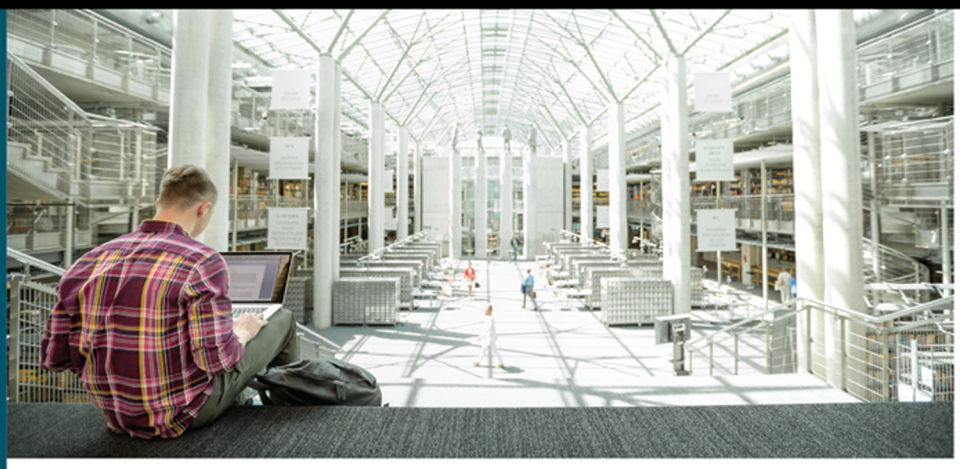

Deploying ACI

The complete guide to planning, configuring, and managing Application Centric Infrastructure

Frank Dagenhardt, CCIE® No. 42081
Jose Moreno, CCIE® No. 16601
With contributions from **Bill Dufresne,** CCIE® No. 4375

ciscopress.com

Deploying ACI

The complete guide to planning, configuring, and managing Application Centric Infrastructure

Frank Dagenhardt, CCIE No. 42081,
Jose Moreno, CCIE No. 16601,

With contributions from
Bill Dufresne, CCIE No. 4375

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

- **Hybrid mode:** A mode that borrows features from network-centric and application-centric modes. Enterprises running in this mode are using additional features and levels of security. Your ACI network may be running in network-centric mode with the addition of integrated services and/or more granular contracts. You may be running some of your network in network-centric mode, and other parts where groups and contracts are defined on an application-by-application basis.
- **Application-centric mode:** Application-centric mode gives ACI users the highest level of visibility and security. In this mode, we define groups and contracts based on the individual applications they service. In this mode, groups may contain an entire application, or we may break the applications up into tiers based on function. We can then secure communications by only allowing the traffic that should be allowed between certain tiers or devices. We also have the ability to insert services on a hop-by-hop or tier-by-tier basis. This is the zero-trust model that most enterprises are trying to take advantage of for some or all of their applications. This mode also offers us the ability to track application health and performance on an application-by-application basis.

The point of the preceding explanations is that each of these modes uses the group and contract model. Sometimes there is confusion when we say mode X, Y, or Z. It sounds like we are flipping a switch to change ACI into something different. In reality, it's a friendly name given to the amount of configuration and work effort that is required when you choose to implement one configuration type/mode or the other.

Note One of the wonderful things about ACI is its dynamic nature. If you are considering adding additional features, you have the ability to spin up a tenant on the fly and test proposed changes without affecting production. This is a great way test configuration elements when moving from network-centric mode to hybrid mode to application-centric mode. It is always prudent to reference the scalability guide for the software release you are working with when designing or making changes to your ACI network. It is recommended that you use ACI Optimizer to examine the resource requirements of your specific design or proposed changes. We will explore ACI Optimizer in chapter 12. Additional information on ACI Optimizer can also be found here: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_ACI_Optimizer.html.

