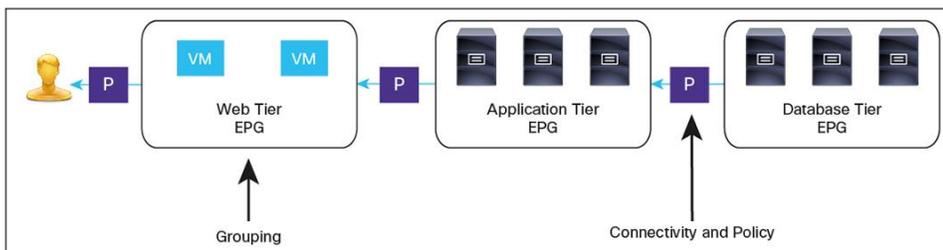


Figure 7 shows two ToR switches being added to support an additional rack of servers. This addition can be an expansion of an existing application environment or the deployment of a new application, or application tier or component. In these cases, the new switches are added at the top of the new racks and cabled redundantly to the existing aggregation tier. Integration of Cisco ACI into an existing data center network can follow this same model.

Using this model, an entry-level Cisco ACI fabric or bundle would be inserted in place of other ToR switches, but connected in the same way. This small leaf-and-spine topology can be placed at the top of a new rack or pair of racks and attached to the existing aggregation layer. Figure 8 illustrates this insertion of a Cisco ACI starter bundle as an integrated fabric.

Figure 8. Expanding the Environment Using an Entry-Level Cisco ACI Fabric

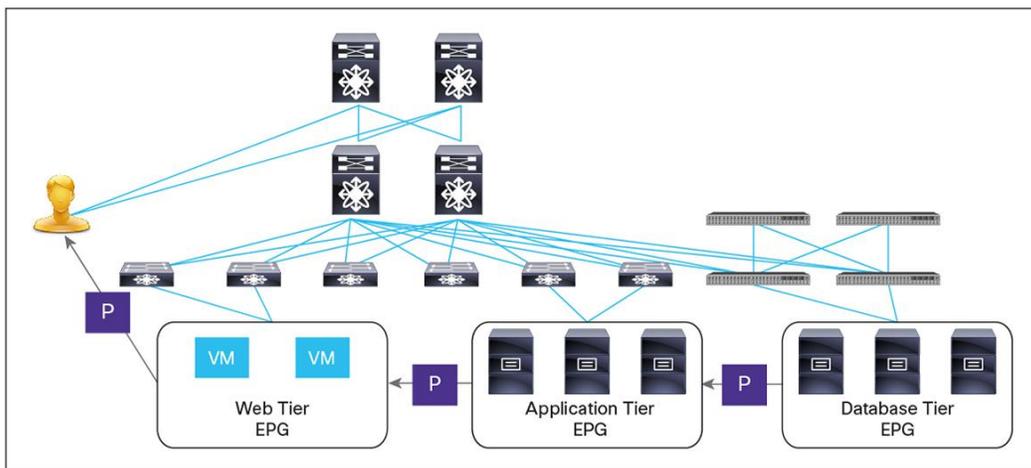


Figure 8 shows expansion of the access layer using an entry-level Cisco ACI fabric or bundle. The new Cisco Nexus 9000 Series leaf switches (Nexus 9396PX and 9396TX) attach to Cisco Nexus 9000 Series fixed spine switches (Nexus 9336PQ). The same leaf switches act as the connections to the existing aggregation tier. The new racks of servers connect to the leaf switches in the same way as they would have in a ToR expansion in a three-tier design.

With the physical topology in place, you now need to consider the logical topology. The example discussed here uses a three-tier web application that consists of a front-end web tier serving users, a middle application tier, and back-end database tier providing the data. Using the VLAN structure shown in Figure 5, this example uses VLAN 10 as the web tier, VLAN 20 as the application tier, and the new server group as the database tier. Figure 9 depicts this structure.

