Cisco Virtual Office High-Scalability Design

Contents

Scope of Document

This document describes the recommended headend designs for a high-scale Cisco® Virtual Office deployment. The different components needed along with their corresponding functions are discussed, and sample configurations are provided.

Please refer to the Cisco Virtual Office overview (http://www.cisco.com/go/cvo) for more information about the solution, its architecture, and all the related components.

Introduction

Cisco Virtual Office is a VPN-based solution that consists of remote spoke routers connecting to a centrally located headend infrastructure. Cisco Virtual Office can scale to support many thousands of devices simultaneously. To achieve this high level of scalability, a load balancer must be deployed at the headend that distributes the incoming connections from the different spokes to a group of servers (server farm) connected behind it, as shown in Figure 1. VPN termination and routing are all configured on the servers in the server farm.

Figure 1. High-level load balancing architecture

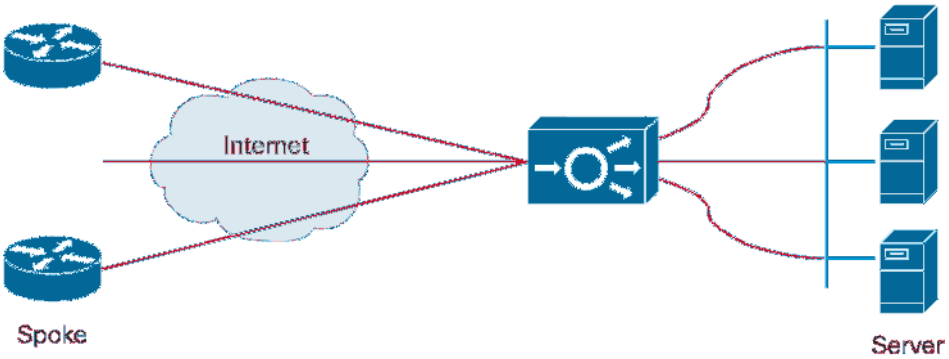

In this guide, three possible high-scale headend designs are presented, differentiated by the level of redundancy that they provide.

The Cisco Application Control Engine (ACE) module plays the role of the load balancer. It can deliver up to 16-Gbps throughput (4 Gbps is the minimum) and is supported on a Cisco Catalyst 6500 or Catalyst 7600 chassis with a Supervisor Engine 720. For more information about the ACE module, please check the ACE module datasheet.

If the Cisco Virtual Office network does not have high-throughput requirements, a standalone Cisco ACE 4710 Application Control Engine appliance can be used as the load balancer instead of a module. The appliance provides the same functions as an ACE module, with a maximum throughput of 4 Gbps (0.5 Gbps is the minimum). For more information about the ACE appliance, please check the ACE 4710 appliance datasheet.

Platforms and Images

Table 1 lists the hardware and software components used in this document. For a complete list of supported and recommended products and images, please refer to the Cisco Virtual Office Hardware and Software list.

Table 1. Hardware and Software Components

	Component	Hardware	Software
Host	Switch	Cisco 7609 router with Supervisor Engine 720 with 10-Gbps Ethernet ports	Cisco IOS® Software Release 12.2(18) SXF
Load balancer	Cisco Application Control Engine (ACE) module	ACE20 (4 Gbps)	Software Release A2 (1.6a)
Server farm	VPN aggregation	ASR 1006 (RP1)	Cisco IOS®-XE Software Release 3.1.0S

Design A

In this design, one ACE module and one host chassis are used. The redundancy is limited to the server farm; one additional real server is used to achieve N + 1 redundancy (the moment one server fails, the remaining servers will be able to handle the resulting load).

Figure 2 shows the basic high-scalability headend architecture of design A.

Figure 2. Headend in design A.

