Cisco Application Control Engine and Cisco Nexus 7000 Series Integration

White Paper

September, 2011

The Cisco® Application Control Engine (ACE) optimizes application delivery in the data center. The Cisco ACE integrated solution with the innovative Cisco NX-OS technology for the Cisco Nexus® 7000 Series Switches further increases the overall value proposition by delivering scalability, availability, and infrastructure virtualization.

What You Will Learn

As the IT landscape rapidly changes, cost reduction pressures, focus on time to market, and employee empowerment are leading enterprises and IT service providers to develop innovative strategies to address these challenges. In the past, environments that needed additional security and isolation required separate sets of equipment and infrastructure. Although this approach provided the required separation, efficiency dropped significantly due to lower utilization, while costs increased due to higher management and capital expenditures. The Cisco multitenant architecture addresses these concerns through a secure, isolated multitenant solution at every network layer.

In addition to multitenant architecture adoption, today's data centers are evolving from a hierarchical architecture model to a flatter Layer 2, geographically distributed model. One of the main factors behind these changes is the rapid adoption of server virtualization technology in data centers.

In a virtualized data center environment, virtual machines can move anywhere in the geographically distributed data center, or they can move to cloud-based data centers hosted by cloud service providers. Geographically dispersed data centers provide added application resiliency and workload allocation flexibility. To support these features, the network must provide Layer 2 connectivity between data centers. Connectivity must be provided without compromising the autonomy of data centers or the stability of the overall network to enable efficient use of the remote computing resources when local resources are scarce.

As more and more organizations begin to chart out plans for accommodating IPv6, many are still having difficulty defining a clear adoption strategy. Now Cisco is offering a complete and comprehensive set of IPv6 solutions to help enterprises and service providers adopt IPv6 in seamless fashion.

Challenge

As enterprise IT environments have dramatically grown in scale, complexity, and diversity of services, they have typically deployed application and customer environments in silos of dedicated infrastructure. These silos are built around specific applications, customer environments, business organizations, and operational requirements. For example:

- Universities need to separate student user services from business operations, student administrative systems, and commercial or sensitive research projects.
- Telcos and service providers must separate billing, customer relationship management (CRM), payment systems, reseller portals, and hosted environments.
- Financial organizations need to securely isolate client records and investment, wholesale, and retail banking services.

To migrate to a multitenant architecture, enterprises and external service providers must be able to securely isolate all customer data, communications, and application environments from other tenants while still delivering the management and flexibility benefits of shared resources. With external providers, the separation must be so complete and secure that the tenants can have no visibility to each other. Enterprises must deliver the secure separation required for their organizational structure, applications, and regulatory compliance.

Businesses face the challenge of providing very high availability for applications while keeping operating expenses low. Applications must be available anytime and anywhere with optimal response times.

The deployment of geographically dispersed data centers allows the IT designer to put in place mechanisms that effectively increase the availability of applications. Geographic dispersion allows flexible mobility of workloads across data centers to avoid demand-hotspots and fully utilize available capacity.

Cisco Solution

Multitenancy is the capability to logically partition a single physical device into many virtual devices. Each virtual device must have all the capabilities of the actual physical device, and each virtual device must be independent and isolated, so that it appears to be a unique physical device from the viewpoint of the network and the network administrator. With virtualization, each virtual device can be allocated its own resources and quality of service (QoS). Each virtual device can also be assigned its own configuration files, management interfaces, and access-control policies - policiesin which access control privileges are assigned to users based on their administrative roles.

To enable all the benefits of geographically dispersed data centers, the network must extend Layer 2 connectivity across diverse locations. LAN extensions may be required at different layers to enable workload mobility between data centers. Existing mechanisms for the extension of Layer 2 connectivity are less than optimal in addressing connectivity and independence requirements and present many challenges and limitations that Cisco Overlay Transport Virtualization (OTV) technology effectively overcomes.

The CiscoApplication Control Engine (ACE) and Cisco Nexus Switch Families offer features tailored to virtual environments, allowing consistent visibility, control, and isolation of the application stack within a multitenant architecture. Cisco ACE integrated with Cisco Nexus OTV technology delivers dynamic workload scaling to enable faster and more efficient deployment of distributed data centers.

Cisco Nexus 7000 Series and Cisco ACE Deployment in Multitenancy and Virtualization Environment

Overview

Cisco Application Control Engine (ACE) and the Cisco Nexus Family of switches offer features tailored to virtual environments, allowing consistent visibility, control, and isolation of applications in a multitenant architecture. The unique multitenancy capabilities of Cisco ACE and the Cisco Nexus 7000 Series enable the customer to accelerate application rollout while reducing consumption and costs through virtualization.

Cisco ACE is the industry-proven, virtualized application delivery solution designed to meet the today's requirements for application delivery. Cisco ACE is a state-of-the-art virtualized load balancer and an application delivery solution that improves application scalability and availability while also improving the utilization of infrastructure resources through offloading and compression technologies.

The Cisco Nexus 7000 Series offers unique virtualization features that enable greater flexibility in network design so that existing or new data center space can be fully utilized. The virtual devices on the Cisco Nexus 7000 Series Switch segregate different service groups inside the same physical switch, consolidating the aggregation layer while preserving the organizational structure of operations and service delivery.

Cisco ACE and Cisco Nexus 7000 Series multitenancy provides scalable, reliable, and cost-effective application delivery and infrastructure services in the virtual data center. These virtualization capabilities provide secure isolation of application environments while delivering the following capabilities:

- **Performance and scale:** Unique virtualization capabilities add new dimensions to application delivery to provide guaranteed resources to applications.
- **Simplification:** Deployment and ongoing maintenance of application services are streamlined through the virtualization capabilities of the Cisco ACE and Cisco Nexus 7000 Series.
- **Flexibility:** Multitenant architecture allows flexibility in overall network design to improve application delivery response time.

Cisco Nexus 7000 Series Virtual Device Context (VDC) Overview

The Cisco Nexus 7000 Series inherits a number of virtualization technologies present in Cisco IOS® Software. From a Layer 2 perspective, virtual LANs (VLAN) virtualize bridge domains in the Nexus 7000 chassis. Virtualization support for Layer 3 is provided through virtual route forwarding (VRF) instances. A VRF can be used to virtualize the Layer 3 forwarding and routing tables.

The virtualization aspect of the Cisco NX-OS Software platform has been extended to support the notion of virtual device contexts (VDCs). A VDC can be used to virtualize the device itself, presenting the physical switch as multiple logical devices (Figure 1). Within that VDC, it can contain its own unique and independent set of VLANs and VRFs. Each VDC can have physical ports assigned to it, thus allowing the hardware data plane to be virtualized as well. Within each VDC, a separate management domain can manage the VDC itself, thus allowing the management plane to also be virtualized.

Figure 1. Cisco Nexus 7000 Series Virtualization: One-to-Many Through VDCs

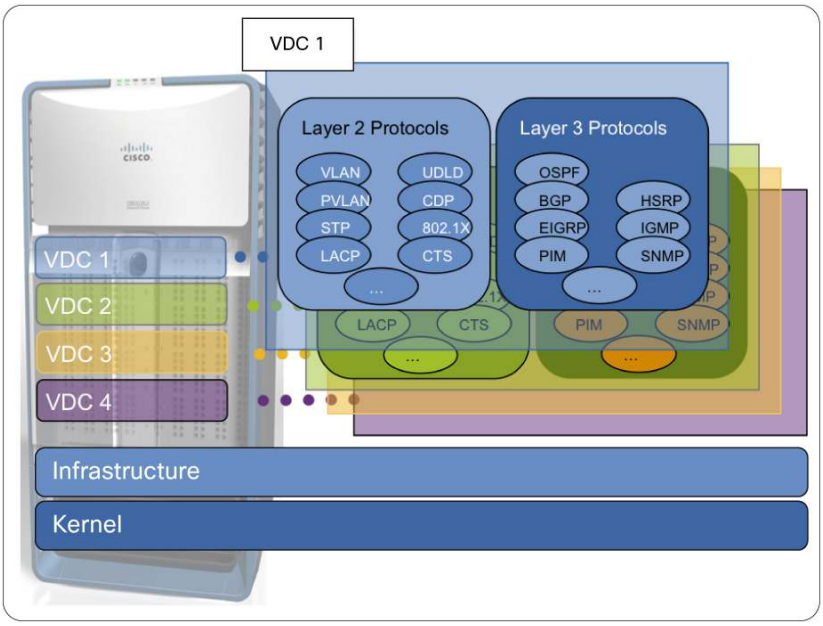

Cisco ACE Virtual Context (VC) Overview

The virtualization environment is divided into objects called virtual contexts (VCs). Each context behaves like an independent Cisco ACE appliance within its own policies, interface, domains, server farms, real servers, and administrators. Each context also has its own management VLAN that you can access using Telnet or Secure Shell (SSH).

As the global administrator (Admin), you can configure and manage all contexts through the Admin context, which contains the basic settings for each virtual device or context. When you log in to the ACE using the console, Telnet, or SSH, you are authenticated in the Admin context.

The Admin context is similar to other contexts. The difference is that when you log in to the Admin context (for example, using SSH), you have full system administrator access to the entire Cisco ACE and all contexts and objects within it (Figure 2). The Admin context provides access to networkwide resources - for example, a syslog server or context configuration server. All global commands for the ACE settings, contexts, resource classes, and so on, are available only in the Admin context.

Each context, including the Admin context, has its own configuration file and local user database that are stored in the local disk partition on the flash disk or that can be downloaded from a File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP), or HTTP(S) server. The startup-config file for each context is stored as the startup configuration file on the flash disk.

