Newer Cisco SBA Guides Available

This guide is part of an older series of Cisco Smart Business Architecture designs. To access the latest Cisco SBA Guides, go to http://www.cisco.com/go/sba

Cisco strives to update and enhance SBA guides on a regular basis. As we develop a new series of SBA guides, we test them together, as a complete system. To ensure the mutual compatibility of designs in Cisco SBA guides, you should use guides that belong to the same series.

SBA

DATA CENTER

DEPLOYMENT GUIDE

Data Center Deployment Guide

SMART BUSINESS ARCHITECTURE

August 2012 Series

Preface

Who Should Read This Guide

This Cisco® Smart Business Architecture (SBA) guide is for people who fill a variety of roles:

- Systems engineers who need standard procedures for implementing solutions
- Project managers who create statements of work for Cisco SBA implementations
- Sales partners who sell new technology or who create implementation documentation
- Trainers who need material for classroom instruction or on-the-job training

In general, you can also use Cisco SBA guides to improve consistency among engineers and deployments, as well as to improve scoping and costing of deployment jobs.

Release Series

Cisco strives to update and enhance SBA guides on a regular basis. As we develop a series of SBA guides, we test them together, as a complete system. To ensure the mutual compatibility of designs in Cisco SBA guides, you should use guides that belong to the same series.

The Release Notes for a series provides a summary of additions and changes made in the series.

All Cisco SBA guides include the series name on the cover and at the bottom left of each page. We name the series for the month and year that we release them, as follows:

month year Series

For example, the series of guides that we released in August 2012 are the "August 2012 Series".

You can find the most recent series of SBA guides at the following sites:

Customer access: http://www.cisco.com/go/sba

Partner access: http://www.cisco.com/go/sbachannel

How to Read Commands

Many Cisco SBA guides provide specific details about how to configure Cisco network devices that run Cisco IOS, Cisco NX-OS, or other operating systems that you configure at a command-line interface (CLI). This section describes the conventions used to specify commands that you must enter.

Commands to enter at a CLI appear as follows:

```
configure terminal
```

Commands that specify a value for a variable appear as follows:

```
ntp server 10.10.48.17
```

Commands with variables that you must define appear as follows:

```
class-map [highest class name]
```

Commands shown in an interactive example, such as a script or when the command prompt is included, appear as follows:

```
Router# enable
```

Long commands that line wrap are underlined. Enter them as one command:

```
wrr-queue random-detect max-threshold 1 100 100 100 100 100 
100 100 100
```

Noteworthy parts of system output or device configuration files appear highlighted, as follows:

```
interface Vlan64
  ip address 10.5.204.5 255.255.255.0
```

Comments and Questions

If you would like to comment on a guide or ask questions, please use the SBA feedback form.

If you would like to be notified when new comments are posted, an RSS feed is available from the SBA customer and partner pages.

Table of Contents

What’s In This SBA Guide

Cisco SBA Data Center

Cisco SBA helps you design and quickly deploy a full-service business network. A Cisco SBA deployment is prescriptive, out-of-the-box, scalable, and flexible.

Cisco SBA incorporates LAN, WAN, wireless, security, data center, application optimization, and unified communication technologies—tested together as a complete system. This component-level approach simplifies system integration of multiple technologies, allowing you to select solutions that solve your organization’s problems—without worrying about the technical complexity.

Cisco SBA Data Center is a comprehensive design that scales from a server room to a data center for networks with up to 10,000 connected users. This design incorporates compute resources, security, application resiliency, and virtualization.

Route to Success

To ensure your success when implementing the designs in this guide, you should first read any guides that this guide depends upon—shown to the left of this guide on the route below. As you read this guide, specific prerequisites are cited where they are applicable.

About This Guide

This *deployment guide* contains one or more deployment chapters, which each include the following sections:

- Business Overview—Describes the business use case for the design. Business decision makers may find this section especially useful.
- Technology Overview—Describes the technical design for the business use case, including an introduction to the Cisco products that make up the design. Technical decision makers can use this section to understand how the design works.
- Deployment Details—Provides step-by-step instructions for deploying and configuring the design. Systems engineers can use this section to get the design up and running quickly and reliably.

You can find the most recent series of Cisco SBA guides at the following sites:

Customer access: http://www.cisco.com/go/sba

Partner access: http://www.cisco.com/go/sbachannel

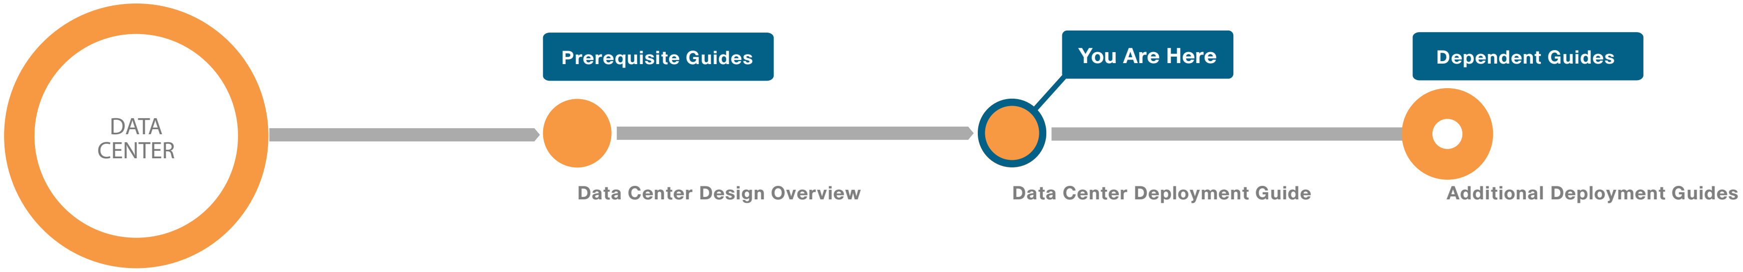
DATA CENTER
Prerequisite Guides
Data Center Design Overview
You Are Here
Data Center Deployment Guide
Dependent Guides
Additional Deployment Guides

Introduction

The Cisco Smart Business Architecture (SBA) data center foundation is a comprehensive architecture designed to provide data center Ethernet and storage networking, security, and load balancing for up to 300 server ports with a mix of physical and logical servers. This out-of-the-box approach is simple, easy to use, affordable, scalable, and flexible. The architecture for the Cisco SBA data center builds upon the server room deployment detailed in the *Server Room Deployment Guide*.

The *Cisco SBA—Data Center Deployment Guide* incorporates Ethernet, storage network, servers, security, and application resiliency tested together as a solution. This solution-level approach to building out an architecture simplifies the system integration normally associated with multiple technologies, allowing you to select the modules that meet your organization's requirements rather than worrying about matching components and worrying about interoperability.

Cisco SBA was designed to be easy to configure, deploy, and manage. This architecture:

- Provides a solid foundation
- Makes deployment fast and easy
- Accelerates your ability to easily deploy new servers and additional services
- Avoids the need to reengineer the network as your organization grows

This guide includes the following chapters:

- The first chapter covers elements of the data center design regarding the physical environment. The aspects of power, cooling, mounting racks, and space required are outlined for consideration in your data center design.
- The "Ethernet Infrastructure" chapter establishes the foundation connectivity for your data center network as it outgrows the server farm size. This section focuses on building a central connection point for the application servers that drive the organization and the services that surround them. The Ethernet chapter explains how to configure Layer 2 and Layer 3 connectivity in the data center and the communications path to the rest of the organization.
- The "Storage Infrastructure" chapter shows how the foundation Ethernet design accommodates IP-based network storage for network attached storage (NAS). The storage infrastructure chapter shows in depth how to deploy a Fibre Channel storage area network (SAN) using the Cisco Nexus 5500UP switches as the SAN core.
- The "Compute Connectivity" chapter explains the various host connectivity options that you can use in the data center. The chapter covers dual-homed and single-homed servers, and blade server systems' connectivity to the network.
- The "Network Security" chapter focuses on the deployment of firewalls to protect the critical and sensitive information assets of your organization. The intrusion prevention system (IPS) section explains how to deploy Cisco IPS to monitor your network for intrusions and attacks.
- The "Application Resiliency" chapter shows how server load balancing can be used to quickly grow server application farms, monitor server and application operation, and balance loads across multiple servers for better performance.
- The appendices provide the complete list of products used in the lab testing of this architecture, as well as the software revisions used on the products, and a list of major changes to the guide since it was last published.

To enhance the architecture, there are also a number of supplemental guides that address specific functions, technologies, or features from Cisco and Cisco partners that may be important to solving your organization's requirements.

Design Goals

The Cisco SBA program follows a consistent design process of building a network based on layers of services. The primary building block is the foundation layer upon which all other services rely. The data center foundation must be resilient, scalable, and flexible to support data center services, which add value, performance, and reliability. The ultimate goal of the design is to support the user services that drive the organization's success. Figure 1 illustrates the Cisco SBA data center design layered services.

Figure 1 - Cisco SBA data center pyramid of service layers

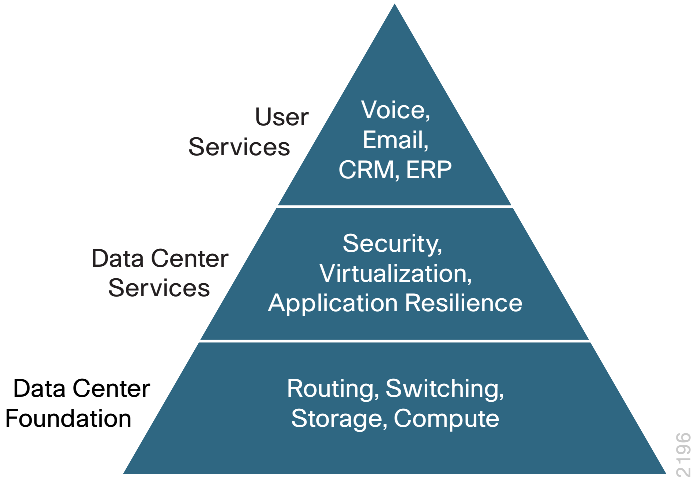
User Services
Voice, Email, CRM, ERP
Data Center Services
Security, Virtualization, Application Resilience
Data Center Foundation
Routing, Switching, Storage, Compute
2196

The Cisco SBA deployment guides are all designed to use a modular concept of building out a network. Each module is focused on the following principles:

- Ease of use—A top requirement was to develop a design that could be deployed with the minimal amount of configuration and day-two management.
- Cost-effective—Another critical requirement in the selection of products was to align with the requirements for a data center that scales to up to 300 server ports.
- Flexibility and scalability—As the company grows, so too must its infrastructure. Products selected needed to have the ability to grow or be repurposed within the architecture.
- Reuse—The goal, when possible, was to reuse the same products throughout the various modules to minimize the number of products required for spares.

Business Overview

Organizations encounter many challenges as they work to scale their information-processing capacity to keep up with demand. In a new organization, a small group of server resources may be sufficient to provide necessary applications such as file sharing, email, database applications, and web services. Over time, demand for increased processing capacity, storage capacity, and distinct operational control over specific servers can cause a growth explosion commonly known as "server sprawl." An organization can then use some of the same data center technologies that larger organizations use to meet expanding business requirements in a way that keeps capital and operational expenses in check. This deployment guide provides reference architecture to facilitate rapid adoption of these data center technologies by using a common, best-practices configuration.

The Cisco SBA data center design provides an evolution from the basic "server room" infrastructure. The Cisco SBA data center design is designed to address five primary business challenges:

- Supporting rapid application growth
- Managing growing data storage requirements
- Optimizing the investment in server processing resources
- Securing the organization's critical data
- Increasing application availability

Supporting Rapid Application Growth

As applications scale to support a larger number of users or new applications are deployed, the number of servers required to meet the needs of the organization often increases. The first phase of the server room evolution is often triggered when the organization outgrows the capacity of the existing server room network. Many factors can limit the capacity of the existing facility, including rack space, power, cooling, switching throughput, or basic network port count to attach new servers. The architecture outlined in this guide is designed to allow the organization to smoothly scale the size of the server environment and network topology as business requirements grow.