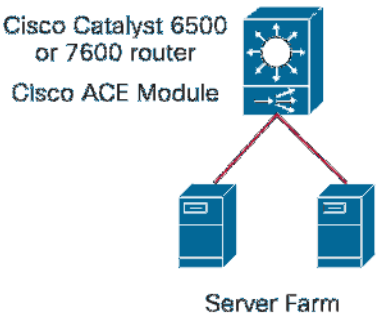

Configuration Tasks

1. Configure the Cisco ACE module.
2. Configure the real servers.

1. Configure the ACE Module

Figure 3 shows the high-level design used in this example. VLAN 9 connects the load balancer to the server farm, and VLAN 99 connects it to the access switches.

Figure 3. Overall design A architecture

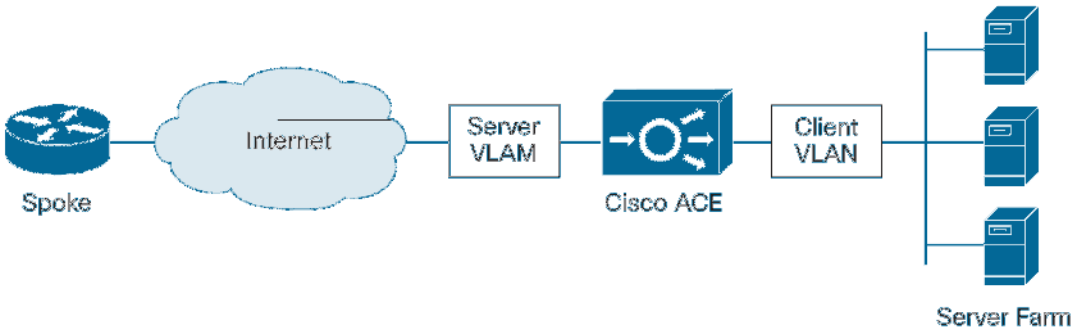

First, the host switch needs to be configured to allow access to and pass the traffic from VLANs 9 and 99 to the ACE module.

```
!! VLAN to the server farm
interface vlan 9
 ip address 192.168.1.5 255.255.255.0
 no shutdown

router eigrp 99
 distribute-list server-farm out
 network 192.168.1.5

ip access-list standard server-farm
 deny any

!! VLAN to the access switches
interface vlan 99
 ip address 192.168.2.1 255.255.255.0
 no shutdown

!! Allow traffic from multiple VLAN interfaces to pass to the ACE !! module
svclc multiple-vlan-interfaces

!! Define the VLAN group that will be passed to the ACE module.
!! VLAN 9 connects to the server farm, and VLAN 99 is the
!! northbound interface

svclc vlan-group 9  9,99

!! Associate the VLAN group with the ACE module
svclc module <ACE-module-number> vlan-group 9
```

After setting up the host, configure the ACE module. To access the ACE module, issue the following command on the host switch:

```
session slot <ACE-module-number> processor 0
```