Figure 2. Cisco Application Control Engine Virtual Context (VC)

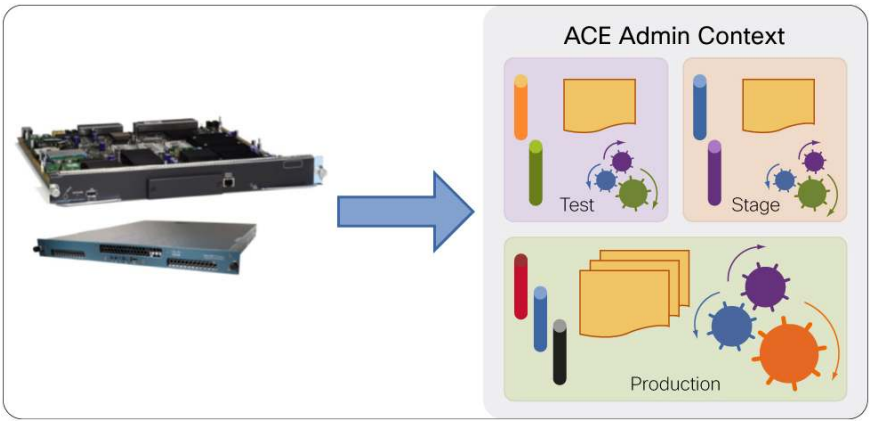

Cisco Nexus 7000 Series VDC Deployment Use Cases with Cisco ACE

Use Case 1: Production and Test/Development Data Center Separation

With Cisco multitenancy, administrators have the flexibility to allocate resources to virtual devices in any way they choose. For example, one administrator may want to allocate one virtual device for the production network and another virtual device for a development network. Cisco ACE and Nexus 7000 Series virtual partitioning provides a novel way of protecting a set of services configured as production virtual devices from accidental mistakes, or from malicious configurations, made in development virtual devices. A configuration failure on a Cisco multitenancy device is limited to the scope of the virtual device in which the configuration was created. A failure in a development virtual device has no effect on a production virtual device, increasing uptime for critical applications.

Figure 3 depicts the physical connectivity of Cisco ACE 4710 Application Control engine and Nexus 7000 Series Switches. Nexus 7000 is configured with Production and Lab VDCs and resources are allocated accordingly. Similarly, ACE 4710 is provisioned with virtual contexts for Production and Lab environments. ACE 4710 has similar resource allocation model as Nexus 7000 which both have ability to allocate resources “as needed” on the fly. Thus, ACE 4710 and Nexus 7000 virtualization technology provide multitenancy in a shared infrastructure.

Figure 3. Physical Connectivity Between the Cisco ACE 4710 Application Control Engine and Cisco Nexus 7000 Series Switches

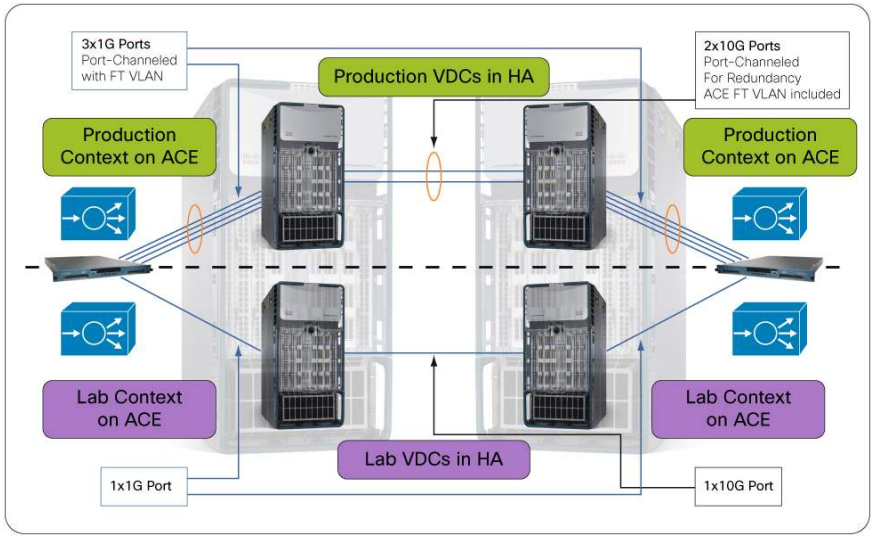

Use Case 2: Segmentation with VDCs and ACE Insertion in Firewalled Security Environments

Users can use the Cisco Nexus 7000 Series VDC and Cisco ACE VC to consolidate multiple devices in a single device. In this use case, typically connectivity to the outside network and inside network is established via separate physical devices, which are interconnected via firewalls. Virtual devices enable network administrators to consolidate these two networks on a single physical device while using logical devices to maintain the security boundary.

In the example shown in Figure 4, the Cisco ACE 4710 Application Control Engine appliance is positioned in the inside network where typically the server farms are deployed protected by perimeter firewall. The ACE 4710 bundles four 1-Gbps interfaces as a port channel to provide maximum aggregate of 4-Gbps throughput connectivity to the inside Cisco Nexus 7000 Series VDC. By creating separate VDCs for logical separation, the network administrator can separate the secure and nonsecure network (that is, the “inside” and “outside” networks) so that they cannot be easily bypassed.

Figure 4. Cisco ACE 4710 Application Control Engine Inserted in Firewalled VDC Environment

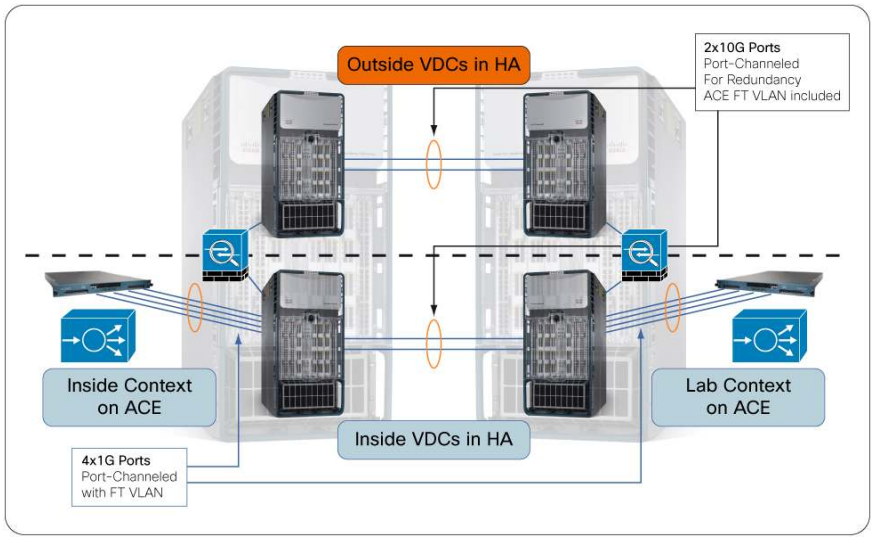

With use of virtual technology of ACE and Cisco Security Firewalls, VLAN, and VRFs, Cisco Nexus 7000 Series VDCs in the data center can support multiple customer environments simultaneously. These virtual services can be added to a given scenario without the need to purchase new hardware, cabling, switching, and so on. A Cisco ACE context can be deployed for server load balancing (SLB), Secure Socket Layer (SSL), or security. Cisco ASA 5500 Series Adaptive Security Appliances can also be used with Cisco ACE for a more robust security feature set.

Use Case 3: Core and Distribution Consolidation Using VDC

A Cisco Nexus 7000 Series VDC can also be used to consolidate physical devices that provide different functional characteristics or roles within a network topology. This type of consolidation provides uses collapsed core and distribution model while using single physical device, which optimizes the use of the interfaces and facility resources such as space, power, and cooling. In the deployment shown in Figure 5, the Cisco ACE 4710 is connected to distribution VDC using the port channel to aggregate client and server traffic and optimize application delivery.

Figure 5. Cisco ACE 4710 Insertion in Core/Distribution Consolidation Design

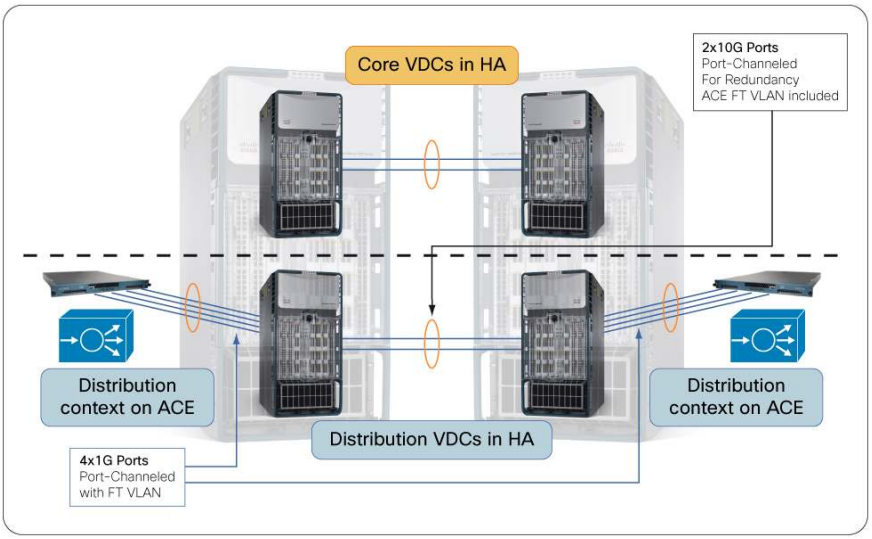

The data center solution must address the high-availability requirements of any device or link failure. It is also an area where more intelligence is required from the network, to perform differentiated services such as application delivery and security for servers and the applications they host. Cisco's data center network design is based on a proven layered approach, which has been tested and improved over the past years in some of the largest data center implementations in the world.

The layers of the data center design are the core, distribution, and access layers. The distribution layer hosts integrated services, primarily the Cisco ACE, to optimize application delivery. The distribution VDC pairs work together to provide redundancy and to maintain the session state while providing valuable services to the access layer, which provides connectivity for the many servers that deliver application and web services.

Cisco Nexus 7000 Series Virtual Device Context Configuration

This section provides an overview of the steps involved in creating a VDC and assigning resources to it.

Initial VDC Setup

The first step in configuring a VDC us to create the VDC. Up to four VDCs can exist in the system at one time. Given that there is always a default VDC active (VDC 1), this means up to three additional VDCs can be created from the command-line interface (CLI).