Figure 9. Example of an Integrated Environment Serving a Three-Tier Web Application

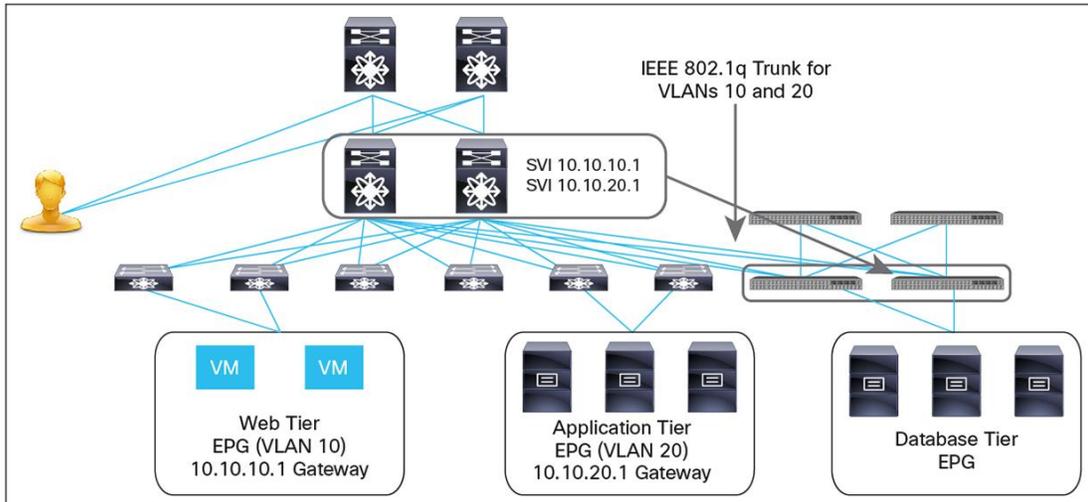


Figure 9 shows the three-tiers required to serve the web application spread across the existing environment and the new Cisco ACI pod. The web and application tiers remain segregated as before using VLANs and subnets, and the database tier sits on the Cisco ACI fabric and is designated by its endpoint group (EPG). Connectivity is required between these tiers to properly serve the end user. Additionally, policy enforcement needs to be provided for the connections. For more information about EPGs and the Cisco ACI logical model, see <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-731310.html>.

The logical model of Cisco ACI provides the capability to group objects based on the services they deliver and the policies and connectivity required. This model 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. Figure 10 shows how this sample three-tier application might look in the Cisco ACI logical model.