The ACE is configured in Asymmetric Server Normalization mode, where the return traffic (from the server farm to the spokes) bypasses the ACE and goes straight through the host switch, allowing for using the full throughput available from the ACE module for load balancing incoming flows. The default gateway for the ACE module is the host switch (VLAN 99 in this example). From the perspective of the spoke, the whole server farm appears as a single virtual server whose IP address is the virtual IP address that is configured on the ACE.

```
!! Modify admin user password on the ACE
username admin password 5 <password> role admin domain default-domain

!! Modify www user password on the ACE
```

```
username www password 5 <password> role admin domain default-domain

!! Configure a ping probe that is used to track the availability !! of the real
servers in the server farm

probe icmp ping
 faildetect 3
 interval 10
 passdetect interval 30
 passdetect count 2

!! Define the real servers
rserver host rs1
 ip address 192.168.1.2
 inservice
rserver host rs2
 ip address 192.168.1.3
 inservice

!! Configure the server farm that groups the real servers. The
!! load balancing algorithm used in this example is Least
!! Connections. The conn-limit should be equal to at least four
!! times the number tunnels that the realserver supports
serverfarm host sf
 transparent
 failaction purge
 predictor leastconns
 probe ping
 rserver rs1
  conn-limit max 4000 min 4000
  inservice
 rserver rs2
  conn-limit max 4000 min 4000
  inservice
!! Allocate resources (minimum is 10% in this example) to
!! stickyness on the ACE. By default, the ACE does not spare any !! resources for
stickyness
resource-class resources
 limit-resource all minimum 0.00 maximum unlimited
 limit-resource sticky minimum 10.00 maximum equal-to-min

!! Configure the Admin context as a member of the customized
!! resource group
context Admin
 member resources

!! Configure a stickyness group with a 10 minute timeout.
```

```
!! Stickyness ensures that connections coming from the same source
!! are sent to the same real server
sticky ip-netmask 255.255.255.255 address source sticky-group
 timeout 10
 serverfarm sf
 replicate sticky

!! Set the UDP timeout value to match the sticky timeout (10
!! minutes). This makes sure that the ACE doesn’t time out the UDP
!! connections coming from the spokes
parameter-map type connection udp_timeout
 set timeout inactivity 600

!! Configure a class map to match ISAKMP, ESP, and NAT-T traffic
!! going to the virtual IP address. Add any additional ports as
!! desired
class-map match-any lb-cm
 match virtual-address <virtual-IP> udp eq 500
 match virtual-address <virtual-IP> udp eq 4500
 match virtual-address <virtual-IP> 50

!! Configure the load balancing policy
policy-map type loadbalance first-match lb-pm-fm
 class class-default
  sticky-serverfarm sticky-group

!! Tie in the load balancing policy with the virtual IP address
!! class-map, and the configured udp_timeout parameter-map. The
!! ACE is also configured to reply to ping requests coming to the
!! virtual IP address
policy-map multi-match lb-pm
 class lb-cm
  loadbalance vip inservice
  loadbalance policy lb-pm-fm
  loadbalance vip icmp-reply
  connection advanced-options udp_timeout
 class class-default
!! Configure an access-list to control traffic destined to the
!! virtual IP address. Only the desired traffic (ISAKMP, ESP,
!! ICMP, and NAT-T) is permitted here
access-list dmvpn line 1 extended permit udp any host <vitual-IP> eq isakmp
access-list dmvpn line 2 extended permit udp any host <virtual-IP> eq 4500
access-list dmvpn line 3 extended permit esp any host <virtual-IP>
access-list dmvpn line 4 extended permit icmp any host <virtual-IP>

!! Configure a management class-map to match the desired
!! management protocols. This includes the protocols used to
```

```
!! access the ACE
class-map type management match-any remote-access-cm
 match protocol icmp any
 match protocol ssh any
 match protocol http any
 match protocol telnet any
 match protocol https any

!! Configure the ACE's management policy, permitting the protocols
!! defined in the management class-map and blocking all other
!! traffic
policy-map type management first-match remote-access-pm
 class remote-access-cm
  permit

!! Configure the ACE's VLAN interfaces and apply the ACL, and the !! management
and load balancing policies.

interface vlan 9
  description "connection to server farm"
  ip address 192.168.1.1 255.255.255.0
  no icmp-guard
  service-policy input remote-access-pm

interface vlan 99
  description "connection to Host Switch"
  ip address 192.168.2.2 255.255.255.0
  no normalization
  no icmp-guard
  access-group input dmvpn
  service-policy input lb-pm
  service-policy input remote-access-pm
  no shutdown

!! Configure the default gateway (VLAN 99 on the Host Switch)
ip route 0.0.0.0 0.0.0.0 192.168.2.1
```