You create a VDC using the **vdc <name of vdc>** command. Following is the example configuration steps to create VDCs.

```
switch# conf t
switch (config)# vdc production
switch(config-vdc)# show vdc
vdc_id vdc_name state mac
------ -------- ----- ----
1 switch active 00:18:ba:d8:4c:3d
2 production active 00:18:ba:d8:4c:3e
switch(config-vdc)# show vdc detail
vdc id: 1
```

```
vdc name: switch
vdc state: active
vdc mac address: 00:18:ba:d8:4c:3d
vdc ha policy: RESET
vdc id: 2
vdc name: production
vdc state: active
vdc mac address: 00:18:ba:d8:4c:3e
vdc ha policy: BRINGDOWN
```

When the VDC has been created, the system places you into VDC configuration mode where further configuration options can be assigned to the VDC. A default set of configuration statements is assigned to the VDC when it is created, as can be seen in the following output:

```
switch# show run | begin vcd
<snip>
vdc production id 2
template default
hap bringdown
limit-resource vlan minimum 16 maximum 4094
limit-resource span-ssn minimum 0 maximum 2
limit-resource vrf minimum 16 maximum 8192
limit-resource port-channel minimum 0 maximum 256
limit-resource glbp_group minimum 0 maximum 4096
<snip>
```

These configuration statements provide a definition of the resource consumption that the VDC is allowed to work within. These resources include VLAN, VRF, Switched Port Analyzer (SPAN), PortChannels, and Gateway Load Balancing Protocol (GLBP) group IDs. Resource limit assignments can, however, be changed via the command line. An example of how to change a resource limit is shown here:

```
switch(config)# vdc production
switch(config-vdc)# limit-resource vlan minimum 32 maximum 4094
switch(config-vdc)# show run | begin vdc
<snip>
vdc production id 2
template default
hap bringdown
limit-resource vlan minimum 32 maximum 4094
limit-resource span-ssn minimum 0 maximum 2
limit-resource vrf minimum 16 maximum 8192
limit-resource port-channel minimum 0 maximum 256
limit-resource glbp_group minimum 0 maximum 4096
<snip>
```

This example shows how the minimum number of VLANs for the production VDC is changed from 16 to 32. Follow the same steps above to create a VDC for lab network.

After the VDC is created, the administrator is placed into VDC configuration mode. The next task is to assign physical ports to this VDC. Ports on the physical line cards cannot be shared between different VDCs. By default,

all physical ports belong to the default VDC. When a VDC is created, ports can be placed under the control of this VDC using the following CLI option:

```
switch(config)# vdc production
switch(config-vdc)# allocate interface ethernet 1/1-4
switch(config)# vdc lab
switch(config-vdc)# allocate interface ethernet 1/5-8
switch(config-vdc)# show vdc membership
vdc_id: 1 vdc_name: switch interfaces:
<snip>
vdc_id: 2 vdc_name: production interfaces:
Ethernet1/1     Ethernet1/2  Ethernet1/3
Ethernet1/4
vdc_id: 3 vdc_name: lab interfaces:
Ethernet1/5     Ethernet1/6  Ethernet1/7
Ethernet1/8
```

Cisco ACE 4710 Virtualization

Virtualization is a method in allocating available resources into two or more contexts for security and management purposes. Up to 250 contexts can be configured on Cisco ACE. Resources can be allocated to each context to avoid a single context consuming the entire pool of resources. This document covers only basic virtualization configuration. For design implementation examples that use virtual contexts, please refer to the ACE documentation at http://www.cisco.com/en/US/products/ps7027/tsd_products_support_series_home.html.

Context Configuration

The following example shows the configuration steps for configuring a virtual context:

Step 1. Configure resource class(es)

```
ACE/Admin(config)# resource-class Silver <- Resource-class name
<cr> Carriage return.

Following are the different resources that can be segmented:
ACE/Admin(config-resource)# limit-resource ?
   acl-memory          Limit ACL memory
   all                 Limit all resource parameters
   buffer              Set resource-limit for buffers
   conc-connections    Limit concurrent connections (thru-the-box traffic)
   mgmt-connections    Limit management connections (to-the-box traffic)
   proxy-connections   Limit proxy connections
   rate                Set resource-limit as a rate (number per second)
   regexp              Limit amount of regular expression memory
   sticky              Limit number of sticky entries
   xlates              Limit number of Xlate entries
ACE/Admin(config-resource)#

Sample configuration:
ACE/Admin# show running-config resource-class
Generating configuration...
```

```
resource-class Gold
   limit-resource all minimum 0.00 maximum unlimited
   limit-resource conc-connections minimum 10.00 maximum unlimited
   limit-resource sticky minimum 10.00 maximum unlimited
resource-class Silver
   limit-resource all minimum 0.00 maximum unlimited
   limit-resource conc-connections minimum 5.00 maximum unlimited
   limit-resource sticky minimum 5.00 maximum unlimited
resource-class Platinum
   limit-resource all minimum 0.00 maximum unlimited
   limit-resource conc-connections minimum 20.00 maximum unlimited
```

Step 2. Configure context

A context is configured by giving it a name, allocating VLANs, and assigning the context to a resource-class (configured in step 1.):

```
context production
allocate-interface vlan 51
allocate-interface vlan 80
member Silver
```

Step 3. To configure per-context features and functionality, use the **change to** command to change to the context created in Step 2. Once you have changed the context, you are accessing the newly deployed virtual context.

```
ACE/Admin# change to production
ACE/production# conf t

Enter configuration commands, one per line. End with CNTL/Z.

ACE/production(config)# ?
Configure commands:
   aaa Configure aaa functions
   access-group Activate context global access-list
   access-list Configure access control list
   arp Configure ARP
   banner Configure banner message
   class-map Configure a Class map
   crypto Configure CSR parameters
   do EXEC command
   .
   .
```

Cisco ACE 4710 Physical Characteristics

The Cisco ACE 4710 Application Control Engine appliance provides four physical Ethernet ports for processing traffic. The four Layer 2 Ethernet ports can be configured to provide an interface for connecting to 10-Mbps, 100-Mbps, or 1000-Mbps networks. Each Ethernet port supports auto-negotiate, full-duplex, or half-duplex operation on an Ethernet LAN and can carry traffic within one or more designated VLANs.

The Cisco ACE 4710 appliance does not have additional ports used specifically for management traffic. The four Ethernet ports (Figure 6) are used to handle all data and management traffic in and out of ACE. They are also used for ACE appliances deployed in a redundant fashion utilizing a fault-tolerant VLAN to maintain high availability.

Figure 6. Cisco ACE 4710 Appliance Front and Rear Chassis Views

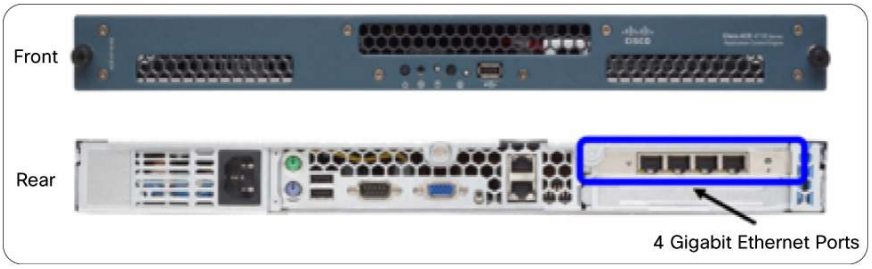

Figure 7 shows the LED link indicators for the Ethernet port and the pin number assignments for the RJ-45 port. As the figure shows for Ethernet port 4, the link LED in the lower-right below each Ethernet port serves as the indicator for the associated port. The ports are numbers from right to left.

Figure 7. Cisco ACE 4710 Appliance: Physical Interfaces Area Labeled from Right to Left

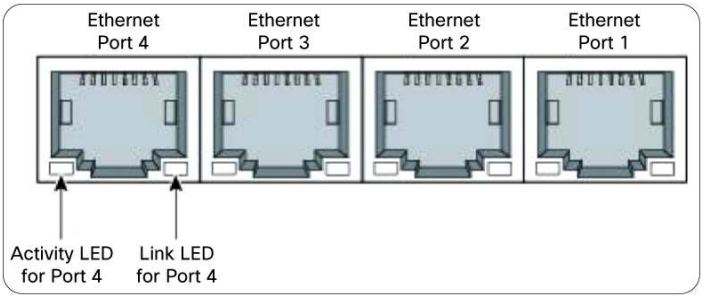

The states of each Ethernet port link LED are as follows:

1. Off when the 10-Mbps Gigabit Ethernet link is connected or when there is no link.
2. Glows steady green when the 100-Mbps Gigabit Ethernet link is connected.
3. Glows steady orange when the 1000-Mbps Gigabit Ethernet link is connected.

The second LED flashes yellow when there is activity.

Connecting the Cisco ACE Appliance to the Network

To maximize application and infrastructure availability, the Cisco ACE 4710 appliance takes advantage of all four Gigabit Ethernet interfaces and ACE virtualization. These interfaces can be configured in a port channel to create a single logical link between the Cisco ACE 4710 appliance and Cisco Nexus 7000 Series Switches. Trunked VLANs can be used to carry all client-server messaging, management traffic, and fault tolerance (FT) communications. For this use case, three interfaces will bundle to become the port channel servicing production network and one interface is dedicated for the lab network (Figure 8).

Note: While the Nexus 7000 Series Switch provides virtual port-channel (vPC) feature for device to create a single logical link to two Cisco Nexus 7000 Series Switches, it is recommended not to leverage the vPC feature for Cisco ACE 4710. The Cisco ACE 4710 is supports up to four 1-Gbps Ethernet interfaces. If interconnecting the

Cisco ACE 4710 to a Cisco Nexus 7000 Series Switch in a vPC deployment, the maximum Cisco ACE 4710 throughput would be limited to 4 Gbps. In order to obtain the maximum bandwidth performance of the Cisco ACE 4710, it is recommend to port-channel all 4 1-Gbps links into a single Cisco Nexus 7000 Series Switch.

Figure 8. Port Channel Carrying Traffic for all VLANs

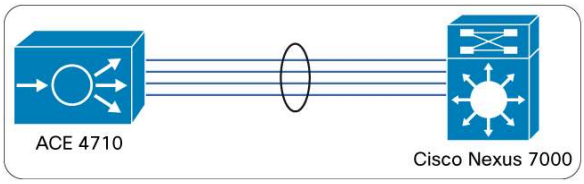

Connecting the Cisco ACE 4710 to a Cisco Nexus 7000 Series Switch in this manner has the following benefits:

- It allows for the creation of a single very high-bandwidth logical link ensuring the highest level (4 Gbps) of throughput possible on the Cisco ACE 4710 appliance. The link gracefully handles asymmetric traffic profiles typical of web architectures.
- It simplifies the interface configuration since the single port channel and 802.1q trunk need to be configured only once and applied to each physical interface.
- Future upgrades - for example, from 1 Gbps to 4 Gbps - can be accomplished in real time by installing a license for increased throughput. This means you don't need to physically re-cable the appliance interfaces.
- Individual ACE contexts are not limited by the throughput of a single 1-Gbps interface. Traffic can be shaped according to the available throughput at the context, virtual IP (VIP) address, or real server level rather than at the interface level.
- It allows the ACE to reach throughput license limits, including the throughput additionally reserved for management traffic. By default, the entry-level ACE appliance has a 1-Gbps through-traffic bandwidth limit and an additional guarantee of 1-Gbps management-traffic bandwidth resulting in a maximum bandwidth of 2 Gbps. Similarly, with the 2-Gbps license, the ACE has a 2-Gbps through-traffic bandwidth and a 1-Gbps management-traffic bandwidth for a total maximum bandwidth of 3 Gbps.
- The port channel provides redundancy if any of the four physical interfaces fail.

