

ACI – Network-Centric Approach

What You Will Learn

Cisco Application Centric Infrastructure is gaining increasing presence within data centers predominantly due to the automation it offers within the network. However, several Enterprises are looking to have a more gradual transition from their currently, well understood classical Ethernet networks to the ACI model.

In this paper you'll learn Cisco's "network-centric approach" with ACI, which allows for that gradual transition. You'll see how you can,

- Retain your current network configurations
- Allow your network engineers and administrators the ability to familiarize themselves to ACI in a model well-known to network administrators, while leveraging many benefits inherent to Cisco ACI.

How to Implement

Physical Network	Application Profiles	Contracts
Multi-tenancy	Layer 2 and 3	Layer 3 External Connection
Layer 4-7 Services	End Points	


When using the network-centric approach to ACI you get the ability to gradually port your current network way of operating to the ACI fabric so that it is working in a familiar, well understood, and classical manner. You can then continue to move towards an application-centric model at a pace appropriate and comfortable to your business. With the network-centric approach, you are still using all the familiar ACI constructs, such as switch profiles and policies, Bridge Domains, and End Point Groups; you are just using them in a slightly different way. Under the covers, they're simply mapping directly to existing subnets, VLANs, and routing protocols.

For those who are more comfortable with the CLI, most if not all of these configurations can be done using the [NX-OS CLI](#).

Physical Network Architecture

Legacy network designs are based on three tiers: core, aggregation, and access. For modern DC network designs a 2-tier spine/leaf design is recommended. This architecture is optimized for east-to-west traffic flows, which are predominant in a modern datacenter. For the purpose of this paper we will assume a new ACI fabric is brought up alongside the existing classical Ethernet network with L2 or L3 connections from the leaf switches to the existing network.

Figure 1. ACI Fabric



As with any approach you take to ACI, all endpoints will be attached to the leaf switches including the APICs. The ACI fabric can be thought of as one logical router, and we can use the GUI, CLI, or APIs to access each individual switch, or the fabric as a whole. In most cases, the common choice is to manage the whole fabric as a system via the APIC Management console. Although the ACI GUI hides a few entities under the covers, we can always issue show commands via the CLI. For example, you can see where VLAN trunks are set up. This isn't absolutely necessary, but may be helpful to a network engineer who is used to monitoring and troubleshooting that way.

Setting up switch profiles and policies will also still be necessary, though not very much different from setting up ports on the command line. From the APIC GUI you go to the Fabric tab and you can either use Advanced or Basic mode to set up ports, port channels, or virtual port channels (vPCs). Basic mode allows for very quick configurations of commonly used constructs. For example, instead of logging on to each switch and using the CLI to configure each individual port or range of ports, you can do this easily and efficiently by creating switch profiles via the APIC. Also, vPCs in ACI no longer require that peer links be configured as it is all done behind the scenes within the fabric. So essentially, from the Basic GUI,

- Click a button to add two switches.
- Click the ports you want to be a part of the vPC.
- Select the correct LACP, CDP, LLDP policies.
- Apply the changes.

Within seconds and with very few steps, you are able to configure vPCs on multiple switches.

Multi-tenancy

Depending on organizational structure, multi-tenancy may or may not be a new concept. Tenants are containers from an administrative perspective. These containers store your logical networking and any end points. Tenants can provide isolation as well as play into the role-based access control model that ACI uses. For example, you could have three tenants: Dev, QA, and Prod. Only some admins are allowed to make configuration changes in Dev and some others in Prod. Some admins may also just have read only access to any of these tenants. There is also a fabric administrator who has access to everything. You can use multitenancy to separate workloads into a network-centric tenant and an Application-Centric tenant. One tenant can be used for existing workloads in the network-centric model, a second tenant can be used for new applications deployed in Application-Centric mode thereby giving you a safe separation between legacy workloads and new workloads brought directly into ACI at inception.


Tenants are used largely by service providers, but the model is gaining prevalence in the Enterprise and Commercial space as well. As corporations are embracing cloud models more, they want to be able to treat their on-premises stack as they would their public cloud, which may even include new charging models to pay for resource usage. Tenants do not actually separate traffic inherently in ACI, but they do allow for the implementation of security domains, end point policy and role based access control both within and between tenants.

Within each tenant a separate VRF can be created allowing overlapping IP space between tenants. VRF stands for Virtual Routing and Forwarding and simply means you can have overlapping logical routers in the same switch (or same switch fabric, in the case of ACI). ACI actually uses VRF-lite within the fabric. For more information see [Connecting Application Centric Infrastructure \(ACI\) to Outside Layer 2 and 3 Networks](#). The VRF allows for the separation of traffic between tenants, or more accurately between VRFs as well. ACI is flexible in that each tenant can have its own VRF or you can share a given VRF among tenants, or any mix of the two.