Managing Growing Data Storage Requirements

As application requirements grow, the need for additional data storage capacity also increases. This can initially cause issues when storage requirements for a given server increase beyond the physical capacity of the server hardware platform in use. As the organization grows, the investment in additional storage capacity is most efficiently managed by moving to a centralized storage model. A centralized storage system can provide disk capacity across multiple applications and servers, providing greater scalability and flexibility in storage provisioning.

A dedicated storage system provides multiple benefits beyond raw disk capacity. Centralized storage systems can increase the reliability of disk storage, which improves application availability. Storage systems allow an organization to provide increased capacity to a given server over the network without needing to physically attach new devices to the server itself. More sophisticated backup and data replication technologies are available in centralized storage systems, which helps protect the organization against data loss and application outages.

Optimizing the Investment in Server Processing Resources

As an organization grows, physical servers are often dedicated to single applications to increase stability and simplify troubleshooting. However, these servers do not operate at high levels of processor utilization for much of the day. Underutilized processing resources represent an investment by the organization that is not being leveraged to its full potential.

Server virtualization technologies allow a single physical server to run multiple virtual instances of a "guest" operating system, creating virtual machines (VMs). Running multiple VMs on server hardware helps to more fully utilize the organization's investment in processing capacity, while still allowing each VM to be viewed independently from a security, configuration, and troubleshooting perspective.

Server virtualization and centralized storage technologies complement one another, allowing rapid deployment of new servers and reduced downtime in the event of server hardware failures. VMs can be stored completely on the centralized storage system, which decouples the identity of the VM from any single physical server. This allows the organization great flexibility when rolling out new applications or upgrading server hardware.

The architecture defined in this guide is designed to facilitate easy deployment of server virtualization, while still providing support for the existing installed base of equipment. Supplemental guides are included with this series which focus specifically on server virtualization.

Securing the Organization's Critical Data

With communication and commerce in the world becoming increasingly Internet-based, network security quickly becomes a primary concern in a growing organization. Often organizations will begin by securing their Internet edge connection, considering the internal network a trusted entity. However, an Internet firewall is only one component of building security into the network infrastructure.

Frequently, threats to an organization's data may come from within the internal network. This may come in the form of onsite vendors, contaminated employee laptops, or existing servers that have already become compromised and may be used as a platform to launch further attacks. With the centralized repository of the organization's most critical data typically being the data center, security is no longer considered an optional component of a complete data center architecture plan.

The Cisco SBA data center design illustrates how to cleanly integrate network security capabilities such as firewall and intrusion prevention, protecting areas of the network housing critical server and storage resources. The architecture provides the flexibility to secure specific portions of the data center or insert firewall capability between tiers of a multi-tier application according to the security policy agreed upon by the organization.

Increasing Application Availability

With the expanding global presence and around-the-clock operations of organizations, key applications that drive the business must be available when the workforce needs them. Availability of applications can be threatened by overloaded servers and server or application failure. Unbalanced utilization can drive unacceptable response times for some users and satisfactory operation for others, making it difficult for IT teams to diagnose.

Figure 2 - Application server farm in various states of operation

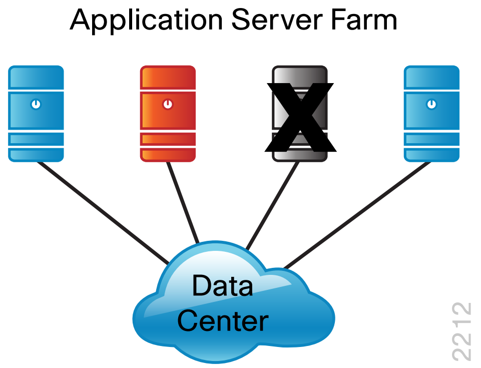

Application availability drives productivity and customer satisfaction, which are critical to an organization's success. IT organizations require the ability to monitor beyond simple server availability to application availability, and to be able to add more servers to an application server farm quickly and transparently.

Technology Overview

The Cisco SBA data center design is designed to allow organizations to take an existing server room environment to the next level of performance, flexibility, and security. Figure 3 provides a high-level overview of this architecture.

Figure 3 - Cisco SBA data center design

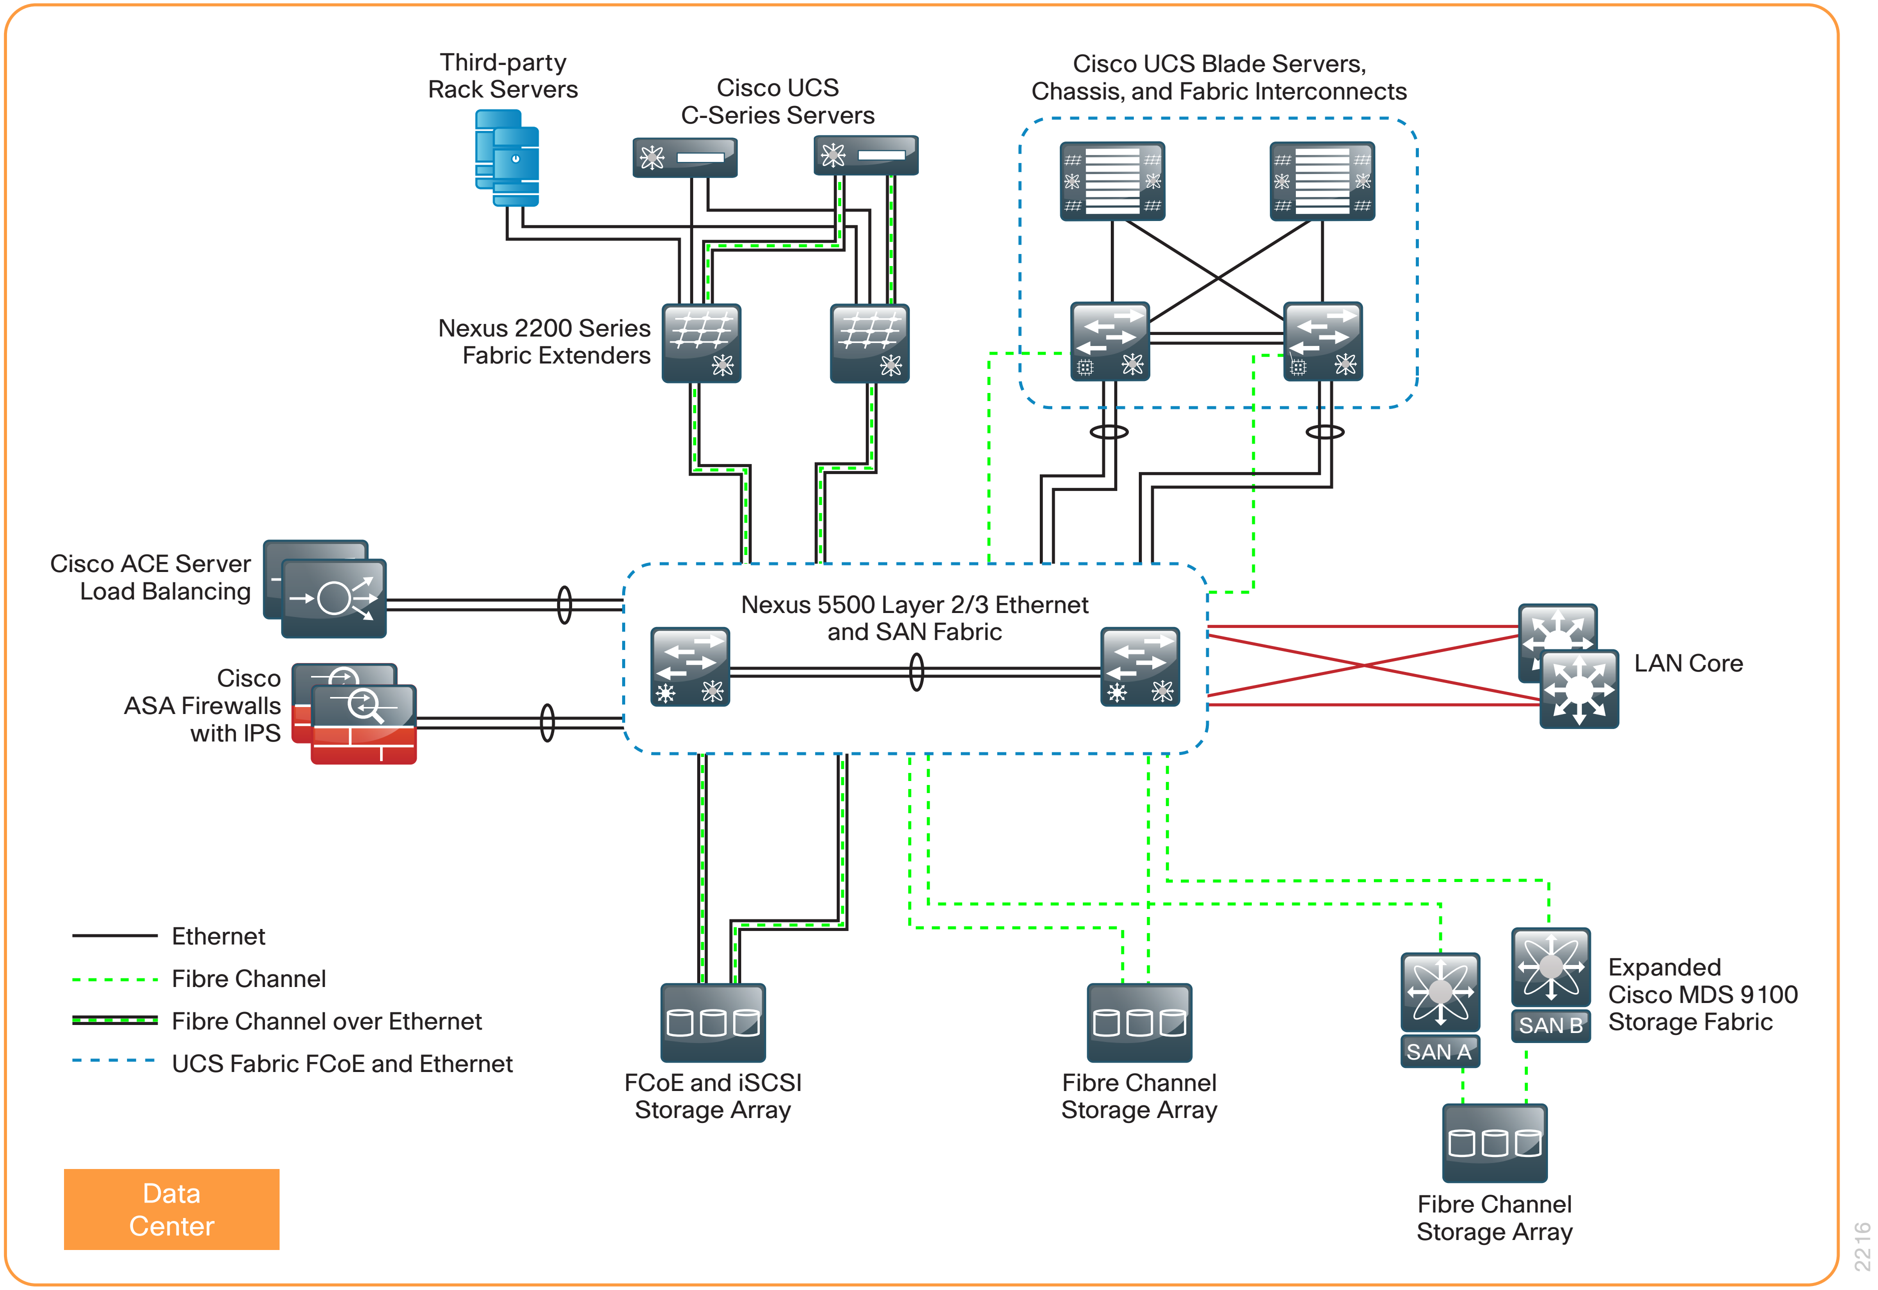

The Cisco SBA data center design is designed to stand alone, if deployed at an offsite facility, or to connect to one of the Cisco SBA Layer-3 Ethernet core solutions as documented in *Cisco SBA—Borderless Networks LAN Design Overview*. The following technology areas are included within this reference architecture.