The single logical link can support all the common deployment modes, including routed, bridged, one-arm, and asymmetric server return while also addressing high availability and stateful connection replication without a problem. Figure 9 shows the network topology when the Cisco ACE 4710 is connected to the Cisco Nexus 7000 Series.

Figure 9. Network Topology Incorporating the Cisco ACE 4710

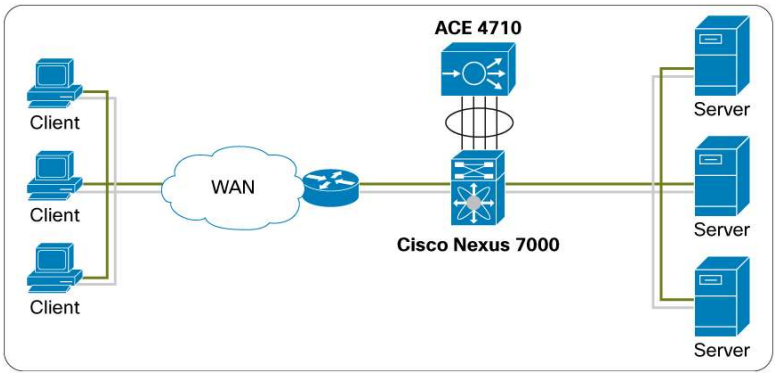

Layer 2 Configuration of the Cisco Nexus 7000 Switch

Once the four physical interfaces on the Cisco ACE 4710 appliance have been physically connected to the Cisco Nexus 7000 Series Switch ports, the first step is to configure the port channel and switch ports on the Cisco Nexus 7000 Series switch.

Switch Port Channel Configuration

In the following example, a Cisco Nexus 7000 Series is configured with a port channel utilizing an 802.1q trunk allowing the associated VLANs. The native VLAN of the trunk is VLAN 10. It is recommended not to use the default VLAN 1 for the native VLAN since this VLAN is used internally on the ACE 4710 appliance.

Port -channel load balancing is used to distribute the traffic load across each of the links in the port channel, helping to ensure efficient utilization of each link. Port-channel load balancing on the Cisco Nexus 7000 Series can use MAC addresses or IP addresses, Layer 4 port numbers, source addresses, destination addresses, or both source and destination addresses. By default, the Cisco ACE 4710 uses **src-dst-mac** to make a load balancing decision. The recommended best practice is to use the source and destination Layer 4 port for the load balancing decision.

```
N7K(config)#port-channel load-balance src-dst l4port
N7K(config)# interface Port-channel1
N7K(config-if)# description ACE 4710
N7K(config-if)# switchport
N7K(config-if)# switchport mode trunk
N7K(config-if)# switchport trunk encapsulation dot1q
N7K(config-if)# switchport trunk native vlan 10
N7K(config-if)# switchport trunk allowed vlan 10,20,30,31, 40,50

N7K(config-if)# do sho run | begin Port
interface Port-channel1
 description to ACE 4710
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport trunk allowed vlan 10,20,30,31,40,50
 switchport mode trunk
```

```
 no ip address
!
```

Once the port channel is configured on the switch, it can be added to the configuration of the four interfaces.

Note: The Cisco ACE 4710 appliance does not support Port Aggregation Protocol (PAgP) or Link Aggregate Control Protocol (LACP). As a result, the port channel is configured using “mode on.”

Switch Interface Configuration

On the Cisco ACE 4710 appliance you can configure the Ethernet port speed for a setting of 10, 100, or 1000 Mbps by using the **speed** command in interface configuration mode. Although the default for the Cisco ACE 4710 appliance is auto-negotiate interface speed, we recommend that you avoid relying on auto negotiation by explicitly configuring the speed to 1000 on both the switch and the appliance. This will avoid the possibility of the interface operating below the expected Gigabit Ethernet speed and help to ensure that the port-channel can reach the maximum 4 Gbps throughput.

The Cisco ACE 4710 does not implement Spanning Tree Protocol and therefore does not take part in spanning-tree root bridge election process. PortFast is configured on the switch to reduce the time required for spanning tree to allow the port connected to the ACE interface to immediately move to forwarding state, bypassing block, listening, and learning states. The average time for a switch port moving into a forward state is approximately 30 seconds. Using PortFast reduces this time to approximately 5 seconds.

Note: In virtual partitions operating in bridge mode, the Cisco ACE 4710 offers an option to bridge spanning-tree Bridge Protocol Data Units (BPDUs) between two VLANs in order to prevent the possibility of a loop. Such a loop may occur when two partitions end up actively forwarding traffic. While this should not happen during normal operation, the option to bridge BPDUs provides a safeguard against this condition. When the switch connected to the ACE 4710 sees BPDUs circling around, the switch will immediately block the port/VLAN the loop originated from. The following ethertype ACL should be configured on ACE and applied to Layer 2 interfaces in bridge mode:

```
access-list BPDU ethertype permit bpdu
```

The following commands are used to configure the switch ports:

```
N7K-1(config-if)# interface ethernet 3/9 - 12
N7K-1(config-if-range)# channel-group 1 force mode on
N7K-1(config-if-range)# spanning-tree port type edge trunk
N7K-1(config-if-range)# spanning-tree trunk native vlan 10
N7K-1(config-if-range)# speed 1000
N7K-1(config-if-range)# no shut
```

The port channel configuration is then added to each of the interfaces, resulting in the following configuration:

```
N7K-1(config-if)# do sho run | beg Ethernet3/9
Building configuration...

interface Ethernet3/9
 description ACE 4710 int1
```

```
 switchport
 switchport trunk native vlan 10
 switchport trunk allowed vlan 10,20,30,31,40,50
 switchport mode trunk
 speed 1000
 spanning-tree port type edge trunk
 channel-group 1 mode on
!
interface GigabitEthernet3/10
 description ACE 4710 int2
 switchport
 switchport trunk native vlan 10
 switchport trunk allowed vlan 10,20,30,31,40,50
 switchport mode trunk
 speed 1000
 spanning-tree port type edge trunk
 channel-group 1 mode on
!
interface GigabitEthernet3/11
 description ACE 4710 int3
 switchport
 switchport trunk allowed vlan 10,20,30,31,40,50
 switchport mode trunk
 speed 1000
 spanning-tree port type edge trunk
 channel-group 1 mode on
!
interface GigabitEthernet3/12
 description ACE 4710 int4
 switchport
 switchport trunk allowed vlan 10,20,30,31,40,50
 switchport mode trunk
 speed 1000
 spanning-tree port type edge trunk
 channel-group 1 mode on
```

Configuring the Cisco ACE 4710 Appliance

Once the switch is configured, the next task is to configure the Cisco ACE 4710 Gigabit Ethernet interfaces and port channel. In this design, we configure the four Ethernet ports as 1000-Mbps full-duplex and associate each of the four ports as a member of a Layer 2 port channel. The port channel bundles the individual physical Ethernet ports into a single logical link associated as an 802.1q trunk.

Cisco ACE Port Channel Configuration

In the following example, the Cisco ACE 4710 is configured with a port channel utilizing an 802.1q trunk, allowing the associated VLANs. Similar to the Cisco Nexus switch configuration, the native VLAN of the trunk is VLAN 10. It is recommended not to use the default VLAN 1 for the native VLAN since this VLAN is used internally on the Cisco ACE 4710 appliance.

Note that the port-channel number on ACE can be different from that of the switch. For example, in an HA configuration, it would be possible to define the distribution switch port-channel 1 for the primary Cisco ACE 4710 appliance and port-channel 2 for the backup. During HA replication, the port channel is replicated to the backup device. This means one of the Cisco ACE 4710 appliances will always have a different port-channel number than that of the switch. Since the port-channel numbers are not required to be consistent between devices, there will be no problem.

```
ACE/Admin(config)# interface port-channel 3
ACE/Admin(config-if)# switchport trunk native vlan 10
   ACE/Admin(config-if)# switchport trunk allowed vlan 10,20,30,31,40,50
   ACE/Admin(config-if)# port-channel load-balance src-dst-port
   ACE/Admin(config-if)# no shutdown
```

Cisco ACE Ethernet Interface Configuration

In the following example, the configuration for the Cisco ACE 4710 appliance is similar to the configuration for the Cisco Nexus 7000 Series Switch. The interface speed on ACE is set to 1000M Full Duplex, and each of the four interfaces is associated with the port channel using the **channel-group** command. We recommend that you configure a carrier delay of 30 seconds for deployments in which ACE is configured with fault tolerance and preemption.

Note: Refer to the section "HA with Preemption and Carrier Delay" at **http://www.cisco.com/en/US/docs/app_ntwk_services/data_center_app_services/ace_appliances/vA5_1_0/configuration/quick/guide/redundancy.html** for more information regarding carrier delay.

Additionally, the Cisco ACE appliance is configured to prioritize incoming HA heartbeat traffic (CoS value of 7 by default) on each of the ports.