Let's explore groups in more detail. Cisco ACI can classify three types of endpoints:

- Physical endpoints
- Virtual endpoints
- External endpoints (endpoints that send traffic to the Cisco ACI fabric from the outside)

The administrator determines how the hardware and software classifies the traffic. Some versions of hardware and software have different classification capabilities, which we will

discuss later. A group can contain one device or many devices. A single device can be a member of one or more groups based on its connectivity to the network. A server with two network interface cards may have its production-facing interface in a group called "web servers" and another interface dedicated to backups in a group called "backup." You can then control the communication of these groups individually through the use of contracts. By default, all devices inside of the same group can talk to each other freely. This behavior can be modified with a feature called *intra*-EPG isolation, which is similar to a private VLAN where communication between the members of a group is not allowed. Or, intra-EPG contracts can be used to only allow specific communications between devices in an EPG. In all configurations, members are always allowed to talk to the Switchted Virtual Interfaces or gateways that exist within their associated bridge domains. A group can only be associated with one bridge domain at a time. Multiple EPGs can exist in the same bridge domain.

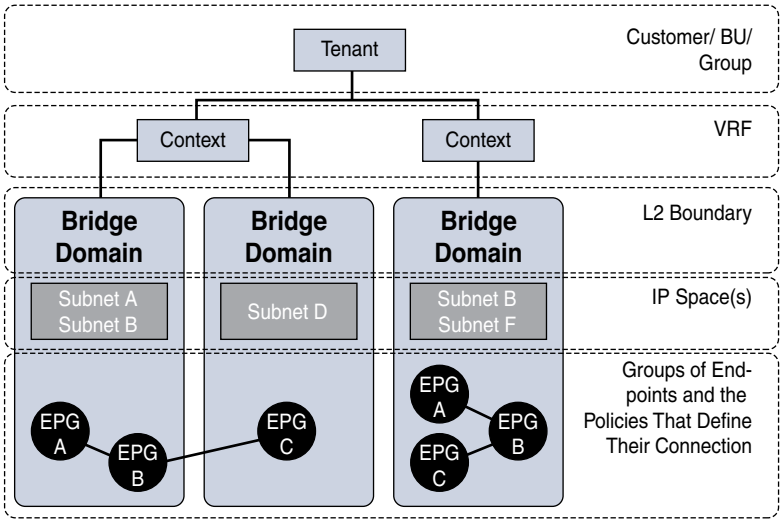

Figure 5-1 *ACI Policy Hierarchy*

Notice the groups depicted in Figure 5-1. EPGs A and B will only be able to use the SVIs for subnets A and B. You will also be able to add a device from either subnet to either group. This highlights the flexibility of having more than one subnet associated with a bridge domain. EPG C will only be able to use the SVI for subnet D, and therefore you can only add devices from subnet D to EPG C. Another, less frequently used option is to associate the subnet with the EPG. Only devices from that group would be able to use that SVI versus every group in the bridge domain having access to the SVI. From the group standpoint, this association of bridge domains to groups only influences which devices can be put into which groups. It does not have any effect on the capability for groups to communicate with one another. We will cover the network implications of the bridge domain with which a group is associated in the next section.

Three different types of groups are available within ACI. The first group is a traditional end point group (EPG). The second is a microsegmented EPG. The third is an external EPG that is created when we connect to an external network, outside of the fabric.

When you create an external routed connection (L3 Out) or an external bridged connection (L2 Out), you also create a new EPG associated with that connection. The devices that are reachable via these external connections become associated with the newly created EPG. Communication to these devices will be controlled based on the configuration of the external EPG and the contracts between it and other devices on the network.

Using traditional and microsegmented EPGs, you can assign a workload to an EPG as follows:

- Map an EPG statically to a port and VLAN.
- Map an EPG statically to a VLAN switchwide on a leaf.
- Map an EPG to a virtual machine manager (VMM) domain (followed by the assignment of vNICs to the associated port group).
- Map a base EPG to a VMM domain and create microsegments based on virtual machine attributes (followed by the assignment of vNICs to the base EPG).
- Map a base EPG to a bare-metal domain or a VMM domain and create microsegments based on the IP address (followed by the assignment of vNICs to the base EPG).