Application Network Profiles

Don't be thrown off by Application Network Profile as it accentuates the term application. It is a container for End Point Groups, Contracts, and Layer 4-7 Devices.

Figure 2. Application Network Profile




We'll dive into each of these components throughout the rest of the white paper.

Layer 2 and Layer 3

Datacenters built prior to ACI use VLANs for the purpose of isolation. VLANs are broadcast domains that allow frames to be sent out all ports of a switch tagged with that VLAN, if the frame has no awareness of the destination. This is called flooding. VLANs are generally mapped to one subnet. For example, you may have VLAN 21 which contains all of your database servers. It is likely that these servers will only be assigned to one subnet, perhaps 192.168.21.0/24. Usually a black list model is used, meaning traffic is allowed by default within subnets. Security rules are typically assigned at the Layer 3 boundary or default gateway using Access Control Lists (ACL) or Firewall rules.

Figure 3. Traditional VLAN/Subnet Model

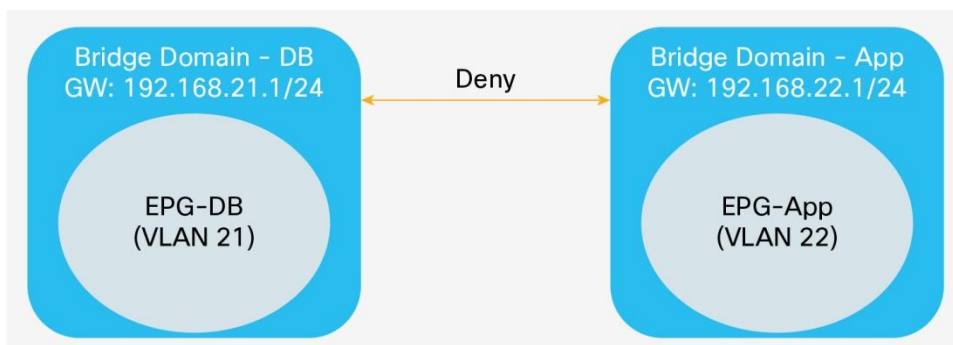


ACI uses layer 2 and 3 constructs called Bridge Domains and End Point Groups (EPGs). You can think of Bridge Domains as being the same as VLANs in this case. The bridge domain contains a gateway, or SVI. The SVI acts as a pervasive gateway for our endpoints. In network-centric mode you have only one SVI, or subnet, contained within a bridge domain.

End point groups are just that – groups of end points. It's simply a container for virtual and physical servers, or Linux containers; end points. There may or may not be extra configurations, but let's keep it simple for right now. In network-centric mode, you specify end points all belonging to the same VLAN, contained within an EPG. So there is a one-to-one-to-one mapping between the Bridge Domain, EPG, Subnet, and VLAN. You may see this described as VLAN=BD=EPG.

Note: A word on scalability: With routing turned on, there is a limit of 1,750 Bridge Domains/VLANs/EPGs per leaf. Without routing, [legacy mode](#) becomes an option and scale doubles. Please check the latest scalability guides from <http://cisco.com/go/aci> for more information on static bindings per EPG.

Figure 4. Relationship between BD, EPG, VLAN, and Subnet in Network-Centric Approach




Notice Figures 3 and 4 look very similar. VLAN 21 is mapped to Subnet 192.168.21.0 and all endpoints within that subnet, containing Database servers. VLAN 22 is mapped to subnet 192.168.22.0 and all endpoints within that subnet, containing application servers. In the network-centric approach, it's not necessary to know which applications, or application tiers exist on a given VLAN. But it is necessary to know which VLANs should be allowed to communicate. There is one noticeable difference between traditional networks and ACI – ACI uses the white list model. Hence, by default no traffic is allowed between these subnets, though traffic is allowed within the bridge domains, by default. This default behavior can be easily changed to match traditional networking concepts. To address the change in the white list model, we need to discuss contracts.

Contracts

Contracts are stateless firewall rules. They allow or deny EPGs to communicate over specified protocols such as http or https. They go between EPGs. But in the network-centric case, since an EPG, subnet, and VLAN are basically all the same, contracts behave the same as ACLs. There is an option, when configuring a contract, to allow any traffic to traverse between EPGs. This is the more common approach in this mode, unless extensive ACLs are used in the legacy environment.

Figure 5. Creating a contract in APIC GUI



A contract in ACI is made up of two parts. Just like an ACL includes ACEs (access control entries), a contract in ACI is made up of filters (like ACEs) that are grouped into Contracts (like an ACL). Because ACI is built around a whitelist model where the entire fabric acts as a pervasive and integrated stateless firewall, it is important to understand that contracts are required (even if they allow everything) for communication between EPGs. Unlike ACLs, that use IP subnets and ports to define communication, contracts use only ports and protocols, and not subnets. Enforcement of contracts is done at the EPG level. With simple EPG membership (regardless of IP address), you can apply policy with ease in any part of an ACI Fabric.