Note: Refer to the section "Enabling Quality of Service for High Availability" at http://www.cisco.com/en/US/docs/app_ntwk_services/data_center_app_services/ace_appliances/vA5_1_0/configuration/quick/guide/redundancy.html for more information regarding QoS and HA traffic.

```
ACE/Admin(config)# interface gigabitEthernet 1/1
ACE/Admin(config-if)#  speed 1000M
ACE/Admin(config-if)#  duplex FULL
ACE/Admin(config-if)#  channel-group 3
ACE/Admin(config-if)#  carrier-delay 30
ACE/Admin(config-if)#  qos trust cos
ACE/Admin(config-if)#  no shutdown

ACE/Admin(config)# interface gigabitEthernet 1/2
ACE/Admin(config-if)#  speed 1000M
ACE/Admin(config-if)#  duplex FULL
ACE/Admin(config-if)#  channel-group 3
ACE/Admin(config-if)#  carrier-delay 30
```

```
ACE/Admin(config-if)#  qos trust cos
   ACE/Admin(config-if)#  no shutdown

ACE/Admin(config)# interface gigabitEthernet 1/3
ACE/Admin(config-if)#  speed 1000M
ACE/Admin(config-if)#  duplex FULL
ACE/Admin(config-if)#  channel-group 3
ACE/Admin(config-if)#  carrier-delay 30
ACE/Admin(config-if)#  qos trust cos
ACE/Admin(config-if)#  no shutdown

ACE/Admin(config)# interface gigabitEthernet 1/4
ACE/Admin(config-if)#  speed 1000M
ACE/Admin(config-if)#  duplex FULL
ACE/Admin(config-if)#  channel-group 3
ACE/Admin(config-if)#  carrier-delay 30
ACE/Admin(config-if)#  qos trust cos
   ACE/Admin(config-if)#  no shutdown
```

The port channel configuration is then added to each of the interfaces, resulting in the following configuration:

```
ACE/Admin(config)# do show run int
Generating configuration....
    interface gigabitEthernet 1/1
    speed 1000M
    duplex FULL
    channel-group 3
    carrier-delay 30
        no shutdown
interface gigabitEthernet 1/2
        speed 1000M
        duplex FULL
        channel-group 3
        carrier-delay 30
        no shutdown
interface gigabitEthernet 1/3
        speed 1000M
        duplex FULL
        channel-group 3
        carrier-delay 30
        no shutdown
interface gigabitEthernet 1/4
        speed 1000M
        duplex FULL
        channel-group 3
```

```
carrier-delay 30
no shutdown
```

Verify Layer 2 Network Connectivity

At this point, the port channel and trunk should be up on both the switch and the Cisco ACE 4710 appliance. There are several **show** commands that can verify the port channel and trunk status. For example, to view the configuration status for port-channel interface 3, enter:

```
ACE/Admin# show int port-channel 3
PortChannel 3:
----------------------------
Description:
mode: Trunk     native vlan: 10
status: (UP), load-balance scheme: unknown
PortChannel 3 mapped phyport:1/1 1/2 1/3 1/4
PortChannel 3 mapped active phyport:1/1 1/2 1/3 1/4
PortChannel 3 allow vlan:
 vlan<10>  vlan<20>  vlan<30>-<31>  vlan<40>  vlan<50>
    11606094 packets input, 970121368 bytes, 0 dropped
    Received 694844 broadcasts (10877868 multicasts)
    0 runts, 0 giants
    0 FCS/Align errors, 0 runt FCS, 0 giant FCS
    85431 packets output, 12278955 bytes
    22334 broadcast, 0 multicast, 0 control output packets
    0 underflow, 0 single collision, 0 multiple collision output packets
    0 excessive collision and dropped, 0 Excessive Deferral and dropped
```

It is important to note that the status should indicate UP and that the all four of the interfaces appear in the “mapped” output. Also verify that the mode is Trunk with the correct VLANs associated. Similarly, the status of each physical interface can be verified using the **show interface** command:

```
ACE/Admin# show interface gigabitEthernet 1/4
   GigabitEthernet Port 1/4 is UP, line protocol is UP
   Hardware is ACE Appliance 1000Mb 802.3, address is 00.00.00.00.20.62
MTU 0 bytes
Full-duplex, 1000Mb/s
0 packets input, 0 bytes, 0 dropped
Received 0 broadcasts (0 multicasts)
   0 runts, 0 giants
   0 FCS/Align errors, 0 runt FCS, 0 giant FCS
0 packets output, 0 bytes
0 broadcast, 0 multicast, 0 control output packets
0 underflow, 0 single collision, 0 multiple collision output packets
0 excessive collision and dropped, 0 Excessive Deferral and dropped
```

You can also inspect the interface counters on ACE using the following command:

```
ACE/Admin# show interface gigabitEthernet 1/1 counters
```

On the Cisco Nexus 7000 Series Switch, the following **show** commands can be used to verify the port- channel and interface configuration:

```
N7K-1# show port-channel summary interface port-channel 1
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      NONE      Eth3/9(P)   Eth3/10(P)   Eth3/11(P)
   Eth3/12(P)

N7K-1# show interface trunk

--------------------------------------------------------------------------------
Port          Native  Status        Port
              Vlan                  Channel
--------------------------------------------------------------------------------
Port          Vlans allowed on trunk
Po1           1       trunking      --
--------------------------------------------------------------------------------
Port          Vlans Allowed on Trunk
--------------------------------------------------------------------------------
Po1           10,20,30,31,40,50

N7K-1# show port-channel traffic interface port-channel 1
ChanId      Port Rx-Ucst Tx-Ucst Rx-Mcst Tx-Mcst Rx-Bcst Tx-Bcst
------ --------- ------- ------- ------- ------- ------- -------
     1   Eth3/9    0.31%  42.14%  70.00%   2.79%  99.63%  43.62%
     1   Eth3/10   7.21%  55.11%  30.00%   2.59%   0.18%  10.60%
     1   Eth3/11  92.46%   2.73%    0.0%  94.61%   0.18%  45.76%
     1   Eth3/12  92.46%   2.73%    0.0%  94.61%   0.18%  45.76%
```

Bandwidth Reservation for Management Traffic

By default, the entry-level ACE has a 1-Gbps through-traffic bandwidth and a 1-Gbps management-traffic bandwidth for a total maximum bandwidth of 2 Gbps. However, when the 4-Gbps throughput license is applied, the ACE cannot reserve additional bandwidth beyond the four 1-Gbps port limit. Therefore, some fraction of the total

available bandwidth must be reserved at the context level for management traffic sent to the ACE appliance using the **limit-resource** command. In the following example, a resource class is created and 2 percent of the total 4-Gbps bandwidth is reserved for management traffic in the ACE Admin context:

```
ACE/Admin(config)# resource-class GLOBAL
ACE/Admin(config-resource)# limit-resource rate mgmt-traffic minimum 2 maximum
equal-to-min
ACE/Admin(config)# context Admin
ACE/Admin(config)# resource-class GLOBAL
```

When you allocate a minimum percentage of bandwidth to management traffic, Cisco ACE subtracts that value from the maximum available management traffic bandwidth for all contexts in the ACE. By default, management traffic is guaranteed a minimum bandwidth rate of 0 and a maximum bandwidth rate of 1 Gbps, regardless of the bandwidth license that you install in the ACE. The best practice recommendation is to reserve roughly 100 Mbps for management traffic per context.

HA with Preemption and Carrier Delay

The **carrier-delay** command was introduced in Cisco ACE 4710 Software Release 1.8. This command was added to handle a very specific scenario involving fault-tolerant configurations and preemption. In this scenario, two ACE 4710 appliances are connected to each other through a common LAN switch, such as a Nexus 7000 Series. ACE A is Active while ACE B is Standby. Suppose ACE B takes over because of a failure of ACE A. Moments later, ACE A comes back and wishes to reclaim its active role (it is configured to preempt). When the ACE A comes back up, it brings up its Ethernet interfaces and assumes shortly thereafter that the switch is ready to accept and process traffic. However, this may not be the case due to timing differences. For example, the spanning-tree process may still be determining whether the port can safely be put in the forwarding state on the switch side. In the meantime, ACE A has already sent gratuitous Address Resolution Protocol (ARP) commands to refresh the switch fabric's MAC addresses. To prevent this timing discrepancy, we recommend that you configure a carrier-delay of 30 seconds on the ACE 4710 that is configured to preempt.

Enabling Quality of Service for High Availability

By default, quality of service (QoS) is disabled for each physical Gigabit Ethernet port on the ACE. You can enable QoS for a configured physical Ethernet port that is based on VLAN class of service (CoS) bits (priority bits that segment the traffic in eight different classes of service). If a VLAN header is present, the CoS bits are used by the ACE to map frames into class queues for ingress only. If the frame is untagged, it falls back to a default port QoS level for mapping.

When you enable QoS on a port (a trusted port), ingress traffic is mapped into different ingress queues based on their VLAN CoS bits. If there are no VLAN CoS bits, or QoS is not enabled on the port (untrusted port), the traffic is then mapped into the lowest priority queue.

You can enable QoS for an Ethernet port configured for fault tolerance. In this case, heartbeat packets are always tagged with CoS bits set to 7 (a weight of High). We recommend that you enable QoS on all ports utilizing the FT VLAN to provide a higher priority for incoming FT heartbeats.

ACE Dynamic Workload Scaling (DWS)

Overview

Today's data centers are evolving from a hierarchical architecture model to a flatter, Layer 2-based, geographically distributed model. One of the main factors behind these changes is the rapid adoption of server virtualization technology in data centers. In a virtualized data center environment, virtual machines (VMs)can move anywhere in geographically distributed data centers or they can also move to cloud-based data centers, hosted by cloud service providers. This makes it possible to efficiently use remote compute resources when local resources are scarce.

Cisco ACE, along with Cisco Nexus 7000 Series and VMware vCenter, provides a complete solution for private cloud capacity expansion for data centers. In this solution, Cisco ACE actively monitors the CPU and memory information of the local VMs and computes the average load of the local data center. During normal operations, when the average load is below a preconfigured threshold, Cisco ACE load balances the incoming traffic to only local VMs. However, during peak hours, local VMs may be overloaded and additional capacity may be required to service the incoming requests. When the average load of the local data center crosses a configured threshold, Cisco ACE adds the remote VMs to its load balancing rotation pool, adding more compute resources to service the increased load.

Use Case

Suppose an IT department wants to add additional web servers to an application - for example, a (http://www.bxb.com) when the traffic load exceeds the threshold, but doesn't have any capacity in on-premises data centers. In this example, we'll also assume that the organization has few other remote data centers that have enough computing resources that can be used when on-premise data center is overloaded. Cisco's Overlay Transport Virtualization (OTV) technology is a new data centers interconnect (DCI) technology that can connect multiple geographically distributed data centers over an existing IP network. Using OTV, the IT department can extend its on-premises data center's Layer 2 network to other data centers.

Let's assume that the IT department has interconnected its data centers using OTV. Now IT administrators can clone several VMs from a template and move those VMs using OTV to other data centers. The cloned VMs in remote data centers are an extension of the on-premises infrastructure and are exposed to users only through a load balancer that provides a virtual IP (VIP) address. From the ACE perspective, VMs from both on-premises and remote data centers are just a pool of resources that can be used to provide the capacity to scale an application. However, load balancing to remote VMs can add latency as a result of the roundtrip to remote data centers. For this reason, the IT organization wants to maximize the local computing resources before bursting traffic to remote data centers.

Cisco ACE provides an intelligent solution to this capacity expansion use case, as illustrated in Figure 10. Cisco ACE dynamically determines the locality of the VMs and also monitors the CPU and memory load of the VMs. When the average load of the local VMs exceeds a preconfigured threshold, Cisco ACE takes advantage of remote resources and starts load balancing the incoming client traffic to remote VMs. When the average load comes down, Cisco ACE again starts load balancing the client traffic to only local VMs to maximize the use of local resources.