Note If you configure EPG mapping to a VLAN switchwide (using a static leaf binding), Cisco ACI configures all leaf ports as Layer 2 ports. If you then need to configure an L3 Out connection on this same leaf, these ports cannot then be configured as Layer 3 ports. This means that if a leaf is both a computing leaf and a border leaf, you should use EPG mapping to a port and VLAN, not switchwide to a VLAN.

The administrator may configure classification based on virtual machine attributes, and, depending on the combination of software and hardware, that may translate into a VLAN-based classification or MAC-based classification.

Hardware-based switches (depending on the ASIC model) can classify traffic as follows:

- Based on VLAN or VXLAN encapsulation
- Based on port and VLAN or port and VXLAN
- Based on network and mask or IP address for traffic originated outside the fabric (that is, traffic that is considered part of the Layer 3 external traffic)
- Based on source IP address or subnet (with Cisco Nexus E platform leaf nodes and EX platform leaf nodes)
- Based on source MAC address (Cisco Nexus EX platform leaf nodes)

It is possible to configure classification of the incoming traffic to the leaf as follows:

- Based on VLAN encapsulation.
- Based on port and VLAN.
- Based on network and mask or IP address for traffic originating outside the fabric (that is, traffic that is considered part of the Layer 3 external traffic).
- Based on explicit virtual NIC (vNIC) assignment to a port group. At the hardware level, this translates into a classification based on a dynamic VLAN or VXLAN negotiated between Cisco ACI and the VMM.
- Based on source IP address or subnet. For virtual machines, this function does not require any specific hardware if you are using Application Virtual Switch (AVS). For physical machines, this function requires the hardware to support source IP address classification (Cisco Nexus E platform leaf nodes and later platforms).
- Based on source MAC address. For virtual machines, this function does not require specific hardware if you are using AVS. For physical machines, this requires the hardware to support MAC-based classification and ACI version 2.1 or higher.
- Based on virtual machine attributes. This option assigns virtual machines to an EPG based on attributes associated with the virtual machine. At the hardware level, this translates into a classification based on VLAN or VXLAN (if using AVS software on the virtualized host or, more generally, if using software that supports the OpFlex protocol on the virtualized host) or based on MAC addresses (Cisco Nexus 9000 EX platform with VMware vDS).

Note Each tenant can include multiple EPGs. The current number of supported EPGs per tenant is documented in the Verified Scalability Guide at Cisco.com (https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html).

An EPG provides or consumes a contract (or provides and consumes multiple contracts). For example, the NFS EPG in Figure 5-2 provides a contract that the External EPG consumes. However, this does not prevent the NFS EPG from providing the same or different contracts to other groups, or consuming contracts from others.

The use of contracts in Cisco ACI has the following goals:

- Define an ACL to allow communications between security zones.
- Define route leaking between VRFs or tenants.

Figure 5-2 shows how contracts are configured between EPGs (for instance, between internal EPGs and external EPGs).

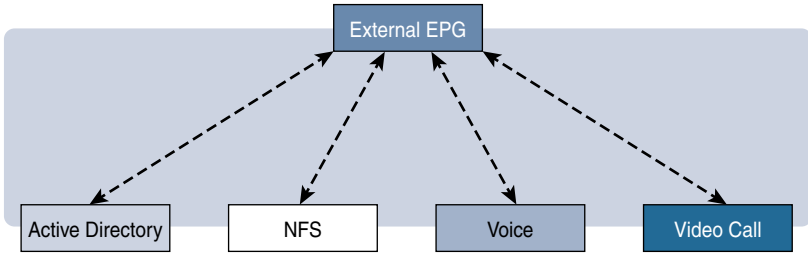

Figure 5-2 *External EPG Contract Example*

Contracts Are ACLs Without IP Addresses

You can think of contracts as ACLs between EPGs. As Figure 5-3 illustrates, the forwarding between endpoints is based on routing and switching as defined by the configuration of VRF instances and bridge domains. Whether the endpoints in the EPGs can communicate depends on the filtering rules defined by contracts.