Figure 10. Three-Tier Web Application Example in the Cisco ACI Logical Model

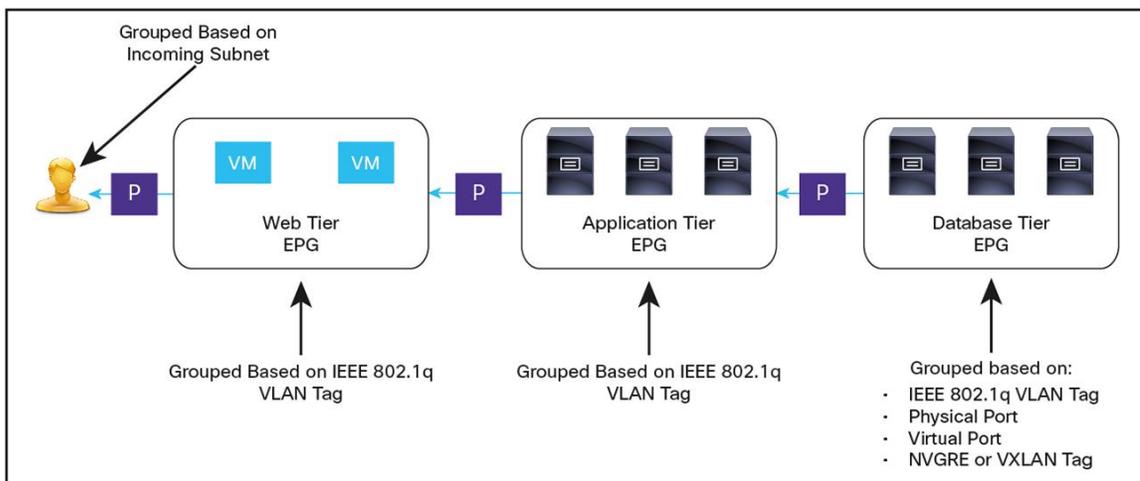
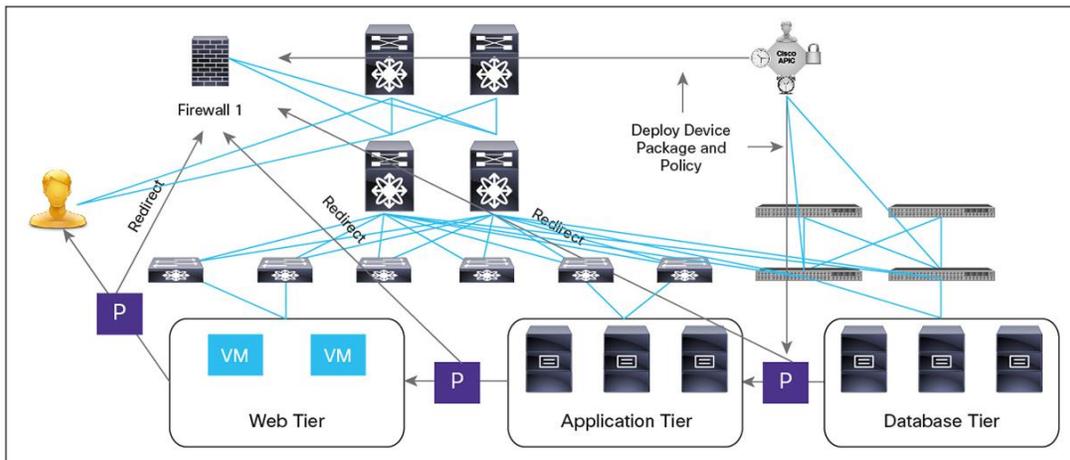


Figure 10 shows a logical diagram of a three-tier web application as an application network profile (ANP), which is the construct used to design and configure applications in Cisco ACI. The groups represent virtual and physical servers or services. The groups are then connected using policy contracts to define communication.

This same model would apply for connectivity and policy instantiation between the Cisco ACI fabric and the existing data center network. Using Cisco ACI as the policy automation and enforcement engine between tiers, traffic from the existing network would be identified and grouped, and then connectivity and policy would be enforced as designed in the Cisco ACI ANP. Figure 11 shows this relationship.

Figure 11. Logical and Physical Mapping Across an Integrated Cisco ACI Network

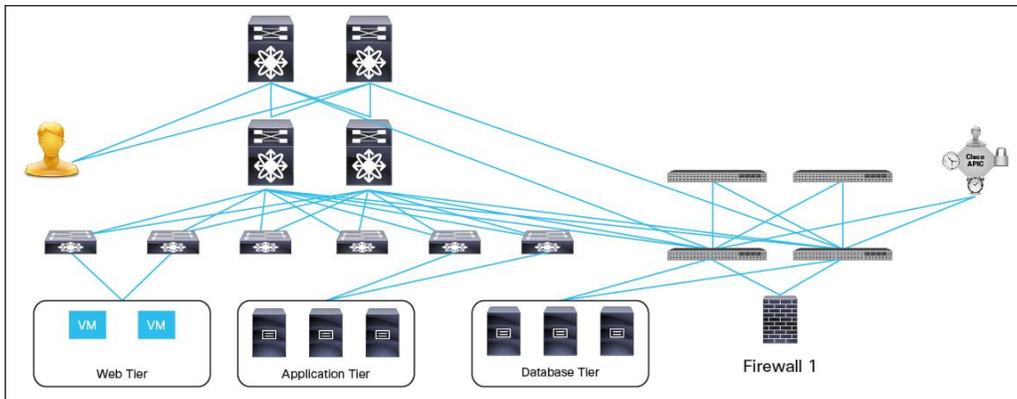


In Figure 11, the gray arrows show the logical connectivity, and the policy enforced for those connections, the blue lines show the physical connectivity. The logical connectivity remains the same from the perspective of Cisco ACI; it is simply the connection of groups of objects and application of policy. From a physical perspective, the traffic flow needs to support the use of Cisco ACI as the policy engine for this network.