Example of ACE Virtual Private Cloud Capacity Expansion

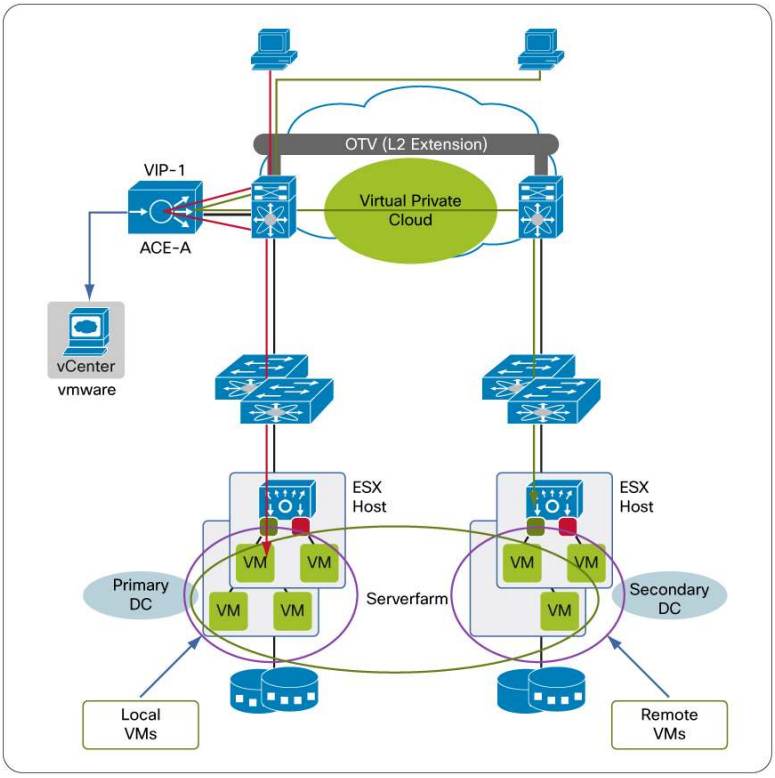

Cisco ACE Configuration for DWS

To configure the virtual private cloud capacity expansion feature, use the following these steps. Figure 11 shows the topology used for this example.

Figure 10. Topology Used for DWS Configuration Example

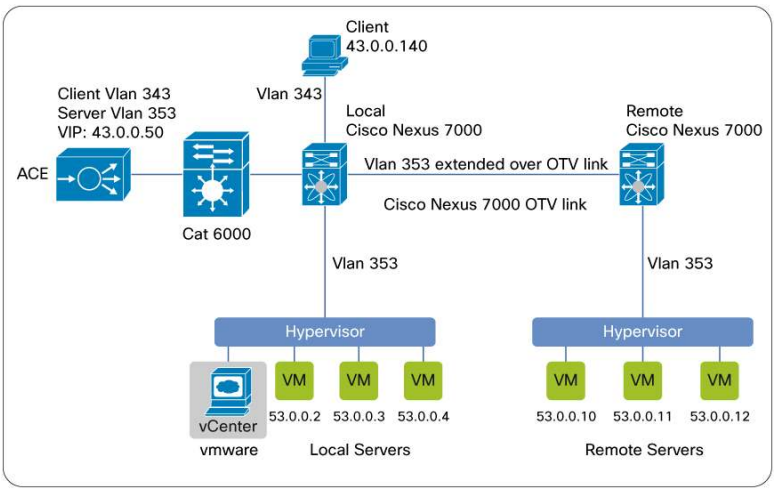

Step 1. Configure the Cisco Nexus 7000 Series Switches and VMware vCenter credentials

Cisco ACE communicates with VMware vCenter using the vCenter web services API. To enable the communication between the Cisco ACE and the vCenter, configure username, password and the URL for the web services API agent of VMware vCenter on the Cisco ACE, as follows:

```
ACE/Admin# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ACE/Admin(config)# vm-controller vcenter
ACE/Admin(config-vm)# credentials Administrator 4Gigabit
ACE/Admin(config-vm)# url https://172.20.33.233/sdk
ACE/Admin(config-vm)#
```

The Cisco ACE module also communicates with Cisco Nexus 7000 Series Switches using SSH protocol. Configure the Cisco Nexus 7000 Series management IP address, username and password on the Cisco ACE to enable the communications between Cisco ACE and Nexus 7000 Series.

```
ACE/Admin# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ACE/Admin(config)# nexus-device n7k
ACE/Admin(config-dci)# ip-address 172.20.33.228
ACE/Admin(config-dci)# credentials admin cisco_1234
ACE/Admin(config-dci)#
```

Step 2. Configure the vm aware probe

Create a new **vmprobe** that queries CPU and memory information from vCenter. Associate the vCenter object that is created in Step1 to the new vm probe. Configure the load parameter and CPU or memory (or both) with the max and min threshold value.

For example, with the following configuration when the average CPU load of the local VMs exceeds 60 (max), Cisco ACE starts load balancing client traffic to remote VMs. When the average load subsides below 40 (min), it stops sending traffic to remote VMs. You can do a similar configuration with memory load as well. When CPU and memory both are configured, Cisco ACE starts load balancing to remote VMs when either the CPU or memory load exceeds the max threshold. ACE stops sending connections to remote VMs when both loads are below the min threshold.

```
ACE/Admin(config)# probe vm vmprobe
ACE/Admin(config-probe-vm)# vm-controller vcenter
ACE/Admin(config-probe-vm)# load cpu burst-threshold max 60 min 40
ACE/Admin(config-probe-vm)#
```

Step 3. Enable the burst to remote datacenters

Create a **serverfarm** with real servers from local and remote data centers. For example, the following serverfarm has three remote real servers (vm10, vm11, and vm12) and three local servers (vm2, vm3, and vm4):

```
serverfarm host sf1
  rserver vm10
    inservice
  rserver vm11
    inservice
  rserver vm12
    inservice
  rserver vm2
    inservice
  rserver vm3
    inservice
```

```
rserver vm4
  inservice
```

Once the serverfarm is created, apply the vmprobe that we created in Step 2 under the serverfarm sf1 object to enable the burst to remote VMs using **dci** command. When DWS burst is configured and the average load exceeds max threshold, Cisco ACE load balances the incoming connections to both local VMs and remote VMs. When the local DWS is configured, Cisco ACE load balances the incoming connections to only local VMs.

```
ACE/Admin# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ACE/Admin(config)# serverfarm host sf1
ACE/Admin(config-sfarm-host)# dws burst probe vmprobe
ACE/Admin(config-sfarm-host)#
ACE/Admin(config-sfarm-host)# exit
ACE/Admin(config)#
```

Related Show Commands

```
ACE/Admin# show probe vmprobe

 probe       : vmprobe
 type        : VM
 state       : ACTIVE
----------------------------------------------
   interval            :  6000    vm-controller    :  vcenter
   cpu-load:
      burst-threshold:
       max threshold :  60      min threshold :  40
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  1
                ------------------ probe results ------------------
   associations ip-address  cpu-load  mem-load  probes  failed   passed   health
   ------------ --------------+-----+--------+--------+--------+--------+------
   serverfarm  : sf1
              aggregate-stats 2     51       1        0        1        SUCCESS

     real      : vm2[0]
              53.0.0.2        2     51       0        0        0        INIT

     real      : vm3[0]
              53.0.0.3        2     50       0        0        0        INIT

     real      : vm4[0]
              53.0.0.4        2     52       0        0        0        INIT

ACE/Admin# show serverfarm sf1
```

```
  Codes: L - local,   R - remote

 serverfarm     : sf1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DCI state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)
 ---------------------------------
                                                 ----------connections-----------
      real                  weight state        current    total      failures
   ---+--------------------+-----+------------+----------+----------+---------
   rserver: vm10
       53.0.0.10:0          8   OPERATIONAL [R] 0          0          0
   rserver: vm11
       53.0.0.11:0          8   OPERATIONAL [R] 0          0          0
   rserver: vm12
       53.0.0.12:0          8   OPERATIONAL [R] 0          0          0
   rserver: vm2
       53.0.0.2:0           8   OPERATIONAL [L] 0          32         0
   rserver: vm3
       53.0.0.3:0           8   OPERATIONAL [L] 0          37         0
   rserver: vm4
       53.0.0.4:0           8   OPERATIONAL [L] 0          31         0


ACE/Admin# show nexus-device n7k

 dci-device       : n7k
-----------------------------------------------
 Total Connections Successful       : 45
 Total Connections Failed           : 0
 Last successful Attempt            : Mon Oct 18 22:56:35 2010
 Last failed Attempt                :
 Last Failure Reason                :

ACE/Admin# show vm-controller vcenter

 vm-controller          : vcenter
 state                  : ACTIVE
-----------------------------------------------
 Total Connections Successful       : 30
 Total Connections Failed           : 0
 Last successful Attempt            :
 Last failed Attempt                :
 Last Failure Reason                :
```

Cisco ACE DWS in OTV Multi-Homing Configuration

OTV Multi-Homing Overview

A critical function built in the OTV protocol is multi-homing, where two (or more) OTV edge devices provide LAN extension services to a given site. The concept of the authoritative edge device (AED) is introduced to provide a more resilient LAN extension solution. The AED has two main tasks:

1. Forwarding Layer 2 traffic (unicast, multicast, and broadcast) between the site and the overlay (and vice versa)
2. Advertising MAC reachability information to the remote edge devices

The AED role is negotiated, on a per-VLAN basis, between all the OTV edge devices belonging to the same site (that is, characterized by the same site ID). To decide what device should be elected as an AED for a given site, the OTV edge devices establish an internal OTV control protocol peering, as shown in Figure 12.