Note: Note: When one of the real servers in the server farm is not a Cisco Aggregation Services Router (ASR) 1000 (e.g., the Cisco 3945E Integrated Services Router G2), the Secure Device Provisioning (SDP) function can be configured on that server, along with the VPN configuration. In that case, an additional server farm is configured on the Cisco ACE module to handle SDP traffic (HTTP and HTTPS), with that server only as a member real server. For the sake of illustration, assume that server rs2 is not a Cisco ASR 1000 router, and that TCP port 8000 is used for HTTP. The maximum and minimum connection limits are limited to twice the maximum number of simultaneous HTTP connections that Cisco IOS Software can handle (15).

```
!! Configure the server farm for SDP. Only rs2 is a member
```

```
 serverfarm host sdp
 transparent
 failaction purge
 predictor leastconns
 probe ping
 rserver rs2
  conn-limit max 30 min 30
  inservice

!! Configure a class map to match HTTP(TCP 8000) and HTTPS traffic
!! going to the virtual IP address
class-map match-any sdp-cm
 match virtual-address <virtual-IP> tcp eq 8000
 match virtual-address <virtual-IP> tcp eq https

!! Configure the load balancing policy
policy-map type loadbalance first-match sdp-fm
 class class-default
  serverfarm sdp

!! Add the rules for class-map sdp-cm under the load balancing
!! policy-map
policy-map multi-match lb-pm
 class sdp-cm
  loadbalance vip inservice
  loadbalance policy sdp-fm
  loadbalance vip icmp-reply
  connection advanced-options udp_timeout
```

2. Configure the Real Servers

Each real server in the server farm must be configured with an interface that connects to the host switch. The IP address configured on that interface must be on the same subnet as the server-farm VLAN that was configured on the host switch (VLAN 9 in the example). The default gateway of the real servers is the host switch (VLAN 9). Each real server is also configured with a loopback interface that has as its IP address the virtual IP address of the server farm. The DMVPN tunnel to the spokes on the real server is sourced from that loopback interface. All the tunnel interfaces on the real servers must also share the same IP address, which is configured as the next-hop server on the DMVPN spokes.

Following is a sample of the relevant DMVPN configuration for one of the real servers. For a full DMVPN hub sample configuration, please refer to the Cisco Virtual Office Converged VPN Deployment Guide at http://www.cisco.com/go/cvo.

```
!! Configure the interface connecting to the Host Switch
interface <interface-name>
 ip address 192.168.1.2 255.255.255.0
 no shutdown

!! Configure the default gateway (VLAN 9 on the Host Switch)
ip route 0.0.0.0 0.0.0.0 192.168.1.5

!! Configure the loopback interface with the virtual IP address
interface loopback 0
 ip address <virtual-IP> 255.255.255.255

!! Configure the DMVPN tunnel interface. The tunnel subnet used in
!! this example is 172.16.0.0/24. 172.16.0.1 is the common tunnel
!! IP address
interface Tunnel 0
 bandwidth 2000
 ip address 172.16.0.1 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim dr-priority 10
 ip pim nbma-mode
 ip pim sparse-dense-mode
 ip nhrp map multicast dynamic
 ip nhrp network-id 78600
 ip nhrp holdtime 300
 ip nhrp server-only
 ip nhrp shortcut
 ip nhrp redirect
 no ip split-horizon eigrp 99
 ip tcp adjust-mss 1360
 delay 1500
 qos pre-classify
 tunnel source Loopback0
```

```
 tunnel mode gre multipoint
 tunnel key 786
 tunnel protection ipsec profile dmvpn

!! Configure the routing protocol on the required interfaces.
!! Since all the real servers in the server farm share the same
!! address on the loopback interface, the router ID is manually
!! configured on each one of them
router eigrp 99
 no auto-summary
 network 172.16.0.1 0.0.0.0
 network 192.168.1.2 0.0.0.0
 eigrp router-id 1.1.1.1
```

In order to share the routes learned from the spokes between the different real servers, an additional mGRE tunnel interface that connects all the real servers to each other is configured. An iBGP routing protocol instance (bgp 99) is built over this tunnel, and the EIGRP routes that were learned from the spokes are redistributed into the iBGP instance, allowing every real server in the server farm to learn about all the spokes that are on the DMVPN cloud and be able to reach them. The mGRE tunnel interfaces (hub-to-hub and hub-to-spokes) are on different subnets, and they have the same NHRP network-id, enabling the building of direct spoke-to-spoke tunnels (a redirect message is sent to the originating spoke **if ip nhrp redirect** is configured on the incoming tunnel interface (hub-to-spoke tunnel), **and** if the incoming and outgoing tunnel interfaces have the same network-id). As a result of using the same network-id, no tunnel keys can be configured, and thus both tunnels must be sourced from different physical interfaces.