Ethernet Infrastructure

The Ethernet infrastructure forms the foundation for resilient Layer 2 and Layer 3 communications in the data center. This layer provides the ability to migrate from your original server farm to a scalable architecture capable of supporting Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet connectivity for hundreds of servers in a modular approach.

The core of the Cisco SBA data center is built on the Cisco Nexus 5500UP series switches. Cisco Nexus 5500UP series is a high-speed switch capable of Layer 2 and Layer 3 switching with the Layer 3 daughter card tested in this design. Cisco Nexus 5500UP series 48-port model is used in this design, with an available 96-port model for higher density requirements. Cisco Nexus 5500UP supports Fabric Extender (FEX) technology, which provides a remote line card approach for fan out of server connectivity to top of rack for Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet requirements. The physical interfaces on the Cisco FEX are programmed on the Cisco Nexus 5500UP switches, simplifying the task of configuration by reducing the number of devices you have to touch to deploy a server port.

The Cisco Nexus 5500UP series features Virtual Port Channel (vPC) technology, which provides a loop-free approach to building out the data center in which any VLAN can appear on any port in the topology without spanning-tree loops or blocking links. The data center core switches are redundant with sub-second failover so that a device failure or maintenance does not prevent the network from operating.

Storage Infrastructure

Storage networking is key to solving the growing amount of data storage that an organization has to struggle with. Centralized storage reduces the amount of disk space trapped on individual server platforms and eases the task of providing backup to avoid data loss. The Cisco SBA data center design uses Cisco Nexus 5500UP series switches as the core of the network. The importance of this model switch is that it has universal port (UP) capabilities. A universal port is capable of supporting Ethernet, Fibre Channel, and Fibre Channel over Ethernet (FCoE) on any port. This allows the data center core to support multiple storage networking technologies like Fibre Channel storage area network (SAN), Internet Small Computer System Interface (iSCSI), and network attached storage (NAS) on a single platform type. This not only reduces costs to deploy the network but saves rack space in expensive data center hosting environments.

Cisco Nexus 5500UP Fibre Channel capabilities are based on the Cisco NX-OS operating system and seamlessly interoperate with the Cisco MDS Series SAN switches for higher-capacity Fibre Channel requirements. This deployment chapter includes procedures for interconnecting between Cisco Nexus 5500UP series and Cisco MDS series for Fibre Channel SAN. Cisco MDS series can provide an array of advanced services for Fibre Channel SAN environments where high-speed encryption, inter-VSAN routing, tape services, or Fibre Channel over IP extension might be required.

Compute Connectivity

There are many ways to connect a server to the data center network for Ethernet and Fibre Channel transport. This chapter provides an overview of connectivity ranging from single-homed Ethernet servers to a dual-homed Fabric Extender, and dual-homed servers that might use active/standby network interface card (NIC) teaming or EtherChannel for resiliency. Servers that use 10-Gigabit Ethernet can collapse multiple Ethernet NICs and Fibre Channel host bus adapters (HBAs) onto a single wire using converged network adapters (CNAs) and FCoE. Dual-homing the 10-Gigabit Ethernet servers with FCoE provides resilient Ethernet transport and Fibre Channel connections to SAN-A/SAN-B topologies. This chapter also provides an overview of how the integrated connectivity of Cisco Unified Computing System (UCS) blade server systems work and considerations for connecting a non–Cisco blade server system to the network.

Network Security

Within a data center design, there are many requirements and opportunities to include or improve security for customer confidential information and the organization's critical and sensitive applications. The data center design is tested with the Cisco ASA 5500 series firewall. Cisco ASA 5500 provides high-speed processing for firewall rule sets and high bandwidth connectivity with multiple 10-Gigabit Ethernet ports for resilient connectivity to the data center core switches. Cisco ASA 5500 also has a slot for services, and in this design provides an IPS module to inspect application layer data, to detect attacks and snooping, and to block malicious traffic based on the content of the packet or the reputation of the sender. The Cisco ASA 5500 firewalls with IPS modules are deployed in a pair, which provides an active/standby resiliency to prevent downtime in the event of a failure or platform maintenance.

Application Resiliency

Application performance and availability directly affect employee productivity and customer satisfaction, which drives the bottom line of an organization. As organizations expand to do business in a 24-hour, globally available environment, they find it even more important to make sure that critical applications are operating at peak performance.

This architecture includes Cisco Application Control Engine (ACE) to provide the latest technology for Layer 4 through Layer 7 switching and server load balancing (SLB). Server load balancers can spread the load across multiple servers for an application, and actively probe the servers and applications for load and health statistics to prevent overload and application failures. Cisco ACE also offers TCP processing offload, Secure Sockets Layer (SSL) offload, compression, and various other acceleration technologies. The Cisco ACE 4710 appliances used in this architecture are scalable to multi-Gigabit operation and are deployed as an active/standby pair to prevent outage from device failure or maintenance.

This architecture is designed to allow an organization to position its network for growth while controlling both equipment costs and operational costs. The deployment processes documented in this chapter provide concise, step-by-step instructions for completing the configuration of the components of the architecture to get your network up and running. This approach allows you to take advantage of some of the newer technologies being used in the data centers of very large organizations without encountering a steep learning curve for the IT staff. Although this architecture has been designed and validated as a whole, the modular nature of this chapter allows you to perform a gradual migration by choosing specific elements of the architecture to implement first.

Notes

Physical Environment

Business Overview

When building or changing a network, you have to carefully consider the location where you will install the equipment. When building a server room, a switch closet, or even a data center, take three things into consideration: power, cooling, and racking. Know your options in each of these categories, and you will minimize surprises and moving of equipment later on.

Technology Overview

The Cisco SBA data center design provides a resilient environment with redundant platforms and links; however, this cannot protect your data center from a complete failure resulting from a total loss of power or cooling. When designing your data center, you must consider how much power you will require, how you will provide backup power in the event of a loss of your power feed from your provider, and how long you will retain power in a backup power event. You also need to consider that servers, networking equipment, and appliances in your data center dissipate heat as they operate, which requires that you develop a proper cooling design that includes locating equipment racks to prevent hotspots.

Power

Know what equipment will be installed in the area. You cannot plan electrical work if you do not know what equipment is going to be used. Some equipment requires standard 110V outlets that may already be available. Other equipment might require much more power.

Does the power need to be on all the time? In most cases where servers and storage are involved, the answer is yes. Applications don't react very well when the power goes out. To prevent power outages, you need an uninterruptable power supply (UPS). During a power interruption, the UPS will switch over the current load to a set of internal or external batteries. Some UPSs are online, which means the power is filtered through the batteries all the time; others are switchable, meaning they use batteries only during power loss. UPSs vary by how much load they can carry and for how long. Careful planning is required to make sure the correct UPS is purchased, installed, and managed correctly. Most UPSs provide for remote monitoring and the ability to trigger a graceful server shutdown for critical servers if the UPS is going to run out of battery.

Distributing the power to the equipment can change the power requirements as well. There are many options available to distribute the power from the outlet or UPS to the equipment. One example would be using a power strip that resides vertically in a cabinet that usually has an L6-30 input and then C13/C19 outlets with the output voltage in the 200–240V range. These strips should be—at a minimum—metered so one does not overload the circuits. The meter provides a current reading of the load on the circuit. This is critical, because a circuit breaker that trips due to being overloaded will bring down everything plugged into it with no warning, causing business downtime and possible data loss. For complete remote control, power strips are available with full remote control of each individual outlet from a web browser. These vertical strips also assist in proper cable management of the power cords. Short C13/C14 and C19/C20 power cords can be used instead of much longer cords to multiple 110V outlets or multiple 110V power strips.

Cooling

With power comes the inevitable conversion of power into heat. To put it simply: power in equals heat out. Planning for cooling of one or two servers and a switch with standard building air conditioning may work. Multiple servers and blade servers (along with storage, switches, etc.) need more than building air conditioning for proper cooling. Be sure to at least plan with your facilities team what the options are for current and future cooling. Many options are available, including in-row cooling, overhead cooling, raised floor with underfloor cooling, and wall-mounted cooling.

Equipment Racking

It's important to plan where to put the equipment. Proper placement and planning allow for easy growth. After you have evaluated power and cooling, you need to install racking or cabinets. Servers tend to be fairly deep and take up even more space with their network connections and power connections. Most servers will fit in a 42-inch deep cabinet, and deeper cabinets give more flexibility for cable and power management within the cabinet.

Be aware of what rails are required by your servers. Most servers now come with rack mounts that use the square hole–style vertical cabinet rails. Not having the proper rails can mean that you have to use adapters or shelves, which makes managing servers and equipment difficult if not sometimes impossible without removing other equipment or sacrificing space. Data center racks should use the square rail mounting options in the cabinets. Cage nuts can be used to provide threaded mounts for such things as routers, switches, shelves, etc. that you may need.

Summary

The physical environmental requirements for a data center require careful planning to provide for efficient use of space, scalability, and ease of operational maintenance. Working toward deployment of Cisco SBA allows you to plan the physical space for your data center with a vision towards the equipment you will be installing over time, even if you begin on a smaller scale. For additional information on data center power, cooling, and equipment racking, contact Cisco partners in the area of data center environmental products such as Panduit and APC.

Notes

Ethernet Infrastructure

Business Overview

As your organization grows, you may outgrow the capacity of the basic "server-room" Ethernet switching stack illustrated in the *Cisco SBA—Data Center Design Overview*. It is important to be prepared for the ongoing transition of available server hardware from 1-Gigabit Ethernet attachment to 10-Gigabit Ethernet. Multi-tier applications often divide browser-based client services, business logic, and database layers into multiple servers, increasing the amount of server-to-server traffic and driving performance requirements higher. As the physical environment housing the organization's servers grows to multiple racks, it also becomes more challenging to elegantly manage the cabling required to attach servers to the network. The use of 10-Gigabit Ethernet connections helps to improve overall network performance, while reducing the number of physical links required to provide the bandwidth.

In some organizations, the data center may be located at a facility other than the headquarters building. Some organizations will locate their data center at a remote facility where power or cooling more suitable for a data center is located; others may rent floor space, racks, and power from a communications service provider to lower their capital costs. The ability to locate the data center in a number of different locations requires a data center architecture that is flexible to adapt to different locations while still providing the core elements of the architecture.

Technology Overview

The foundation of the Ethernet network in the Cisco SBA data center is a resilient pair of Cisco Nexus 5500UP Series switches. These switches offer the ideal platform for building a scalable, high-performance data center supporting both 10-Gigabit and 1-Gigabit Ethernet attached servers. The Cisco SBA data center is designed to allow easy migration of servers and services from your original server room to a data center that can scale with your organization's growth.

The Cisco Nexus 5500UP switches with universal port (UP) capabilities provide support for Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel ports on a single platform. The Nexus 5500UP can act as the Fibre Channel SAN for the data center and connect into an existing Fibre Channel SAN. The Cisco Nexus 5000 Series also supports the Cisco Nexus 2000 Series Fabric Extenders. Fabric Extenders allow the switching fabric of the resilient switching pair to be physically extended to provide port aggregation in the top of multiple racks, reducing cable management issues as the server environment expands.

The Cisco SBA data center design leverages many advanced features of the Cisco Nexus 5500UP Series switch family to provide a central Layer 2 and Layer 3 switching fabric for the data center environment:

- The Layer 3 routing table can accommodate up to 8000 routes.
- The Layer 3 engine supports up to 8000 adjacencies or MAC addresses for the Layer 2 domain.
- The solution provides for up to 1000 IP Multicast groups when operating in the recommended Virtual Port Channel (vPC) mode.