Figure 11. OTV Multi-Homing Deployment with ACE DWS

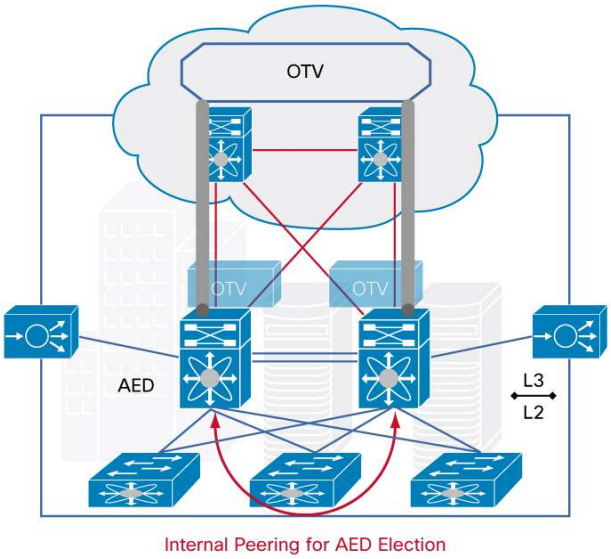

Once two OTV edge devices are configured, they discover each other via the site VLAN and become authoritative for a subset of VLANs. The establishment of internal control plane adjacencies can be verified using the **show otv site** CLI command, as follows:

```
OTV-VDC-A# sh otv site
Site Adjacency Information (Site-VLAN: 15) (* - this device)
Overlay1 Site-Local Adjacencies (Count: 2)
  Hostname                         System-ID         Ordinal
  -------------------------------  ----------------  -------
  OTV-VDC-B                        001b.54c2.3dc2          0
* OTV-VDC-A                        0022.5579.36c2          1
```

OTV elects an AED for each VLAN, which is the only device that can forward the traffic for the extended VLAN inside and outside the data center. Today the extended VLANs are split into odd-numbered and even-numbered VLANs and automatically assigned to the site's OTV edge devices. This results in an edge device being the AED

for the odd VLANs and the other edge device being the AED for the even VLANs. This assignment is not configurable at this time and is done automatically by NX-OS.

The following command can be used to determine the AED selection for the VLANs. In this example, the OTV-VDC-B edge device is the AED for even VLANs:

```
OTV-VDC-B# show otv vlan
OTV Extended VLANs and Edge Device State Information (* - AED)
VLAN   Auth. Edge Device                    Vlan State       Overlay
----   ----------------------------------   ----------       -------
520*   OTV-VDC-B                            active           Overlay1
521    OTV-VDC-A                            inactive(Non AED)Overlay1
```

If the OTV-VDC-B edge device fails, OTV-VDC-A edge device takes over to become the AED for even-numbered VLANs as well as for odd ones.

Cisco ACE and Nexus 7000 Series Configuration for Multi-Homing Deployment

The current version of ACE does not support communicating to multiple Cisco Nexus OTV edge devices, but ACE has MAC reachability information to either odd or even VLANs. Use the following configuration guidelines to overcome this limitation and make ACE deployable in an OTV multi-homing topology. Once the dual-OTV edge devices are configured to become AEDs for subset of VLANs, there are several additional configuration steps required and several caveats:

- The Hot Standby Router Protocol (HSRP) feature is used to provide the HSRP virtual IP address (that is, standby IP address) shared among OTV edge devices in the pair.
- HSRP object tracking is enabled to adjust the priority based on the availability of the OTV join interface. Enable preemption so that HSRP triggers failover when the OTV join interface availability changes.
- Determine either odd or even VLANs to be used for the serverfarm and configure ACE to reflect the same VLAN ID chosen for the server farm.

Cisco Nexus 7000 Series Configuration

The Nexus 7000 Series is configured with a switched virtual interface (SVI) and enables HSRP with object tracking on line-protocol state changes on the OTV join interface. The object tracking on the OTV join interface allows the failover of HSRP when the OTV adjacently fails. For this example, OTV-VDC-B is chosen to be active device for HSRP group 1 and its AED role is for even VLANs.

OTV-VDC-B

```
feature hsrp
track 2 interface Ethernet2/25 line-protocol

interface Vlan510
  no shutdown
  ip address 172.16.110.5/24
  hsrp 1
    preempt
    priority 120
    ip 172.16.110.7
    track 2 decrement 20
```

```
interface Ethernet2/25
  mtu 9216
  ip address 17.1.1.2/24
  ip igmp version 3
  no shutdown

OTV-VDC-B# show otv vlan
OTV Extended VLANs and Edge Device State Information (* - AED)
VLAN   Auth. Edge Device                    Vlan State       Overlay
----   -----------------------------------  ----------       -------
520*   OTV-VDC-B                            active           Overlay1
521    OTV-VDC-A                            inactive(Non AED)Overlay1
```

OTV-VDC-A

```
feature hsrp
interface Vlan510
  no shutdown
  ip address 172.16.110.6/24
  hsrp 1
    preempt
    priority 110
    ip 172.16.110.7
```

Cisco ACE DWS Configuration

Configure the Admin context of ACE to connect to HSRP virtual IP address configured previously on OTV edge devices.

```
ACE/Admin# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ACE/Admin(config)# nexus-device n7k
ACE/Admin(config-dci)# ip-address 172.16.110.7
ACE/Admin(config-dci)# credentials admin cisco_1234
```

See complete DWS configuration in the section “Cisco ACE Configuration for DWS” at http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA5_1_0/configuration/slb/guide/dws.html

Related Syslog Messages

Vmprobe Syslog

The following syslog is generated when the vmprobe cannot connect to or log in to a vCenter:

```
%ACE-3- 256001: Failed to login to vcenter, reason: <%s>.
 %ACE-3- 256002: Failed to retrieve load value for vm <%s>, reason: <%s>.
```

Serverfarm-Related Syslogs

The following syslogs are generated when the average load on local VMs crosses the max or min thresholds:

```
%ACE-5-441004: Average load on local VMs for serverfarm <%s> has crossed above
the max threshold level, starting the loadbalancing on remote VMs.
```

```
%ACE-5-441005: Average load on local VMs for serverfarm <%s> has crossed below
the min threshold level, stopping the loadbalancing on remote VMs.
%ACE-4-441006: Average load on local VMs for serverfarm <%s> has crossed above
the max threshold level, no remote VMs found for remote loadbalancing.
```

Real Server Locality Syslog

The following syslog is generated when the locality of an rserver changes:

```
%ACE-5-442006: Locality state for rserver <%s> has changed to <remote/local.
```

Cisco Nexus Polling Syslog

The following syslog is generated when there are Nexus 7000 Series poling errors:

```
%ACE-2-XXXXXX, N7K polling failed with error: <%s>.
```

MIBs

Serverfarm Related Objects

The following new management information base (MIB) objects are added to the cslbxServerFarmTable in CISO-SLB-EXT-MIB.my table for configuration- and operation-related information:

```
cslbxServerFarmVpcCfgState:
cslbxServerFarmVpcOpState:
```

Real Server Related Objects

The following new MIB object is added to the cesRserver. Table in CISO-ENHANCED-SLB.my table for rserver locality information.

```
cesRserverLocality:
```

VIP Level VPC Info

The following MIB object is added to CISCO-SLB-EXT-MIB.my. This object provides information on whether the feature is enabled or disabled.

```
cslbxVServerVpcCfgState
```

The following MIB object will be added to CISCO-SLB-EXT-MIB.my to represent the VPC operational state at VIP level:

```
cslbxVServerVpcOpState
```

SNMP Traps

Serverfarm Trap

This trap will be generated when the load on the VPC that is enabled crosses the max or min thresholds: cslbxServerFarmVpcOpStateChange

The trap will have the following varbinds:

```
Serverfarm name
cslbxServerFarmVpcOpState
```

Rserver Locality Trap

This trap is generated when the locality of the global rserver changes:

```
cesRserverLocalityChange
```

The trap will have the following varbinds:

```
Rserver name
cesRserverLocality
```

ACE IPv6

Introduction

As the Internet gears up to transition to new IPv6 addressing protocol, enterprises have to adopt IPv6 as well in order to communicate with customers, partners, and suppliers. For a safe transition to IPv6, enterprises must plan carefully for the coexistence of IPv4 and IPv6.

Cisco ACE provide ways to help enterprises make a seamless transition to IPv6. Cisco ACE has following capabilities:

- Dual stack
  - IPv4-to-IPv4 and IPv6-to-IPv6
  - HTTP and DNS inspection support for native IPv6-IPv6 traffic
- Translation
  - SLB64, SLB46 for all the Layer 4 load balancing, which does not need payload modifications or pinholing
  - NAT64, NAT46 for all TCP and User Datagram Protocol (UDP) protocols, which do not need payload modifications or pinholing
  - No DNS64 or DNS46 support on ACE
  - SLB64 and SLB46 support Layer 7 load balancing for HTTP and SSL protocols
- Mixed IPv4 and IPv6 rserver support
- Duplicate address discovery
- Neighbor discovery
- Application awareness
  - HTTP, HTTPS, and DNS

ACE IPv6 Network Deployment Modes

The network deployment mode of ACE can be any of the three traditional modes, which are:

- Routed mode: Here the real servers default gateway is ACE's VLAN interface. Otherwise, return traffic from the real servers would hit ACE's server side VLAN interface.
- Bridged mode: Here ACE bridges the VLANs between the client side and server side. The default gateway of the real servers is the upstream router's address, and in order to reach the upstream router, traffic has to pass through ACE.
- Source-NAT mode: Here ACE is configured with a source-NAT pool and client traffic is load balanced to real servers is source-NAT'ed and hence the return traffic from the real servers is sent back to ACE's Network Address Translation (NAT) pool address.

Cisco ACE IPv6 End-to-End (Dual-Stack) Routed Mode Configuration

In a deployment that uses IPv6 end-to-end configuration routed mode, traffic sent to the IPv6 virtual server will load be balanced to IPv6 real servers, while traffic destined to IPv4 virtual server will be load balanced to IPv4 real servers.

Figure 13 shows a sample network topology for an IPv6 end-to-end implementation along with a sample configuration highlighting the IPv6 part of the configuration. Here the default gateway for real servers is the ACE VLAN's 3001 addresses.