In this example, tunnel 1 is the hub-to-hub tunnel interface, and it is sourced from the Cisco EtherChannel connected to the VSS. (Tunnel 0 is the spoke-to-hub tunnel interface, and its configuration is repeated for the sake of clarity.)

```
!! Spoke-to-hub DMVPN tunnel interface. The tunnel subnet used in
!! this example is 172.16.0.0/24. 172.16.0.1 is the common tunnel
!! IP address. No tunnel key is configured
interface Tunnel 0
 bandwidth 2000
 ip address 172.16.0.1 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim dr-priority 10
 ip pim nbma-mode
 ip pim sparse-dense-mode
 ip nhrp map multicast dynamic
 ip nhrp network-id 78600
 ip nhrp holdtime 300
 ip nhrp server-only
 ip nhrp shortcut
 ip nhrp redirect
 no ip split-horizon eigrp 99
```

```
 ip tcp adjust-mss 1360
 delay 1500
 qos pre-classify
 tunnel source Loopback0
 tunnel mode gre multipoint
 tunnel protection ipsec profile dmvpn

!! Hub-to-hub DMVPN tunnel interface. The tunnel subnet used in
!! this example is 172.20.0.0/24. No tunnel key is configured. An !! ‘ip nhrp map’, ‘ip nhrp map multicast’, and ‘ip nhrp nhs’
!! commands are needed for each one of the other servers in the
!! server farm
interface Tunnel 1
 bandwidth 2000
 ip address 172.20.0.1 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim dr-priority 10
 ip pim nbma-mode
 ip pim sparse-dense-mode
 ip nhrp map 172.20.0.2 192.168.1.3
 ip nhrp map multicast 192.168.1.3
 ip nhrp network-id 78600
 ip nhrp holdtime 300
 ip nhrp nhs 172.20.0.2
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 delay 1500
 qos pre-classify
 tunnel source <interface-to-Host-Switch>
 tunnel mode gre multipoint

!! Configure the routing protocols on the required interfaces.
!! Neighbor-related commands under the bgp configuration are added !! for each one of the other real servers. EIGRP 99 is also
!! redistributed into BGP 99, controlled by a route-map that only !! allows the routes that were learned from the spokes
router eigrp 99
 no auto-summary
 network 172.16.0.0 0.0.0.255
 network 192.168.1.0 0.0.0.255
 eigrp router-id 1.1.1.1

router bgp 99
 no auto-summary
 no synchronization
 neighbor 172.20.0.2 remote-as 99
```

```
 neighbor 172.20.0.2 update-source tunnel 1
 neighbor 172.20.0.2 next-hop-self
 redistribute eigrp 99 route-map spokes

route-map spokes permit 10
 match ip address spokes-acl

!! Add as many permit statements as necessary, to cover all spokes
ip access-list standard spokes-acl
 permit <spokes-subnets>
```

One of the advantages that Cisco Virtual Office brings is the ability to converge multiple Cisco IOS Software VPN technologies on the same headend aggregation routers. As a result, the same headend infrastructure can be used for multiple solutions, thus increasing the ROI and reducing the CapEx.

For example, Easy VPN tunnels can also be terminated on the server farm, where the ACE module uses the same policy to load balance both Easy VPN and DMVPN connections (ISAKMP and IPsec traffic). The real servers are configured with an Easy VPN policy, in addition to DMVPN. Following is a sample of the relevant Easy VPN configuration for a real server. For full Easy VPN hub-and-spoke sample configurations, please refer to the Cisco Virtual Office Express Deployment Guide at http://www.cisco.com/go/cvo.

```
crypto isakmp client configuration group vpn-group
 dns <DNS-server>
 domain <domain-name>
 pool client_pool

crypto isakmp profile isa-prof
 ca trust-point pki-tp
 match identity group vpn-group
 isakmp authorization list default
 client configuration address respond
 virtual-template 1


!! The EzVPN tunnel shares the source interface (loopback0) and
!! IPSec profile (dmvpn) with the DMVPN tunnel
interface Virtual-Template1 type tunnel
 description Corporate data-traffic Virtual-Tunnel
 ip unnumbered loopback0
 ip virtual-reassembly
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile dmvpn shared

ip local pool client_pool 172.30.0.1 172.30.15.254
```

Design B

In this design, two ACE modules and one host chassis are used. In addition to the server-farm redundancy from design A, this design provides ACE module redundancy. Two ACE modules are used in Fault Tolerance (FT) mode.