A second generation of the Layer 3 engine for the Cisco Nexus 5548 and 5596 switches is now available. This second generation hardware version of the Layer 3 module will operate with the same scalability numbers listed above using current Cisco NX-OS software for the Cisco Nexus 5500, and will double the scalability for routing and adjacencies when enabled in a future Cisco NX-OS software release.

Reader Tip

More specific scalability design numbers for the Cisco Nexus 5500 Series platform can be found at: http://www.cisco.com/en/US/customer/docs/switches/datacenter/nexus5000/sw/configuration_limits/limits_513/nexus_5000_config_limits_513.html#wp328407.

The Layer 3 data center core connects to the Layer 3 LAN core designed in the *Cisco SBA—Borderless Networks LAN Deployment Guide* as shown in Figure 4.

Figure 4 - Data center core and LAN core change control separation

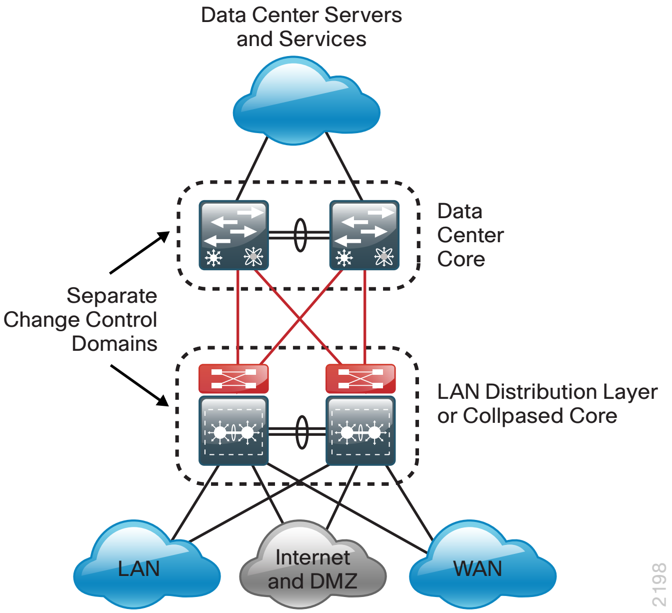
Data Center Servers and Services
Data Center Core
Separate Change Control Domains
LAN Distribution Layer or Collpased Core
LAN
Internet and DMZ
WAN
2198

The result of using Layer 3 to interconnect the two core layers is:

- A resilient Layer 3 interconnect with rapid failover.
- A logical separation of change control for the two core networks.
- A LAN core that provides a scalable interconnect for LAN, WAN, and Internet Edge.
- A data center core that provides interconnect for all data center servers and services.
- Intra-data center Layer 2 and Layer 3 traffic flows between servers and appliances that are switched locally on the data center core.
- A data center that has a logical separation point for moving to an offsite location while still providing core services without redesign.

This section provides an overview of the key features used in this topology and illustrates the specific physical connectivity that applies to the example configurations provided in the "Deployment Details" section.

Resilient Data Center Core

The data center needs to provide a topology where any data center VLAN can be extended to any server in the environment to accommodate new installations without disruption, and also the ability to move a server load to any other physical server in the data center. Traditional Layer 2 designs with LAN switches use spanning tree, which creates loops when a VLAN is extended to multiple access layer switches. Spanning Tree Protocol blocks links to prevent looping, as shown in Figure 5.

Figure 5 - Traditional design with spanning tree blocked links

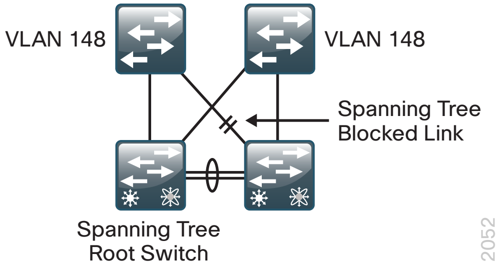

The Cisco Nexus 5500UP Series switch pair providing the central Ethernet switching fabric for the Cisco SBA data center is configured using vPC. The vPC feature allows links that are physically connected to two different Cisco Nexus switches to appear to a third downstream device to be coming from a single device, as part of a single Ethernet port channel. The third device can be a server, switch, or any other device or appliance that supports IEEE 802.3ad port channels. This capability allows the two data center core switches to build resilient, loop-free Layer 2 topologies that forward on all connected links instead of requiring Spanning Tree Protocol blocking for loop prevention.

Cisco NX-OS Software vPC used in the data center design and Cisco Catalyst Virtual Switching Systems (VSS) used in the *Cisco SBA—Borderless Networks LAN Design Overview* are similar technologies in that they allow the creation of Layer 2 port channels that span two switches. For Cisco EtherChannel technology, the term "multichassis EtherChannel" (MCEC) refers to either technology interchangeably. MCEC links from a device connected using vPC to the data center core and provides spanning-tree loop–free topologies, allowing VLANs to be extended across the data center while maintaining a resilient architecture.

A vPC consists of two vPC peer switches connected by a peer link. Of the vPC peers, one is primary and one is secondary. The system formed by the switches is referred to as a *vPC domain*.

Figure 6 - Cisco NX-OS vPC design

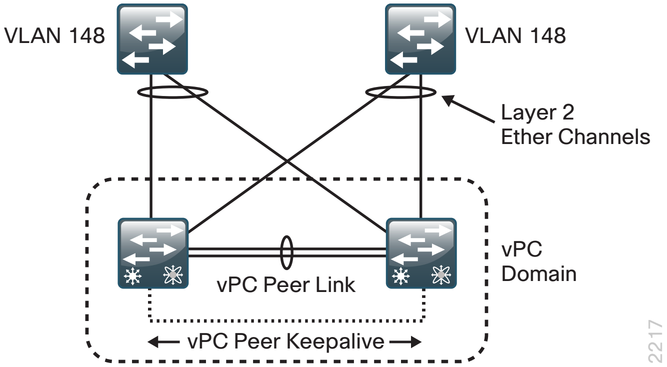

This feature enhances ease of use and simplifies configuration for the data center-switching environment.

Reader Tip

For more information on vPC technology and design, refer to the documents "Cisco NX-OS Software Virtual PortChannel: Fundamental Concepts" and "Spanning-Tree Design Guidelines for Cisco NX-OS Software and Virtual PortChannels" on www.cisco.com.

The Cisco SBA data center design uses Hot Standby Router Protocol (HSRP) for IP default gateway resiliency for data center VLANs. When combining HSRP with vPC, there is no need for aggressive HSRP timers to improve convergence, because both gateways are always active and traffic to either data center core will be locally switched for improved performance and resiliency.

Ethernet Fabric Extension

The Cisco Nexus 2000 Series Fabric Extender (FEX) delivers cost-effective and highly scalable 1-Gigabit Ethernet and 10-Gigabit Ethernet environments. Fabric extension allows you to aggregate a group of physical switch ports at the top of each server rack, without needing to manage these ports as a separate logical switch. The Cisco FEX behaves as a remote line card to the Cisco Nexus 5500UP switches. All configuration for Cisco FEX–connected servers is done on the data center core switches, which provide a centralized point to configure all connections for ease of use. Because the Cisco FEX acts as a line card on the Cisco Nexus 5500UP switch, extending VLANs to server ports on different Cisco FEXs does not create spanning-tree loops across the data center.

You can provide network resiliency by dual-homing servers into two separate fabric extenders, each of which is single-homed to one member of the Cisco Nexus 5500UP Series switch pair. To provide high availability for servers that only support single-homed network attachment, the Cisco FEX itself may instead be dual-homed using vPC into the two members of the data center core switch pair. Both the single-homed and dual-homed topologies provide the flexibility to have a VLAN appear on any port without loops or spanning-tree blocked links.

Figure 7 - Cisco FEX and vPC combined

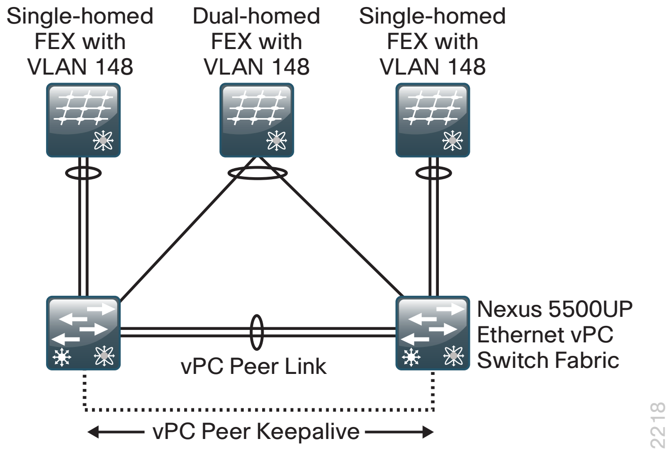

Our reference architecture example shown in Figure 8 illustrates single-homed and dual-homed Cisco FEX configurations with connected servers. Each Cisco FEX includes dedicated fabric uplink ports that are designed to connect to upstream Cisco Nexus 5500UP Series switches for data communication and management. Any 10-Gigabit Ethernet port on the Cisco Nexus 5500UP switch may be used for a Cisco FEX connection.

Tech Tip

When the Cisco Nexus 5500UP Series switches are configured for Layer 3 operation, they support up to sixteen connected Cisco FEXs as of Cisco NX-OS release 5.1(3)N1(1a). The Cisco FEX will support up to four or eight uplinks to the Cisco Nexus 5500UP parent switches, depending on the model of Cisco FEX in use and the level of oversubscription you want in your design. It is recommended, when possible, to configure the maximum number of fabric uplinks leveraging either twinax (CX-1) cabling or the Fabric Extender Transceiver (FET) and OM3 multi-mode fiber. At least two Cisco FEX uplinks to the data center core are recommended for minimum resiliency.

Figure 8 - Ethernet switching fabric physical connections

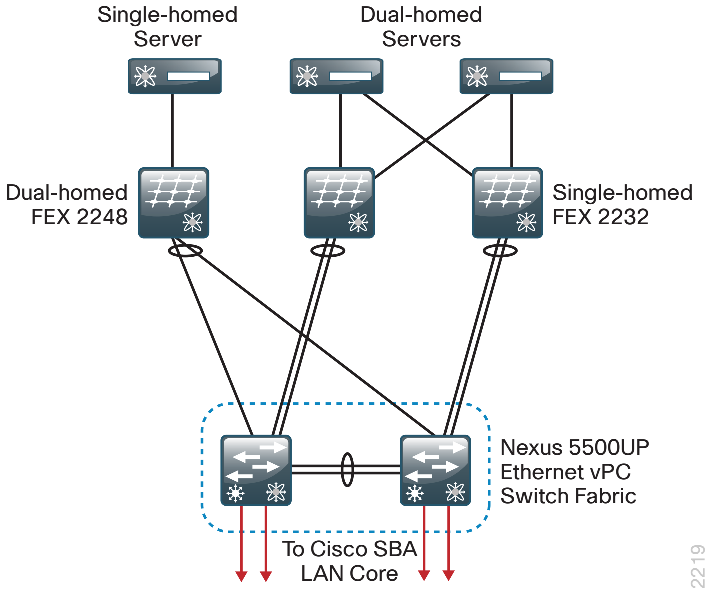

Quality of Service

To support the lossless data requirement of FCoE on the same links as IP traffic, the Nexus 5500 switches and the Nexus 2000 fabric extenders as a system implement an approach that uses Quality of Service (QoS) with a data center focus. Much of the QoS for classification and marking in the system is constructed through the use of the IEEE 802.1Q Priority Code Point, also known as Class of Service (CoS) bits in the header of the Layer 2 frame from hosts supporting FCoE and other trunked devices. As IP traffic arrives at an Ethernet port on the Cisco Nexus 5500 Series switch, it can also be classified at Layer 3 by differentiated services code point (DSCP) bits and IP access control lists (ACLs).