To enforce traffic and policy using Cisco ACI in this example, two tasks must be accomplished: the traffic must pass through the Cisco ACI fabric, and the appropriate group for the traffic must be identified. By having traffic between tiers pass the Cisco ACI fabric and be identified appropriately, Cisco ACI can easily automate all the policy for the traffic. This approach is the same that the aggregation layer typically uses today for policy enforcement.

To perform the first task, traffic needs to be routed to the Cisco ACI fabric. This task can be accomplished by moving the default gateway. As Figure 3 earlier in this document showed, the aggregation layer typically acts as the default gateway or Layer 3 boundary. Therefore, traffic between VLANs is pushed to this layer anyway. This pushing is accomplished by creating a switch virtual interface (SVI) and assigning the gateway address to this SVI. Servers in that subnet send traffic destined for other subnets to the SVI. If instead the SVI is moved to the Cisco ACI fabric, traffic between existing groups (VLANs and subnets) is then handled by the Cisco ACI fabric. Figure 12 shows this behavior.

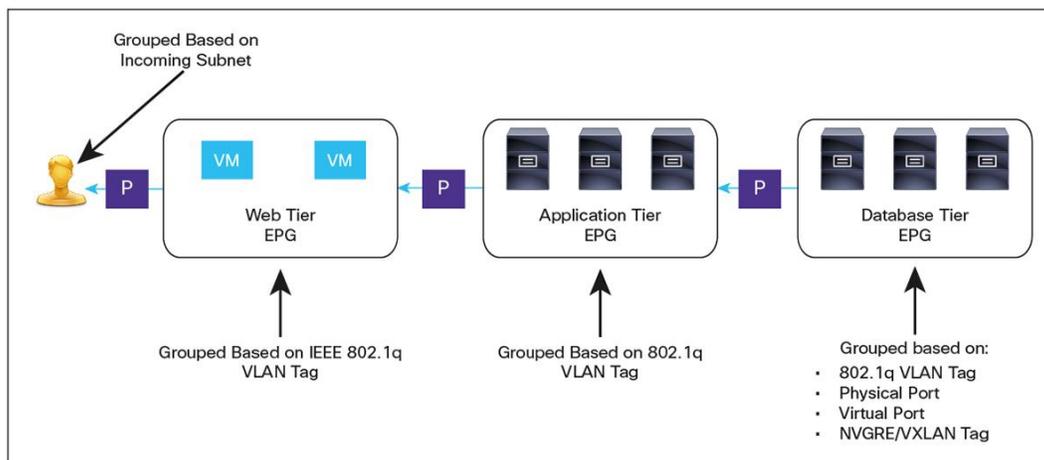
Figure 12. Migrating the Default Gateway to Cisco ACI for Policy Enforcement



In Figure 12, the SVIs act as the default gateways for the web and application tiers on the existing aggregation switches. These SVIs need to be migrated to the Cisco ACI fabric. After the SVIs, and therefore the default gateways, are migrated, all traffic between the web VLAN and application VLAN will be forwarded by the Cisco ACI fabric. This migration allows the Cisco ACI fabric to act as the policy enforcement and automation boundary.

The next step is to identify the traffic from the web and application tiers that resides on the existing network as groups within the Cisco ACI fabric. Cisco ACI provides several mechanisms for determining which objects should be grouped. One of these methods uses the IEEE 802.1q VLAN tag carried in the packet. With this method, the web tier can be identified as VLAN 10, and the application tier as VLAN 20, fitting directly into the Cisco ACI policy model. Figure 13 shows the identification methods for each group in this three-tier web application.