Figure 4 shows the high-scalability headend architecture of design B.

Figure 4. Headend in Design B

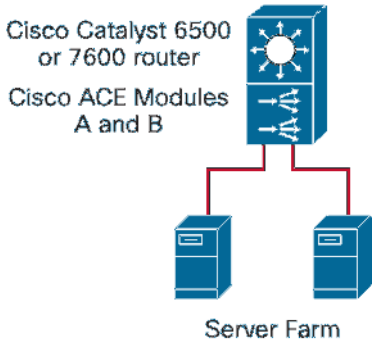

A dedicated VLAN (VLAN 100 in this example) is used as a Fault Tolerance (FT) VLAN to set up the redundancy between the ACE modules. The module with the highest priority takes over as the active module, and the second one becomes the backup. The configuration on the active module, and any changes thereafter, are automatically replicated over the FT VLAN on the standby module. As a result, the configuration mode becomes disabled on the standby module. Heartbeats are also sent between the two modules over the FT VLAN, thus allowing each module to keep a tag on the availability of the other module. This scenario allows for an instant failover when the active module fails. For more information about ACE redundancy, please check the Redundant ACE modules configuration guide.

FT is configured first on both modules. The remainder of the configuration is entered only on the active module, and it gets automatically replicated to the standby module.

FT VLAN 100 is configured on the host switch and added to the group of VLANs that are allowed to reach each of the ACE modules.

```
vlan 100
 state active

svclc multiple-vlan-interfaces
svclc vlan-group 9  9,99,100
svclc module <ACE-module-A-number> vlan-group 9
svclc module <ACE-module-B-number> vlan-group 9
```

In this example, ACE module A is configured as the active module and B as the standby.

ACE module A:

```
!! Define the FT VLAN and the other module as a peer
ft interface vlan 100
```

```
 ip address 192.168.100.1 255.255.255.0
 peer ip address 192.168.100.2 255.255.255.0
 no shutdown

!! Configure the FT peer and the corresponding heartbeat
!! parameters
ft peer 1
 heartbeat interval 300
 heartbeat count 10
 ft-interface vlan 100
 query-interface vlan 99

!! Configure the FT group. ACE module A is configured with a
!! higher priority (20) and so it will become the active module.
!! The Admin context is used on the ACE
ft group 1
 peer 1
 priority 20
 peer priority 2
 associate-context Admin
 inservice
```

ACE module B:

```
!! Define the FT VLAN and the other module as a peer
ft interface vlan 100
 ip address 192.168.100.2 255.255.255.0
 peer ip address 192.168.100.1 255.255.255.0
 no shutdown

!! Configure the FT peer and the corresponding heartbeat
!! parameters
ft peer 1
 heartbeat interval 300
 heartbeat count 10
 ft-interface vlan 100
 query-interface vlan 99

!! Configure the FT group. ACE module B is configured with a lower !! priority
(2) and so it will become the standby module. The
!! Admin context is used on the ACE
ft group 1
 peer 1
 priority 2
 peer priority 20
 associate-context Admin
```

```
 inservice
```

At this point, the configuration mode on ACE module B (standby) will be locked, and any configuration update on ACE module A will automatically be replicated on B.