The traffic classifications are used for mapping traffic into one of six hardware queues, each appropriately configured for desired traffic handling. One queue is predefined for default traffic treatment, while one hardware queue is assigned for use by lossless FCoE traffic. The remaining four queues are available for use to support queuing consistent with the rest of Cisco SBA. For example, a priority queue will be defined for jitter-intolerant multimedia services in the data center.

Lacking the guarantee that all non-FCoE devices in the data center can generate an appropriate CoS marking required for application of QoS policy at ingress to a FEX, the Cisco SBA data center deployment takes the following QoS approach:

- FCoE traffic, as determined by Data Center Bridging Exchange (DCBX) negotiation with hosts, is given priority and lossless treatment end-to-end within the data center.
- Non-FCoE traffic without CoS classification for devices connected to a FEX is given default treatment over available links on ingress toward the Cisco Nexus 5500 switch, with suitable aggregated link bandwidth available to mitigate oversubscription situations. Traffic in the reverse direction toward the FEX is handled by the QoS egress policies on the Cisco Nexus 5500 switch.
- Classification by DSCP is configured at the port level and applied to IP traffic on ingress to the Cisco Nexus 5500 switch, either directly or after traversing a FEX connection. This classification is used to map traffic into the default queue or into one of the four non-FCoE internal queues to offer a suitable QoS per-hop behavior.
- To ensure consistent policy treatment for traffic directed through the Layer 3 engine, a CoS marking is also applied per Cisco Nexus 5500 internal queue. The CoS marking is used for classification of traffic ingress to the Layer 3 engine, allowing application of system queuing policies.

Non-FCoE devices requiring DSCP-based classification with guaranteed queuing treatment can be connected directly to the Cisco Nexus 5500 switch, versus taking the default uplink treatment when connected to a Cisco FEX port.

The QoS policy is also the method for configuring jumbo frame support on a per-class basis. Consistent per-CoS maximum transmission unit (MTU) requirements are applied system-wide for FCoE, as opposed to the port-based MTU configuration typical of devices used outside of the data center. Increasing MTU size can increase performance for bulk data transfers.

Deployment Details

The following configuration procedures are required to configure the Ethernet switching fabric for the Cisco SBA data center design.

Process

Configuring Ethernet Out-of-Band Management

1. Configure platform-specific switch settings
2. Configure switch universal settings
3. Apply the switch global configuration
4. Configure switch access ports
5. Configure switch links to the Layer 3 core

An increasing number of switching platforms, appliances, and servers utilize discrete management ports for setup, monitoring, and keepalive processes. The typical mid-tier data center is an ideal location for an Ethernet out-of-band management network, because the equipment is typically contained within in a few racks and does not require fiber-optic interconnect to reach far-away platforms.

This design uses a fixed-configuration Layer 2 switch for the out-of-band Ethernet management network. A switch like Cisco Catalyst 3560X is ideal for this purpose because it has dual power supplies for resiliency.

The out-of-band network provides:

- A Layer 2 path, independent of the data path of the Cisco Nexus 5500UP data center core switches, for vPC keepalive packets running over the management interface
- A path for configuration synchronization between Cisco Nexus 5500UP switches via the management interfaces
- A common connection point for data center appliance management interfaces like firewalls and load balancers
- A connectivity point for management ports on servers

Although the Layer 2 switch does provide a common interconnect for packets inside the data center, it needs to provide the ability for IT management personnel outside of the data center to access the data-center devices. The options for providing IP connectivity depend on the location of your data center.

If your data center is at the same location as your headquarters LAN, the core LAN switch can provide Layer 3 connectivity to the data center management subnet.

Figure 9 - Core LAN switch providing Layer 3 connectivity

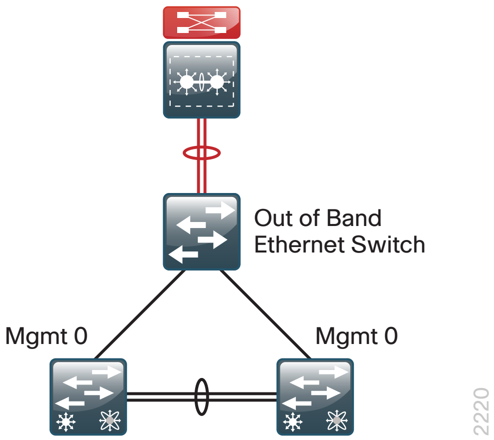

If your data center is located at a facility separate from a large LAN, the WAN router can provide Layer 3 connectivity to the data center management subnet.

Figure 10 - WAN router providing Layer 3 connectivity

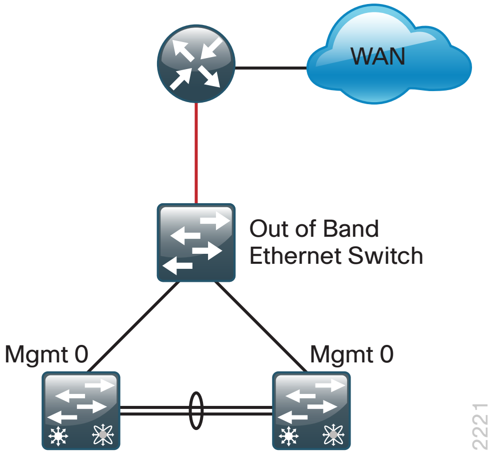

A third option for providing Layer 3 connectivity to the data center management subnet is to use the data center core Cisco Nexus 5500UP switches, as illustrated in Figure 11. This is the configuration described in this guide.

Figure 11 - Providing Layer 3 connectivity by using core Cisco Nexus 5500UP switches

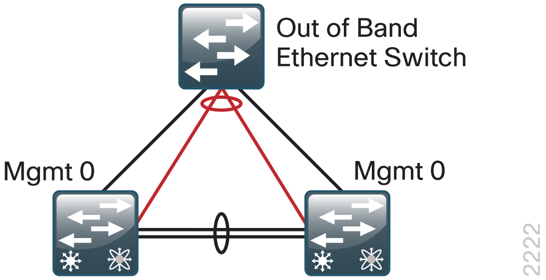

Tech Tip

When you use the data center core Cisco Nexus 5500UP switches for Layer 3 connectivity, the Layer 2 path for vPC keepalive packets will use the Ethernet out-of-band switch, because the Nexus 5500UP management ports are in a separate management Virtual Routing and Forwarding (VRF) path than the global packet switching of the Cisco Nexus 5500UP switches. Also, the management ports are in the same IP subnet, so they do not need a Layer 3 switch for packets between the data center core switches. The Layer 3 switched virtual interface (SVI) will provide connectivity for access outside of the data center.

Procedure 1 Configure platform-specific switch settings

Step 1: Configure the Catalyst 2960-S and 3750-X platform.

```
switch [switch number] priority 15
```

When there are multiple Cisco Catalyst 2960-S or Cisco Catalyst 3750-X Series switches configured in a stack, one of the switches controls the operation of the stack and is called the *stack master*.

When three or more switches are configured in a stack, choose a switch that does not have uplinks configured to configure as the stack master.

Step 2: Ensure the original master MAC address remains the stack MAC address after a failure.

```
stack-mac persistent timer 0
```

The default behavior when the stack master switch fails is for the newly active stack master switch to assign a new stack MAC address. This new MAC address assignment can cause the network to have to reconverge, because Link Aggregation Control Protocol (LACP) and many other protocols rely on the stack MAC address and must restart. As such, the **stack-mac persistent timer 0** command should be used to ensure that the original master MAC address remains the stack MAC address after a failure.

Because AutoQoS may not be configured on this device, you need to manually configure the global QoS settings by defining a macro that you will use in later procedures to apply the platform-specific QoS configuration.

Option 1. Configure QoS for Cisco Catalyst 3750-X and 3560-X

Step 1: Define a macro that you can use later to apply the platform-specific QoS configuration for Cisco Catalyst 3750-X and 3560-X switches.

```
mls qos map policed-dscp 0 10 18 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue input bandwidth 70 30
mls qos srr-queue input threshold 1 80 90
mls qos srr-queue input priority-queue 2 bandwidth 30
mls qos srr-queue input cos-map queue 1 threshold 2 3
mls qos srr-queue input cos-map queue 1 threshold 3 6 7
mls qos srr-queue input cos-map queue 2 threshold 1 4
mls qos srr-queue input dscp-map queue 1 threshold 2 24
mls qos srr-queue input dscp-map queue 1 threshold 3 48 49 50
51 52 53 54 55
mls qos srr-queue input dscp-map queue 1 threshold 3 56 57 58
59 60 61 62 63
mls qos srr-queue input dscp-map queue 2 threshold 3 32 33 40
41 42 43 44 45
mls qos srr-queue input dscp-map queue 2 threshold 3 46 47
mls qos srr-queue output cos-map queue 1 threshold 3 4 5
mls qos srr-queue output cos-map queue 2 threshold 1 2
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 32 33 40
41 42 43 44 45
mls qos srr-queue output dscp-map queue 1 threshold 3 46 47
mls qos srr-queue output dscp-map queue 2 threshold 1 16 17 18
19 20 21 22 23
mls qos srr-queue output dscp-map queue 2 threshold 1 26 27 28
29 30 31 34 35
mls qos srr-queue output dscp-map queue 2 threshold 1 36 37 38
39
mls qos srr-queue output dscp-map queue 2 threshold 2 24
mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50
51 52 53 54 55
mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58
59 60 61 62 63
mls qos srr-queue output dscp-map queue 3 threshold 3 0 1 2 3
4 5 6 7
mls qos srr-queue output dscp-map queue 4 threshold 1 8 9 11
13 15
mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
mls qos queue-set output 1 threshold 1 100 100 50 200
mls qos queue-set output 1 threshold 2 125 125 100 400
mls qos queue-set output 1 threshold 3 100 100 100 3200
mls qos queue-set output 1 threshold 4 60 150 50 200
mls qos queue-set output 1 buffers 15 25 40 20
mls qos
!
macro name EgressQoS
 mls qos trust dscp
 queue-set 1
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
@
!
```

Option 2. Configure QoS for Cisco Catalyst 2960-S

Step 1: Define a macro that you can use later to apply the platform-specific QoS configuration for Cisco Catalyst 2940-S switches.

```
mls qos map policed-dscp 0 10 18 24 46 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue output cos-map queue 1 threshold 3 4 5
mls qos srr-queue output cos-map queue 2 threshold 1 2
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 32 33 40 41 42 43 44 45
mls qos srr-queue output dscp-map queue 1 threshold 3 46 47
mls qos srr-queue output dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23
mls qos srr-queue output dscp-map queue 2 threshold 1 26 27 28 29 30 31 34 35
mls qos srr-queue output dscp-map queue 2 threshold 1 36 37 38 39
mls qos srr-queue output dscp-map queue 2 threshold 2 24
mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55
mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58 59 60 61 62 63
mls qos srr-queue output dscp-map queue 3 threshold 3 0 1 2 3 4 5 6 7
mls qos srr-queue output dscp-map queue 4 threshold 1 8 9 11 13 15
mls qos srr-queue output dscp-map queue 4 threshold 2 10 12 14
mls qos queue-set output 1 threshold 1 100 100 50 200
mls qos queue-set output 1 threshold 2 125 125 100 400
mls qos queue-set output 1 threshold 3 100 100 100 3200
mls qos queue-set output 1 threshold 4 60 150 50 200
mls qos queue-set output 1 buffers 15 25 40 20
mls qos
!
macro name EgressQoS
 mls qos trust dscp
 queue-set 1
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
@
!
```

Procedure 2 Configure switch universal settings

This procedure configures system settings that simplify and secure the management of the switch. The values and actual settings in the examples provided will depend on your current network configuration.

Table 1 - Common network services used in the deployment examples

Service	Address
Domain name	cisco.local
Active Directory, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP) server	10.4.48.10
Cisco Access Control System (ACS) Server	10.4.48.15
Network Time Protocol (NTP) Server	10.4.48.17

Step 1: Configure the device host name to make it easy to identify the device.

```
hostname [hostname]
```

Step 2: Configure VLAN Trunking Protocol (VTP) transparent mode. This deployment uses VTP transparent mode because the benefits of the alternative mode—dynamic propagation of VLAN information across the network—are not worth the potential for unexpected behavior that is due to operational error.