Figure 13. Group Identification in the Cisco ACI Fabric



As shown in Figure 13, the existing VLAN identifier can be used to group objects as EPGs within the Cisco ACI fabric. These can then be tied together or to other new groups that exist directly attached to the Cisco ACI fabric. The figure also shows the capability to group objects based on the external subnet. After groups have been identified based on any given characteristic, connectivity and policy can be automated and applied uniformly in Cisco ACI regardless of whether the groups are directly connected or use virtual or physical resources.

With the physical and logical topology in place, the final step is to apply the necessary policy for traffic connections. Policy consists of application-level requirements - security, service-level agreements (SLAs), user experience, etc.- that can be implemented in many ways. For instance, implementation of policy can take the form of inbound and outbound permit and deny rules, logging, and redirection and instantiation of Layer 4 through 7 service graphs.

Instantiating and Automating Policy in an Integrated Environment

Figures 12 and 13 depict the integration of physical and logical topologies across a mixed environment of Cisco ACI and existing networking equipment. From a connectivity perspective, the job is complete. The last piece to build on top of this structure is policy automation. Because Cisco ACI is now the policy enforcement boundary, all policy can be automated by the Cisco APIC.

In many instances, the policy between existing groups may exist as simple access control lists (ACLs) or QoS marking and enforcement. In these cases, those rules can be left in place for enforcement at the existing aggregation layer, or they can be migrated to policy contracts within the Cisco ACI fabric.

In addition to simple policy enforcement, complex enforcement using Layer 4 through 7 devices may be required. In these cases, traffic needs to be redirected to those devices. Cisco ACI also can integrate with a broad system of Cisco and other Layer 4 through 7 devices so that it can provision and manage rules on those devices.

Continuing with the earlier 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 expansion, 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, providing a single point of policy deployment. Figure 14 shows both methods being used to deploy policy in the sample integrated environment.

Figure 14. Integrating Layer 4 Through 7 Service Devices with the Cisco ACI Network

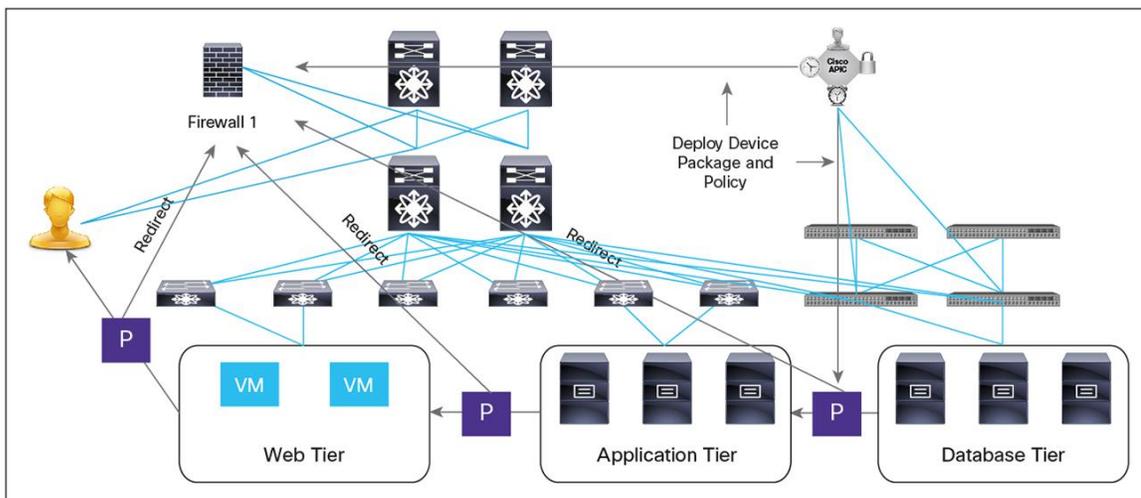


Figure 14 shows policy contracts for the connectivity between the users and the web and application tiers as simple traffic redirects. Redirects are used with the assumption that the policies have already been configured on those devices. The database-tier contract is being deployed for a new application tier on the Cisco ACI fabric, so here policies will not be in place on the device. For this case, Cisco APIC is used in conjunction with a device package to instantiate the required policies on the firewall shown. Traffic will then be redirected to the firewall for enforcement. A firewall is used for this example, but the same would be the case with any Layer 4 through 7 device.