The highlighted option in Figure 5, when checked, shows the example of No Filter (Allow All Traffic). The general nature of contracts is to permit communication. However, it is also possible to use what are called Taboo Contracts. Taboo Contracts are filters within contracts that allow you to deny particular types of traffic. For example, you could create a taboo contract alongside a contract permitting everything, which would deny Telnet traffic from your end user devices to your management devices. In this case, end user devices are contained within a VLAN (or even different VLANs) that equate to an EPG and the same would happen for the management servers. A taboo filter within a contract is applied between the two EPGs (or VLANs) and would deny users from using Telnet to connect to the management devices.


One benefit, with the application-centric or network-centric approach is that both contracts and filters are re-usable. For example, you can create a contract once and then copy it again between two other EPGs. In the case of the filter, you can create a filter that allows SSH specifically within a contract. Then you can reuse that filter again in another contract. This means you are saving administrative time in configuration, and you are eliminating any human error that could occur. Note that contracts can contain any mix or range of ports and protocols defined in filters.

L3 Out

To use existing firewalls (or other layer 3 gateways like routers with Internet connections) you can make use of an ACI concept called Layer 3 Out. It is important to note that an ACI Fabric is effectively a big L3 routed fabric that advertises local routes and learns external routes. For example, you connect a firewall or router to a leaf switch using a single line or a port channel, then use a routing protocol like OSPF, EIGRP or BGP to exchange routes between the device and the ACI fabric. Inside the fabric you build the connection using a combination of SVI and VLAN and then create and apply this via an External Routed Domain. You don't need to make any changes to the default gateways or FHRP on your traditional networks. In other words, leave them where they are currently in the classical Ethernet environment, and later migrate from the old network to the ACI fabric. Once you migrate all workloads completely to ACI you can move the gateways for those VLANs and subnets inside the fabric. It is possible to do this on a per-VLAN basis.

If you bring external routes into the fabric it will be necessary to enable MP-BGP, which is not enabled by default, and use an AS number. Note that this will only run within the fabric and not between the fabric and any external device or network. See below for a diagram of the connectivity. The reason is that it is related to well-known routing concepts. In our example, the OSPF neighbor adjacency is between the Firewall and the first leaf node. The other nodes do not have any neighbor relationship. In order for the OSPF leaf to tell the rest of the fabric about any routes it learns, it will have to use route redistribution. Because ACI is multi-tenant, the protocol of choice is BGP because it carries more information than just routes. We can share information about tenants/VRFs too.

Figure 6. OSPF Neighbor Adjacency Between Leaf Node and Firewall



Layer 4-7 Devices


The ACI fabric allows you to stitch in network services, as mentioned above. Before the 1.2.2 release, a device package was necessary to integrate firewall, load balancing, and other services. You could then use the APIC to manage the device, which is ideal for automation.

However, with the latest release it is possible to stitch these services in without a device package and then continue to manage the device as you would in your current networks, using its native device manager. Either way, traffic will still flow through these devices and be permitted, denied, load balanced, etc., as is necessary. As a note, Layer 4-7 devices are not restricted to only Cisco devices. Cisco has over 50 ecosystem partners interoperable with the ACI fabric. When opting for a network-stitching only approach, no device package is needed, which effectively means you can support any service appliance or VM from any vendor that speaks Ethernet and IP.

End Points

End points, refers to servers (both virtual and physical) or any layer 4-7 device (again virtual or physical), that can be added to the fabric in the same way they are with application-centric mode. We can use static binding to connect bare metal and use either static binding or [VMM integration](#) to connect to our hypervisors.

Figure 7. Logical model hierarchy



Conclusion

The network-centric approach to configuring ACI is a great way to migrate your current network to an ACI fabric. You continue to get many benefits over traditional networking such as:

- Staying with existing IT structure and processes
- Next-generation 10, 25, 40, and 100GB data center network with the use of Nexus 9000 series switches
- Optimized spine/leaf architecture for east-west traffic
- Workload mobility
- One place to manage the fabric

-
- Reusable layer 2 – 7 constructs, mitigating human error, and reducing administrative time spent
 - Real time monitoring with health scores and atomic counters
 - Easier troubleshooting with more visibility and tools such as End Point Tracker
 - Easier scalability through ease of fabric discovery and tools such as ACI Optimizer
 - A programmable fabric

For more information on constructs in ACI please see this [white paper](#) and for all documentation on Cisco ACI please go to <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)