VTP allows network managers to configure a VLAN in one location of the network and have that configuration dynamically propagate out to other network devices. However, in most cases, VLANs are defined once during switch setup with few, if any, additional modifications.

```
vtp mode transparent
```

Step 3: Enable Rapid Per-VLAN Spanning-Tree (PVST+). PVST+ provides an instance of RSTP (802.1w) per VLAN. Rapid PVST+ greatly improves the detection of indirect failures or linkup restoration events over classic spanning tree (802.1D).

Although this architecture is built without any Layer 2 loops, you must still enable spanning tree. By enabling spanning tree, you ensure that if any physical or logical loops are accidentally configured, no actual Layer 2 loops will occur.

```
spanning-tree mode rapid-pvst
```

Step 4: Enable Unidirectional Link Detection (UDLD) Protocol. UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When UDLD detects a unidirectional link, it disables the affected interface and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree loops, black holes, and non-deterministic forwarding. In addition, UDLD enables faster link failure detection and quick reconvergence of interface trunks, especially with fiber-optic cables, which can be susceptible to unidirectional failures.

```
udld enable
```

Step 5: Set EtherChannels to use the traffic source and destination IP address when calculating which link to send the traffic across. This normalizes the method in which traffic is load-shared across the member links of the EtherChannel. EtherChannels are used extensively in this design because they contribute resiliency to the network.

```
port-channel load-balance src-dst-ip
```

Step 6: Configure DNS for host lookup.

At the command line of a Cisco IOS device, it is helpful to be able to type a domain name instead of the IP address for a destination.

```
ip name-server 10.4.48.10
```

Step 7: Configure device management protocols.

Secure HTTP (HTTPS) and Secure Shell (SSH) Protocol are secure replacements for the HTTP and Telnet protocols. They use Secure Sockets Layer (SSL) and Transport Layer Security (TLS) to provide device authentication and data encryption.

The SSH and HTTPS protocols enable secure management of the LAN device. Both protocols are encrypted for privacy, and the unsecure protocols—Telnet and HTTP—are turned off. Specify the **transport preferred none** command on vty lines to prevent errant connection attempts from the CLI prompt. Without this command, if the IP name server is unreachable, long timeout delays may occur for mistyped commands.

```
ip domain-name cisco.local
ip ssh version 2
no ip http server
ip http secure-server
line vty 0 15
  transport input ssh
  transport preferred none
```

Step 8: Enable Simple Network Management Protocol (SNMP) in order to allow the network infrastructure devices to be managed by a network management system (NMS), and then configure SNMPv2c both for a read-only and a read/write community string.

```
snmp-server community cisco RO
snmp-server community cisco123 RW
```

Step 9: If network operational support is centralized in your network, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
access-list 55 permit 10.4.48.0 0.0.0.255
line vty 0 15
  access-class 55 in
!
snmp-server community cisco RO 55
snmp-server community cisco123 RW 55
```

Caution

If you configure an access list on the vty interface, you may lose the ability to use SSH to log in from one router to the next for hop-by-hop troubleshooting.

Step 10: Configure the local login and password

The local login account and password provide basic access authentication to a switch which provides limited operational privileges. The enable password secures access to the device configuration mode. By enabling password encryption, you prevent the disclosure of plaintext passwords when viewing configuration files.

```
username admin password c1sco123
enable secret c1sco123
service password-encryption
aaa new-model
```

By default, HTTPS access to the switch will use the enable password for authentication.

Step 11: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the authentication, authorization, and accounting (AAA) server.

As networks scale in the number of devices to maintain, the operational burden to maintain local user accounts on every device also scales. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root-cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

Reader Tip

The AAA server used in this architecture is Cisco ACS. For details about Cisco ACS configuration, see the *Cisco SBA—Borderless Networks Device Management Using ACS Deployment Guide.*

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined in Step 10 on each network infrastructure device to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
tacacs server TACACS-SERVER-1
address ipv4 10.4.48.15
key SecretKey
!
aaa group server tacacs+ TACACS-SERVERS
  server name TACACS-SERVER-1
!
aaa authentication login default group TACACS-SERVERS local
aaa authorization exec default group TACACS-SERVERS local
aaa authorization console
ip http authentication aaa
```

Step 12: Configure a synchronized clock by programming network devices to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. Configure console messages, logs, and debug output to provide time stamps on output, which allows cross-referencing of events in a network.

```
ntp server 10.4.48.17
!
clock timezone PST -8
clock summer-time PDT recurring
!
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

Procedure 3 Apply the switch global configuration

Step 1: Configure the management VLAN.

The out-of-band management network will use a single VLAN, VLAN 163, for device connectivity.

```
vlan [vlan number]
name DC_ManagementVLAN
```

Step 2: Configure the switch with an IP address so that it can be managed via in-band connectivity, and assign an IP default gateway.

```
interface vlan [management vlan]
ip address [ip address] [mask]
no shutdown
ip default-gateway [default router]
```

Step 3: Configure bridge protocol data unit (BPDU) Guard globally to protect PortFast-enabled interfaces.

BPDU Guard protects against a user plugging a switch into an access port, which could cause a catastrophic, undetected spanning-tree loop.

A PortFast-enabled interface receives a BPDU when an invalid configuration exists, such as when an unauthorized device is connected. The BPDU Guard feature prevents loops by moving a nontrunking interface into an errdisable state when a BPDU is received on an interface when PortFast is enabled.

Figure 12 - Scenario that BPDU Guard protects against

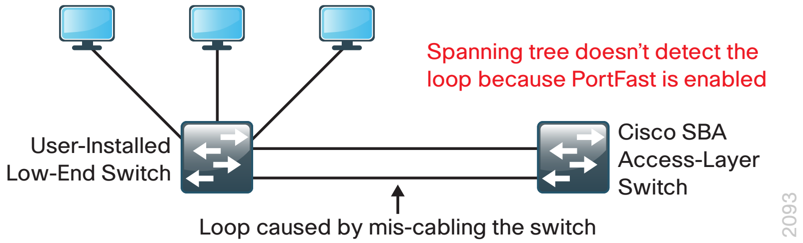
Spanning tree doesn't detect the loop because PortFast is enabled
User-Installed Low-End Switch
Cisco SBA Access-Layer Switch
Loop caused by mis-cabling the switch
2093

Disable the interface if another switch is plugged into the port.

```
spanning-tree portfast bpduguard default
```

Procedure 4 Configure switch access ports

To make configuration easier when the same configuration will be applied to multiple interfaces on the switch, use the **interface range** command. This command allows you to issue a command once and have it apply to many interfaces at the same time, which can save a lot of time because most of the interfaces in the access layer are configured identically. For example, the following command allows you to enter commands on all 24 interfaces (Gig 0/1 to Gig 0/24) simultaneously.

```
interface range Gigabitethernet 1/0/1-24
```

Step 1: Configure switch interfaces to support management console ports. This host interface configuration supports management port connectivity.

```
interface range [interface type] [port number]-[port number]
switchport access vlan [163]
switchport mode access
```

Step 2: Configure the switch port for host mode. Because only end-device connectivity is provided for the Ethernet management ports, shorten the time it takes for the interface to go into a forwarding state by enabling PortFast, disable 802.1Q trunking, and disable channel grouping.

```
switchport host
```

Example: Procedures 3 and 4

```
vlan 163
  name DC_ManagementVLAN
!
interface vlan 163
  description in-band management
  ip address 10.4.63.5 255.255.255.0
  no shutdown
!
ip default-gateway 10.4.63.1
!
spanning-tree portfast bpduguard default
!
interface range GigabitEthernet 1/0/1-22
  switchport access vlan 163
  switchport mode access
  switchport host
```

Procedure 5 Configure switch links to the Layer 3 core

As described earlier, there are various methods to connect to Layer 3 for connectivity to the data center out-of-band management network. The following steps describe configuring an EtherChannel for connectivity to the data center core Cisco Nexus 5500UP switches.

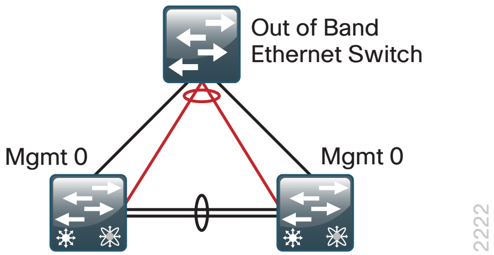

Step 1: Configure two or more physical interfaces to be members of the EtherChannel and set LACP to **active** on both sides. This forms a proper EtherChannel that does not cause any issues.

```
interface [interface type] [port 1]
  description Link to DC Core port 1
interface [interface type] [port 2]
  description Link to DC Core port 2
interface range [interface type] [port 1], [interface type] [port 2]
  channel-protocol lacp
  channel-group 1 mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
```

Step 2: Configure the trunk.

An 802.1Q trunk is used for the connection to this upstream device, which allows it to provide the Layer 3 services to all the VLANs defined on the management switch. The VLANs allowed on the trunk are pruned to only the VLANs that are active on the server room switch.

The Catalyst 2960-S does not require the **switchport trunk encapsulation dot1q** command.

```
interface Port-channel1
  description Etherchannel Link to DC Core for Layer 3
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan [management vlan]
  switchport mode trunk
  logging event link-status
  no shutdown
```

Reader Tip

The configuration on the data center core Cisco Nexus 5500UP switches for Layer 3 connectivity to the out-of-band management network will be covered in Procedure 10, "Configure management switch connection," in the "Configuring the Data Center Core" process later in this chapter.

Step 3: Save your management switch configuration.