Figure 12. Network Topology for IPv6 End-to-End Configuration

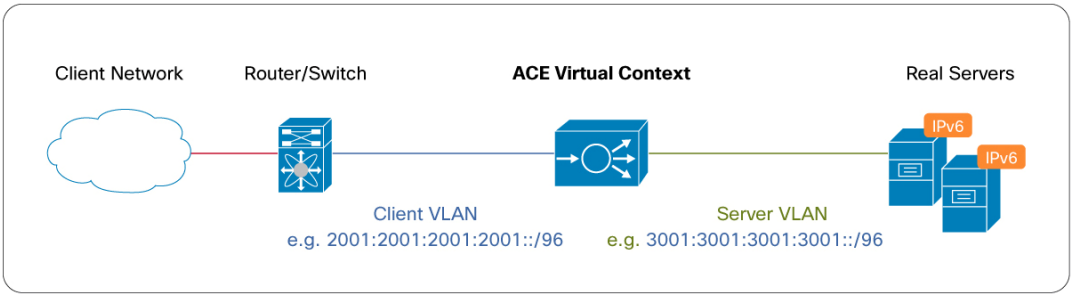


```
access-list all line 8 extended permit ip any any
access-list v6-any line 8 extended permit ip anyv6 anyv6


rserver host v6_rs1
  ip address 3001:3001:3001:3001::3021
  inservice
rserver host v6_rs2
  ip address 3001:3001:3001:3001::3022
  inservice
rserver host v6_rs3
  ip address 3001:3001:3001:3001::3023
  inservice
rserver host v6_rs4
  ip address 3001:3001:3001:3001::3024
  inservice

serverfarm host v6_sf1
  rserver v6_rs1
   inservice
  rserver v6_rs2
   inservice
  rserver v6_rs3
    inservice
  rserver v6_rs4
    inservice
```

```
class-map match-any L4_V6_HTTP-1
  2 match virtual-address 2001:2001:2001:2001::2011 tcp eq www


class-map type management match-all V6-MGMT
  2 match protocol icmpv6 anyv6
class-map type management match-any management
  2 match protocol ssh any
  3 match protocol https any
  4 match protocol icmp any
  5 match protocol http any
  6 match protocol telnet any
  8 match protocol snmp any

policy-map type management first-match MGMT
  class management
    permit
  class V6-MGMT
    permit

policy-map type loadbalance first-match L4_V6_HTTP
  class class-default
    serverfarm v6_sf1

policy-map multi-match V6_Policy1
  class L4_V6_HTTP-1
    loadbalance vip inservice
    loadbalance policy L4_V6_HTTP
    loadbalance vip icmp-reply

service-policy input MGMT

interface vlan 2001
  ipv6 enable
  ip address 2001:2001:2001:2001::2002/96
  ip address 21.1.1.2 255.255.255.0
  access-group input all
  access-group input v6-any
  service-policy input V6_Policy1
  no shutdown
interface vlan 3001
  ipv6 enable
  ip address 3001:3001:3001:3001::3002/96
  no shutdown
```

Cisco ACE IPv6 End-to-End Bridged Mode Configuration

When ACE is deployed in IPv6 end-to-end, one of the network deployment modes can be bridged mode. In this mode, ACE bridges the client and server VLAN. The real server's default gateway is the upstream router, and all the traffic from real servers has to pass through ACE to be routed out.

Figure 14 shows a sample network topology along with sample configuration.

Figure 13. Network Topology for IPv6 End-to-End Bridged Mode

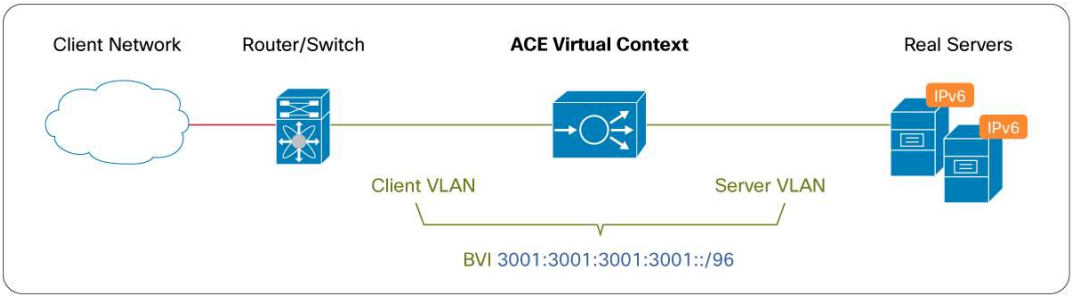


```
access-list 1 line 8 extended permit ip any any
access-list v6-any line 8 extended permit ip anyv6 anyv6



rserver host ipv6-r1
  ip address 2002:2002:2002:2002::3021
  inservice
rserver host ipv6-r2
  ip address 2002:2002:2002:2002::3022
  inservice
rserver host ipv6-r3
  ip address 2002:2002:2002:2002::3023
  inservice
rserver host ipv6-r4
  ip address 2002:2002:2002:2002::3024
  inservice


serverfarm host ipv6-sf1
  rserver ipv6-r1
    inservice
  rserver ipv6-r2
    inservice
  rserver ipv6-r3
    inservice
  rserver ipv6-r4
    inservice
```

```
class-map match-any ipv6-cmap
  3 match virtual-address 2002:2002:2002:2002::2011 tcp any

class-map type management match-all V6-MGMT
  2 match protocol icmpv6 anyv6

class-map type management match-any management
  2 match protocol ssh any
  3 match protocol https any
  4 match protocol icmp any
  5 match protocol http any
  6 match protocol telnet any
  8 match protocol snmp any
  9 match protocol xml-https any

policy-map type management first-match MGMT
  class management
    permit
  class V6-MGMT
    permit

policy-map type loadbalance first-match L7_V6_HTTP
  class class-default
    serverfarm ipv6-sf1

policy-map multi-match V6_Policy1
  class ipv6-cmap
    loadbalance vip inservice
    loadbalance policy L7_V6_HTTP
    loadbalance vip icmp-reply

service-policy input MGMT

interface vlan 2002
  bridge-group 1
  access-group input 1
  access-group input v6-any
  service-policy input V6_Policy1
  no shutdown
interface vlan 3002
  bridge-group 1
  no shutdown

interface bvi 1
  ipv6 enable
  ip address 2002:2002:2002:2002::2002/64
```

```
ip address 22.1.1.2 255.255.0.0
no shutdown
```

Cisco ACE IPv6 SLB-Gateway (Translation) Mode

In IPv6 server load balancing (SLB)-gateway mode, ACE can act as IPv6-to-IPv4 SLB gateway by configuring virtual server with one IP address version (IPv6 or IPv4) and a server farm consisting of real servers with a mix of both IPv4 and IPv6 addresses. Connections that must be translated between IP address versions are translated using the Network Address Translation (NAT) pool address on the appropriate VLAN whose version matches that of the destination address. It is necessary to create specific NAT pool addresses to enable this functionality.

This source address translation makes it possible to use one virtual server to load balance to a server farm containing real servers using both address versions: some members may use IPv4 addresses while others use IPv6 addresses. ACE will load balance as usual to the server farm and translate to the appropriate address when necessary.

Without the NAT pool address of the correct version explicitly configured, the following occurs:

- Connections to an IPv4 virtual server that are forwarded to an IPv6 destination will be translated to the IPv6 NAT pool address of the destination VLAN.
- Connections to an IPv6 virtual server that are forwarded to an IPv4 destination will be translated to the IPv4 NAT pool address of the destination VLAN.
- Connections to a virtual server that are forwarded to a destination of the same IP version will not be translated.

Figure 15 shows a sample network topology along with the configuration.

Figure 14. Network Topology for IPv6 SLB-Gateway (Translation) Mode

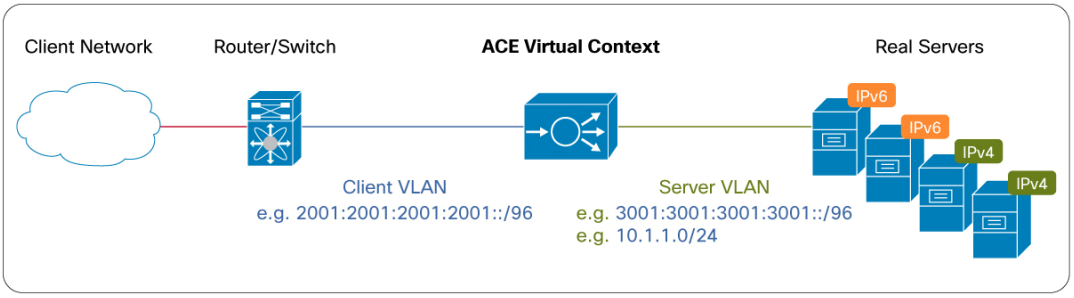


```
access-list all line 8 extended permit ip any any
access-list v6-any line 8 extended permit ip anyv6 anyv6


rserver host rs1
  ip address 10.1.1.21
  inservice
rserver host rs2
  ip address 10.1.1.22
```

```
  inservice

rserver host v6_rs1
  ip address 3001:3001:3001:3001::3021
  inservice
rserver host v6_rs2
  ip address 3001:3001:3001:3001::3022
  inservice

serverfarm host v6_sf1
  rserver rs1
    inservice
  rserver rs2
    inservice
  rserver v6_rs1
    inservice
  rserver v6_rs2
    inservice


class-map match-any L4_V6_HTTP-1
  2 match virtual-address 2001:2001:2001:2001::2011 tcp eq www
class-map type management match-all V6-MGMT
  2 match protocol icmpv6 anyv6
class-map type management match-any management
  2 match protocol ssh any
  3 match protocol https any
  4 match protocol icmp any
  5 match protocol http any
  6 match protocol telnet any
  8 match protocol snmp any

policy-map type management first-match MGMT
  class management
    permit
  class V6-MGMT
    permit

policy-map type loadbalance first-match L4_V6_HTTP
  class class-default
    serverfarm v6_sf1

policy-map multi-match V6_Policy1
  class L4_V6_HTTP-1
    loadbalance vip inservice
    loadbalance policy L4_V6_HTTP
    loadbalance vip icmp-reply
```

```
    nat dynamic 1 vlan 3001

service-policy input MGMT

interface vlan 2001
  ipv6 enable
  ip address 2001:2001:2001:2001::2002/96
  access-group input all
  access-group input v6-any
  service-policy input V6_Policy1
  no shutdown
interface vlan 3001
  ipv6 enable
  ip address 3001:3001:3001:3001::3002/96
  nat-pool 1 10.1.1.100 10.1.1.110 netmask 255.255.255.0 pat
  nat-pool 1 3001:3001:3001:3001::3100 3001:3001:3001:3001::3200/96
  no shutdown
```

Conclusion

As the adoption of IPv6 within the Internet grows, organizations need to create and implement an effective migration strategy. Cisco ACE provides an effective mechanism for operating both IPv4 and IPv6 services simultaneously, offering the ability to move to IPv6 in a phased manner with support of IPv6-to-IPv4 SLB gateway functionality.

For More Information

For more information about the Cisco ACE Application Control Engine Family, visit: http://www.cisco.com/go/ace.

For more information about the Cisco Nexus 7000 Series Switches Family, visit http://www.cisco.com/go/nexus7000.

For more information about Application Networking Services, visit: http://www.cisco.com/en/US/products/hw/contnetw/index.html Or contact your local account representative.

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.

Printed in USA C11-688039-00 11/11