Because policy is the change point in a network as applications are added, removed, expanded, or moved, policies are the most important elements to automate. Existing VLANs, subnets, etc. have already been configured and do not require frequent modification. Therefore, using this method, Cisco ACI can fully automate the point of change - policy - without the need to configure or automate the existing network equipment.

Figure 14 shows stage 1 of migration to the Cisco ACI fabric. At this stage, Cisco ACI is acting as an extension to the existing aggregation tier; it is the central policy enforcement and automation system for any application tiers or components whose traffic is redirected into the fabric as shown.

Stage 2 is the optimization of traffic flows for those applications. To optimize traffic flows, the network topology must be modified to provide Cisco ACI with direct connectivity to Layer 4 through 7 service nodes and the existing network core. To do this, the links to the core devices, Internet and intranet connections, and service devices must be moved to attach to the Cisco ACI fabric leaf switches. Figure 15 shows this second stage of Cisco ACI integration.

Figure 15. Stage 2 of Cisco ACI Integration with Existing Networks

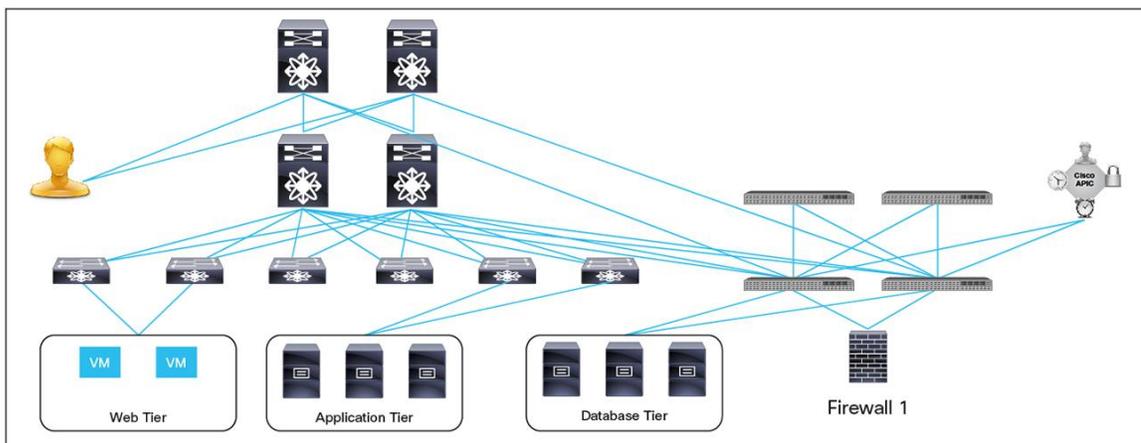


Figure 15 shows that the existing connections for the servers in each tier have not changed. The connections for Layer 4 through 7 services (represented by the firewall) have been moved to the Cisco ACI leaf switches. In addition, new connections have been created directly connecting the existing core tier to the new Cisco ACI fabric leaf switches. These connections provide a direct path for inbound and outbound data center traffic to groups on the Cisco ACI fabric.

Conclusion

Cisco ACI provides an enhanced feature set for the automation of application deployment and the policies required for its connectivity. These features are accelerated by a combination of hardware and software, as well as a network topology designed for today's traffic patterns.

Cisco ACI can be deployed in existing data center environments and is designed to work alongside the infrastructure already in place. By deploying Cisco ACI as an extension to the aggregation layer in existing three-tier topologies, customers can gain full policy automation while maintaining the same policy enforcement point currently in use. Because Cisco ACI can work in existing environments, it also provides investment protection for any infrastructure in use.

For More Information

- Cisco APIC: <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-fabric/white-paper-c11-730021.html>.
- Layer 4 through 7 integration: <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/aci-fabric-controller/white-paper-c11-729998.html>.
- Policy model: <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-731310.html>.



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)