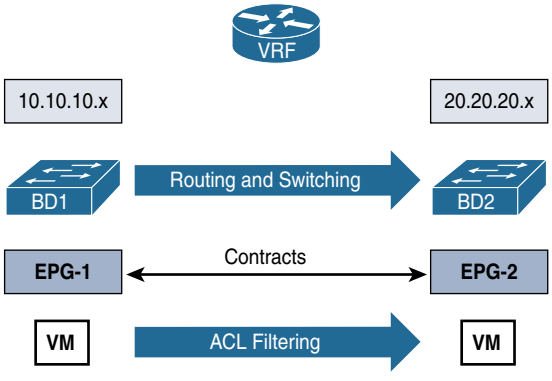

Figure 5-3 *Contracts Are ACLs*

Note Contracts can also control more than just the filtering. If contracts are used between EPGs in different VRF instances, they are also used to define the VRF route-leaking configuration.

Filters and Subjects

A *filter* is a rule specifying fields such as the TCP port and protocol type, and it is referenced within a contract to define the communication allowed between EPGs in the fabric.

A filter contains one or more filter entries that specify the rule. The example in Figure 5-4 shows how filters and filter entries are configured in the APIC GUI.

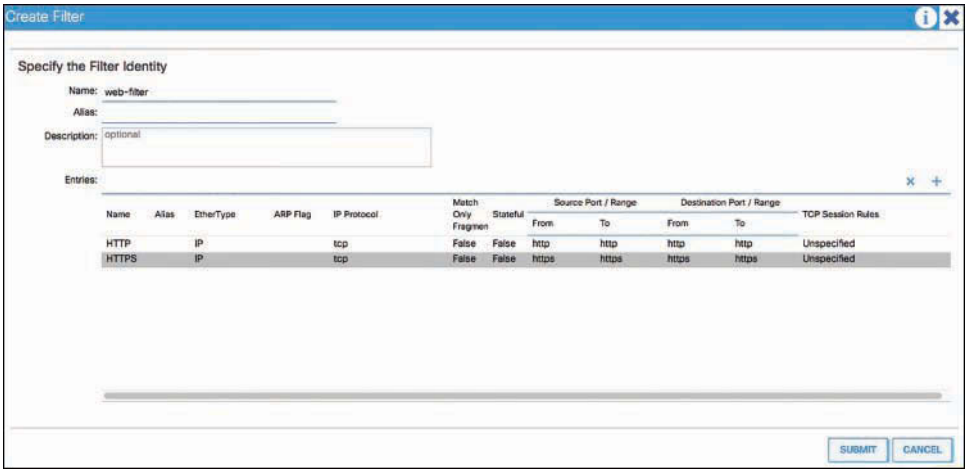

Figure 5-4 *Filters and Filter Entries*

A *subject* is a construct contained within a contract and typically references a filter. For example, the contract Web might contain a subject named Web-Subj that references a filter named Web-Filter.

Concept of Direction in Contracts

As you can see from the previous section, filter rules have a direction, similar to ACLs in a traditional router. ACLs are normally applied to router interfaces. In the case of Cisco ACI, the ACLs (contract filters) differ from classic ACLs in the following ways:

- The interface to which they are applied is the connection line of two EPGs.
- The directions in which filters are applied are "consumer-to-provider" and "provider-to-consumer."
- The ACLs do not include IP addresses because traffic is filtered based on EPG (or source group or class ID, which are synonymous).

Understanding the Bidirectional and Reverse Filter Options

When you create a contract, two options are typically selected by default:

- **Apply Both Directions**
- **Reverse Filter Ports**

The **Reverse Filter Ports** option is available only if the **Apply Both Directions** option is selected (see Figure 5-5).