```
copy running-config startup-config
```

Example

```
interface range GigabitEthernet 1/0/23-24
  description Links to DC Core for Layer 3
  channel-protocol lacp
  channel-group 1 mode active
  logging event link-status
  logging event trunk-status
  logging event bundle-status
  no shutdown
!
interface Port-channel 1
  description Etherchannel to DC Core for Layer 3
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 163
  switchport mode trunk
  logging event link-status
  no shutdown
```

Process

Configuring the Data Center Core

1. Establish physical connectivity
2. Perform initial device configuration
3. Configure QoS policies
4. Configure virtual port channel
5. Configure data center core global settings
6. Configure the IP routing protocol
7. Configure IP routing for VLANs
8. Configure IP Multicast routing
9. Configure connectivity to the SBA LAN core
10. Configure management switch connection
11. Configure Fabric Extender connections
12. Configure end node ports

Cisco Nexus 5500UP Series offers a simplified software management mechanism based on software licenses. These licenses are enforceable on a per-switch basis and enable a full suite of functionalities. The data center core layer is characterized by a Layer 3 configuration, so the Cisco Nexus 5500UP Series switch requires the Layer 3 license, which enables full Enhanced Interior Gateway Routing (EIGRP) functionality. The Fibre Channel license will be required when running native Fibre Channel or FCoE.

Procedure 1 Establish physical connectivity

Complete the physical connectivity of the Cisco Nexus 5500UP Series switch pair according to the illustration below.

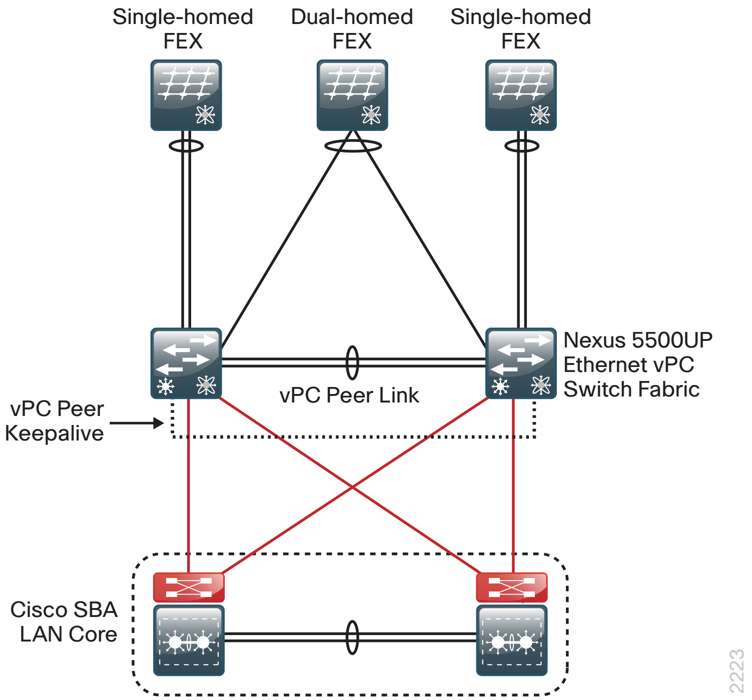

Step 1: Connect two available Ethernet ports between the two Cisco Nexus 5500UP Series switches.

These ports will be used to form the vPC peer-link, which allows the peer connection to form and supports forwarding of traffic between the switches if necessary during a partial link failure of one of the vPC port channels. It is recommended that you use at least two links for the vPC peer-link resiliency, although you can add more to accommodate higher switch-to-switch traffic.

Step 2: Connect two available Ethernet ports on each Cisco Nexus 5500UP Series switch to the Cisco SBA LAN core.

Four 10-Gigabit Ethernet connections will provide resilient connectivity to the Cisco SBA LAN core with aggregate throughput of 40 Gbps to carry data to the rest of the organization.

Step 3: Connect to a dual-homed FEX.

To support a dual-homed FEX with single-homed servers, connect fabric uplink ports 1 and 2 on the Cisco FEX to an available Ethernet port, one on each Cisco Nexus 5500UP Series switch. These ports will operate as a port channel to support the dual-homed Cisco FEX configuration.

Depending on the model Cisco FEX being used, up to four or eight ports can be connected to provide more throughput from the Cisco FEX to the core switch.

Step 4: Connect to a single-homed FEX.

Support single-homed FEX attachment by connecting fabric uplink ports 1 and 2 on each FEX to two available Ethernet ports on only one member of the Cisco Nexus 5500UP Series switch pair. These ports will be a port channel, but will not be configured as a vPC port channel because they have physical ports connected to only one member of the switch pair.

Depending on the model Cisco FEX being used, you can connect up to four or eight ports to provide more throughput from the Cisco FEX to the core switch.

Step 5: Connect to the out-of-band management switch.

This design uses a physically separate, standalone switch for connecting the management ports of the Cisco Nexus 5500 switches. The management ports provide out-of-band management access and transport for vPC peer keepalive packets, which are a part of the protection mechanism for vPC operation.

Procedure 2 Perform initial device configuration

This procedure configures system settings that simplify and secure the management of the solution. The values and actual settings in the examples provided will depend on your current network configuration.

Table 2 - Common network services used in the deployment examples

Service	Address
Domain name	cisco.local
Active Directory, DNS, DHCP server	10.4.48.10
Cisco ACS server	10.4.48.15
NTP server	10.4.48.17
EIGRP Autonomous System (AS)	100

Step 1: Connect to the switch console interface by connecting a terminal cable to the console port of the first Cisco Nexus 5500UP Series switch and then powering on the system to enter the initial configuration dialog box.

Step 2: Run the setup script and follow the Basic System Configuration Dialog for initial device configuration of the first Cisco Nexus 5500UP Series switch. This script sets up a system login password, SSH login, and the management interface addressing. Some setup steps will be skipped and covered in a later configuration step.

```
Do you want to enforce secure password standard (yes/no): y

  Enter the password for "admin":
  Confirm the password for "admin":

         ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic
configuration of the system. Setup configures only enough
connectivity for management of the system.

Please register Cisco Nexus 5000 Family devices promptly with
your supplier. Failure to register may affect response times
```

```
for initial service calls. Nexus devices must be registered to
receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/
no): y
  Create another login account (yes/no) [n]: n
  Configure read-only SNMP community string (yes/no) [n]: n
  Configure read-write SNMP community string (yes/no) [n]: n
  Enter the switch name : dc5548ax
  Enable license grace period? (yes/no) [n]: y
  Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: y
    Mgmt0 IPv4 address : 10.4.63.10
    Mgmt0 IPv4 netmask : 255.255.255.0
  Configure the default gateway? (yes/no) [y]: y
    IPv4 address of the default gateway : 10.4.63.1
  Configure advanced IP options? (yes/no) [n]:n  Enable the
ssh service? (yes/no) [y]: y
    Type of ssh key you would like to generate (dsa/rsa) : rsa
    Number of  key bits <768-2048> : 2048
  Enable the telnet service? (yes/no) [n]: n
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : 10.4.48.17
  Configure default interface layer (L3/L2) [L3]: L2
  Configure default switchport interface state (shut/noshut)
[shut]:shut
  Configure best practices CoPP profile (strict/moderate/
lenient/none) [strict]: moderate
  Configure default switchport trunk mode (on/off/auto) [on]:
auto
  Configure default switchport port mode F (yes/no) [n]:n
  Configure default zone policy (permit/deny) [deny]:y
  Enable full zoneset distribution? (yes/no) [n]: y
  Configure default zone mode (basic/enhanced) [basic]:
```

```
The following configuration will be applied:
password strength-check
switchname dc5548ax
license grace-period
interface mgmt0
ip address 10.4.63.10 255.255.255.0
no shutdown
ip route 0.0.0.0/0 10.4.63.1
  ssh key rsa 2048 force
  feature ssh
  no feature telnet
  no feature http-server
  ntp server 10.4.48.17 use-vrf management
  system default switchport
  system default switchport shutdown
  system default switchport trunk mode auto
  system default zone default-zone permit
  system default zone distribute full
  no system default zone mode enhanced
  policy-map type control-plane copp-system-policy

Would you like to edit the configuration? (yes/no) [n]: n
Use this configuration and save it? (yes/no) [y]: y
[########################################] 100%
dc5548ax login:
```

Step 3: Enable and configure system features.

Because of the modular nature of Cisco NX-OS, processes are only started when a feature is enabled. As a result, commands and command chains only show up after the feature has been enabled. For licensed features, the feature-name command can only be used after the appropriate license is installed. Cisco Nexus 5500UP Series requires a license for Layer 3 operation, Fibre Channel storage protocols, and FCoE N-Port Virtualization (NPV) operation. For more information on licensing, consult the *Cisco NX-OS Licensing Guide* on www.cisco.com.

The example configurations shown in this guide use the following features.

```
feature udld
feature interface-vlan
feature lacp
feature vpc
feature eigrp
feature fex
feature hsrp
feature pim
feature fcoe
```

Tech Tip

Although it is not used in this design, if the Fibre Channel–specific feature NPV is required for your network, you should enable it prior to applying any additional configuration to the switch. The NPV feature is the only feature that when enabled or disabled will erase your configuration and reboot the switch, requiring you to reapply any existing configuration commands to the switch.

Step 4: Configure the name server command with the IP address of the DNS server for the network. At the command line of a Cisco IOS device, it is helpful to be able to type a domain name instead of the IP address.

```
ip name-server 10.4.48.10
```

Step 5: Set local time zone for the device location. NTP is designed to synchronize time across all devices in a network for troubleshooting. In the initial setup script, you set the NTP server address. Now set the local time for the device location.

```
clock timezone PST -8 0
clock summer-time PDT 2 Sunday march 02:00 1 Sunday nov 02:00 60
```

Step 6: Define a read-only and a read/write SNMP community for network management.

```
snmp-server community cisco group network-operator
snmp-server community cisco123 group network-admin
```

Step 7: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the AAA server.

As networks scale in the number of devices to maintain, the operational burden to maintain local user accounts on every device also scales. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root-cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined in the setup script on each Cisco Nexus 5500 switch to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
feature tacacs+
tacacs-server host 10.4.48.15 key SecretKey
aaa group server tacacs+ tacacs
  server 10.4.48.15
  use-vrf default
  source-interface loopback 0
aaa authentication login default group tacacs
```

Step 8: If operational support is centralized in your network, you can increase network security by using an access list to limit the networks that can access your device. In this example, only devices on the 10.4.48.0/24 network will be able to access the device via SSH or SNMP.

```
ip access-list vty-acl-in
 permit tcp 10.4.48.0/24 any eq 22
line vty
 ip access-class vty-acl-in in
!
ip access-list snmp-acl
 permit udp 10.4.48.0/24 any eq snmp
snmp-server community cisco use-acl snmp-acl
snmp-server community cisco123 use-acl snmp-acl
```

Caution

If you configure an access list on the vty interface, you may lose the ability to use SSH to log in from one router to the next for hop-by-hop troubleshooting.

Step 9: Configure port operation mode. In this example, you enable ports 28 through 32 on a Cisco Nexus 5548UP switch as Fibre Channel ports.

```
slot 1
  port 28-32 type fc
```

The Cisco Nexus 5500UP switch has universal ports that are capable of running Ethernet+FCoE or Fibre Channel on a per-port basis. By default, all switch ports are enabled for Ethernet operation. Fibre Channel ports must be enabled in a contiguous range and be the high numbered ports of the switch baseboard and/or the high numbered ports of a universal port expansion module.

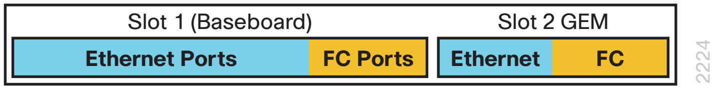

Tech Tip

Changing port type to FC requires a reboot in Cisco Nexus 5500UP NX-OS version 5.1(3)N1(1a) software to recognize the new port operation. This is subject to change in later releases of software. Ports will not show up in the configuration as FC ports if you did not enable the FCoE feature in Step 3.

Step 10: Save your configuration, and then reload the switch. Because the Cisco Nexus switch requires a reboot to recognize ports configured for Fibre Channel operation, this is a good point for you to reload the switch. If you are not enabling Fibre Channel port operation, you do not need to reload the switch at this point.

```
copy running-config startup-config
reload
```

Step 11: On the second Cisco Nexus 5500UP Series switch, repeat all of the steps of this procedure (Procedure 2). In Step 2, use a unique device name (dc5548bx) and IP address (10.4.63.11) for the mgmt0 interface—otherwise, all configuration details are identical.

Procedure 3 Configure QoS policies

QoS policies have been created for the Cisco SBA data center to align with the QoS configurations in the Cisco SBA LAN and WAN to protect multimedia streams, control traffic, and FCoE traffic, that flow through the data center. This is intended to be a baseline that you can customize to your environment if needed. At a minimum, it is recommended that FCoE QoS be configured to provide no-drop protected behavior in the data center.

QoS policies in this procedure are configured for Cisco Nexus 5500 and 2200 systems globally, and then later defined for assignment to Ethernet port and Ethernet port-channel configurations. Cisco Nexus FEX ports can use Layer 2 CoS markings for queuing. The Cisco Nexus 5500 ports can use Layer 2 CoS or Layer 3 DSCP packet marking for queue classification.

The system default FCoE policies are integrated into the overall Cisco SBA policies, to allow for the integration of FCoE-capable devices into the data center without significant additional configuration. If there is not a current or future need to deploy FCoE in the data center, the QoS policy can be adapted to use the standard FCoE qos-group for other purposes.