The remainder of the ACE configuration is the same as in design A, except for the server- and client-side VLAN interfaces. On module A:

```
!! Configure the ACE's VLAN interfaces and apply the ACL, and the !! management
and load balancing policies. An alias IP address
!! must be used. The servers will see both ACE modules as one
!! virtual ACE module with the alias IP as its address. The peer
!! IP address is the one that will be used on VLAN 9 on the
!! standby ACE module
interface vlan 9
  description "connection to server farm"
  ip address 192.168.1.21 255.255.255.0
  alias 192.168.1.1 255.255.255.0
  peer ip address 192.168.1.22 255.255.255.0
  no icmp-guard
  service-policy input remote-access-pm
  no shutdown

interface vlan 99
  description "connection to Host Switch"
  ip address 192.168.2.21 255.255.255.0
  alias 192.168.2.2 255.255.255.0
  peer ip address 192.168.2.22 255.255.255.0
  no normalization
  no icmp-guard
  access-group input dmvpn
  service-policy input lb-pm
  service-policy input remote-access-pm
  no shutdown
```

The configuration for the servers in the server farm is also the same as for design A.

Design C

In this design, two ACE modules and two host chassis are used. In addition to the redundancy from design B, this design provides host-switch redundancy. Two host switches are used in Virtual Switching System (VSS) mode.

Figure 5 shows the high-scalability headend architecture of design B.

Figure 5. High-Scalability Headend Architecture of Design B

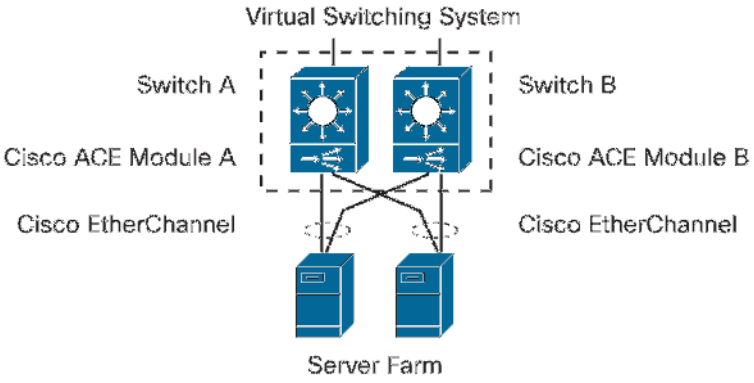

VSS consists of connecting the two switches together using a Virtual Switch Link (VSL), making them appear as one virtual switch. The switch with the higher configured priority becomes active, and the other switch becomes standby. The traffic of the control and management planes is handled by the active switch in the VSS, while the data-plane traffic (packet forwarding) is handled by both switches.

Over the VSL, the active supervisor engine passes the current configuration (including any changes thereafter), state, and forwarding information to the standby supervisor. As a result, the configuration mode becomes disabled on the standby engine. In addition, Stateful Switchover (SSO), Nonstop Forwarding (NSF), and Cisco Express Forwarding are configured on both switches to allow the standby switch to detect any failure on the active switch and then take over immediately, thus minimizing the downtime on the network.

The real servers in the farm connect to the VSS using Multichassis EtherChannel (MEC) links. This link is similar to the regular Cisco EtherChannel link, except that it allows connecting the real servers to both switches, which are part of the same VSS, at the same time. From the perspective of the real servers, this link is treated as a regular Cisco EtherChannel link.

For more information about configuring VSS, please check the VSS configuration guide.

The configurations of the ACE module and server farm are similar to the ones in design B, except for configuring an EtherChannel on each real server in the server farm to connect to the VSS. The EtherChannel must have at least two physical ports: one connecting to each switch in the VSS, as follows:

```
interface port-channel 1
 ip address 192.168.1.2 255.255.255.0
 no shutdown
!! Place each of the physical interfaces connected to the VSS on
!! the EtherChannel
interface <interface-name>
 channel-group 1 mode active
 no shutdown
```

References

1. Cisco ACE module data sheet: http://www.cisco.com/en/US/prod/collateral/modules/ps2706/ps6906/product_data_sheet0900aecd8045861b.html
2. Cisco ACE 4710 appliance data sheet: http://www.ict-partner.net/en/US/prod/collateral/contnetw/ps5719/ps7027/Data_Sheet_Cisco_ACE_4710.html
3. VSS configuration guide: http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/vss.html
4. Redundant Cisco ACE modules configuration guide: http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA2_3_0/configuration/getting/started/guide/redundancy.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C07-684951-00 08/11