The following configurations will be created:

- Overall system classification via **class-map type qos** and **policy-map type qos** configurations will be based on CoS to associate traffic with the system internal qos-groups.
- Interface classification will be based on Layer 3 DSCP values via **class-map type qos** and **policy-map type qos** configurations to associate specific IP traffic types with corresponding internal qos-groups.
- System queue attributes based on matching qos-group are applied to set Layer 2 MTU, buffer queue-limit**,** and CoS mapping (for a Layer 3 daughter card) via **class-map type network qos** and **policy-map type network-qos**.
- System queue scheduling, based on matching qos-group**,** will be applied to set a priority queue for jitter-sensitive multimedia traffic and to apply bandwidth to weighted round-robin queues via **class-map type queuing** and **policy-map type queuing**. The bandwidth assignment for FCoE queuing should be adapted to the deployment requirements to guarantee end-to-end lossless treatment. For example, reallocating bandwidths to allow FCoE to assign **bandwidth percent 40** would be more appropriate for 4Gbps fibre channel traffic over a 10Gbps Ethernet link to a server or storage array.
- System-wide QoS **service-policy** will be configured in the system QoS configuration.
- Interface QoS **service-policy** will be defined for later use when configuring Ethernet end points for connectivity.

Step 1: Configure **class-map type qos** classification for global use, to match specific CoS bits. There is an existing system class-default which will automatically match any unmarked packets, unmatched CoS values, and packets marked with a CoS of zero. The FCoE class-map **type qos class-fcoe** is pre-defined and will be used in the policy map for FCoE traffic to ensure correct operation.

```
class-map type qos match-any PRIORITY-COS
  match cos 5
class-map type qos match-any CONTROL-COS
  match cos 4
class-map type qos match-any TRANSACTIONAL-COS
  match cos 2
class-map type qos match-any BULK-COS
  match cos 1
```

Step 2: Configure **policy-map type qos** policy for global use, to create the CoS-to-internal-qos-group mapping. The system-defined qos-group 0 is automatically created and does not require definition.

```
policy-map type qos DC-FCOE+1P4Q_GLOBAL-COS-QOS
  class type qos PRIORITY-COS
    set qos-group 5
  class type qos CONTROL-COS
    set qos-group 4
  class type qos class-fcoe
    set qos-group 1
  class type qos TRANSACTIONAL-COS
    set qos-group 2
  class type qos BULK-COS
    set qos-group 3
```

Step 3: Configure **class-map type qos** classification for Ethernet interface use, to allow for the mapping of traffic based on IP DSCP into the internal qos-groups of the Cisco Nexus 5500 switch. The **match cos** is used to match inbound Layer 2 CoS marked traffic, and also to map traffic destined for Cisco Nexus 5500 Layer 3 engine for traffic prioritization. All non-matched traffic will be handled by the system-defined class-default queue.

```
class-map type qos match-any PRIORITY-QUEUE
   match dscp ef
   match dscp cs5 cs4
   match dscp af41
   match cos 5
class-map type qos match-any CONTROL-QUEUE
   match dscp cs3
   match cos 4
class-map type qos match-any TRANSACTIONAL-QUEUE
   match dscp af21 af22 af23
   match cos 2
class-map type qos match-any BULK-QUEUE
   match dscp af11 af12 af13
   match cos 1
```

Step 4: Configure **policy-map type qos** policy to be applied to interfaces, for mapping DSCP classifications into internal qos-group. Interface policies are created to classify incoming traffic on Ethernet interfaces which are not members of a port-channel. These policies will also be assigned to port-channel virtual interfaces, but not the port-channel member physical interfaces.

```
policy-map type qos DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  class PRIORITY-QUEUE
    set qos-group 5
  class CONTROL-QUEUE
    set qos-group 4
  class TRANSACTIONAL-QUEUE
    set qos-group 2
  class BULK-QUEUE
    set qos-group 3
```

Step 5: Configure **class-map type queuing** classification for global use, to match to a specific internal qos-group for setting queue attributes. Five internal qos groups are available for assignment, plus an additional system qos-group 0 which is automatically created for default CoS traffic. The internal qos-group number is arbitrarily assigned, and does not necessarily match an equivalent CoS value. The FCoE class-map **type queuing class-fcoe** is pre-defined and will be used in the policy map for FCoE traffic to ensure correct operation.

```
class-map type queuing PRIORITY-GROUP
  match qos-group 5
class-map type queuing CONTROL-GROUP
  match qos-group 4
class-map type queuing TRANSACTIONAL-GROUP
  match qos-group 2
class-map type queuing BULK-GROUP
  match qos-group 3
```

Step 6: Configure **policy-map type queuing** policy for global use, to create appropriate system-wide qos-group attributes of bandwidth, priority, or weight, and FCoE lossless scheduling.

```
policy-map type queuing DC-FCOE+1P4Q_GLOBAL-GROUP-QUEUING
  class type queuing PRIORITY-GROUP
    priority
  class type queuing CONTROL-GROUP
    bandwidth percent 10
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing TRANSACTIONAL-GROUP
    bandwidth percent 25
  class type queuing BULK-GROUP
    bandwidth percent 20
  class type queuing class-default
    bandwidth percent 25
```

Step 7: Configure **class-map type network-qos** class-maps for global use, to match traffic for queue scheduling on a system-wide basis. As with the type queuing class-maps, the **type network-qos** class-maps can use one of five internal groups, along with an additional system configured qos-group 0 which is automatically created for default CoS. The internal qos-group number is arbitrarily assigned and does not necessarily match an equivalent CoS value. The FCoE class-map **type network-qos class-fcoe** is pre-defined and will be used in the policy map for FCoE traffic to ensure correct operation.

```
class-map type network-qos PRIORITY-SYSTEM
  match qos-group 5
class-map type network-qos CONTROL-SYSTEM
  match qos-group 4
class-map type network-qos TRANSACTIONAL-SYSTEM
  match qos-group 2
class-map type network-qos BULK-SYSTEM
  match qos-group 3
```

Step 8: Configure a **policy-map type network-qos** policy for global use, to apply system-wide queue scheduling parameters. The required FCoE queue behavior is configured with the recommended MTU of 2158, no-drop treatment, and the default buffer size of 79,360 bytes. The remaining queues take the default queue-limit of 22,720 bytes with an MTU of 1500, with two exceptions: the BULK-SYSTEM queue is assigned additional buffer space and a jumbo MTU of 9216 to improve performance for iSCSI and large data transfer traffic; by default, the class-default queue is assigned all remaining buffer space.

The Layer 3 routing engine requires CoS bits to be set for QoS treatment on ingress to and egress from the engine. Setting CoS ensures that traffic destined through the engine to another subnet is handled consistently, and the network-qos policy is where the CoS marking by system qos-group is accomplished.

```
policy-map type network-qos DC-FCOE+1P4Q_GLOBAL-SYSTEM-NETWORK-QOS
  class type network-qos PRIORITY-SYSTEM
   set cos 5
  class type network-qos CONTROL-SYSTEM
    set cos 4
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos TRANSACTIONAL-SYSTEM
    set cos 2
  class type network-qos BULK-SYSTEM
     mtu 9216
    queue-limit 128000 bytes
   set cos 1
  class type network-qos class-default
    multicast-optimize
    set cos 0
```

Step 9: Apply the created global policies.

```
system qos
  service-policy type qos input DC-FCOE+1P4Q_GLOBAL-COS-QOS
  service-policy type queuing input DC-FCOE+1P4Q_GLOBAL-GROUP-QUEUING
  service-policy type queuing output DC-FCOE+1P4Q_GLOBAL-GROUP-QUEUING
  service-policy type network-qos DC-FCOE+1P4Q_GLOBAL-SYSTEM-NETWORK-QOS
```

The output queuing applied with **system qos** defines how the bandwidth is shared among different queues for Cisco Nexus 5500 and Cisco Nexus FEX interfaces, and also defines how the bandwidth is shared among different queues on Cisco Nexus 5500 Layer 3 engine.

Step 10: If iSCSI is being used, additional classification and queuing can be added to map iSCSI storage traffic into the appropriate queue for bulk data. Classification of iSCSI traffic can be matched by well-known TCP ports through an ACL. The iSCSI class of traffic can then be added to the existing policy map to put the traffic into the correct qos-group.

```
ip access-list ISCSI
  10 permit tcp any eq 860 any
  20 permit tcp any eq 3260 any
  30 permit tcp any any eq 860
  40 permit tcp any any eq 3260
!
class-map type qos match-all ISCSI-QUEUE
  match access-group name ISCSI
policy-map type qos DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  class ISCSI-QUEUE
    set qos-group 3
```

Tech Tip

Use only **permit** actions in the ACLs for matching traffic for QoS policies on Cisco Nexus 5500. For more details on configuring QoS policies on Cisco Nexus 5500 please refer to: http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/513_n1_1/b_cisco_nexus_5000_qos_config_gd_513_n1_1_chapter_011.html#task_1135158

Use the **show queuing interface** command to display QoS queue statistics.

The Interface QoS service-policy DC-FCOE+1P4Q_INTERFACE-DSCP-QOS created in Step 4 will be assigned later in this guide to:

- Non-FEX Ethernet interfaces on Cisco Nexus 5500.

  Example:

  ```
  interface Ethernet1/1-27
    service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  ```

- Ethernet port-channel interfaces on Cisco Nexus 5500. The port-channel member physical links do not require the policy; they will inherit the service policy from the logical port-channel interface. This service policy is not required on port-channels connected to FEX network uplinks.

  Example:

  ```
  interface port-channel 2-3 , port-channel 5 , port-channel 9
    service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  ```

- FEX host port Ethernet interfaces, which are not port-channel members.

  Example:

  ```
  interface Ethernet102/1/1-48 , interface Ethernet104/1/1-32 , interface Ethernet105/1/1-32
    service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  ```

Procedure 4 Configure virtual port channel

Before you can add port channels to the switch in virtual port channel (vPC) mode, basic vPC peering must be established between the two Cisco Nexus 5500UP Series switches. The vPC peer link provides a communication path between the data center core switches that allows devices that connect to each core switch for resiliency to do so over a single Layer 2 EtherChannel.

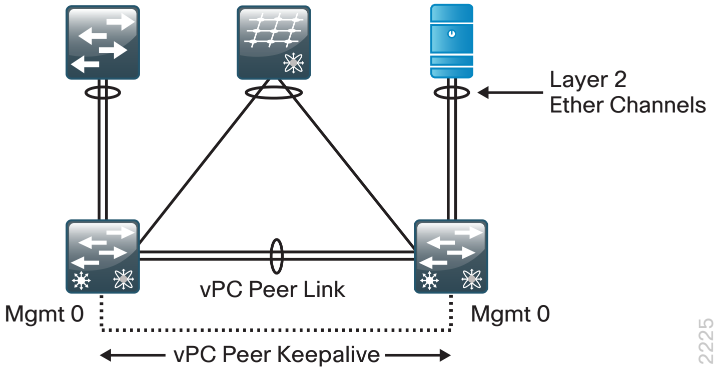

Step 1: Define a vPC domain number to identify the vPC domain to be common between the switches in the pair.

```
vpc domain 10
```

Step 2: Define a lower role priority for the vPC primary switch.

```
role priority 16000
```

The vPC secondary switch will be left at the default value of 32,667. The switch with lower priority will be elected as the vPC primary switch. If the vPC primary switch is alive and the vPC peer link goes down, the vPC secondary switch will suspend its vPC member ports to prevent potential looping while the vPC primary switch keeps all of its vPC member ports active. If the peer link fails, the vPC peer will detect the peer switch's failure through the vPC peer keepalive link.

Step 3: Configure vPC peer keepalive on both Cisco Nexus 5500 switches.

- On the first Cisco Nexus 5500UP switch, configure the peer-keepalive destination and source addresses.

```
peer-keepalive destination 10.4.63.11 source 10.4.63.10
```

- Change destination and source addresses and configure accordingly on the second Cisco Nexus 5500UP switch.

```
peer-keepalive destination 10.4.63.10 source 10.4.63.11
```

The peer-keepalive is ideally an alternate physical path between the two Cisco Nexus 5500UP switches running vPC to ensure that they are aware of one another's health even in the case where the main peer link fails. The peer-keepalive source IP address should be the address being used on the mgmt0 interface of the switch currently being configured. The destination address is the mgmt0 interface on the vPC peer.

Step 4: On the first Cisco Nexus 5500UP switch, configure the following vPC commands in the vPC domain configuration mode. This will increase resiliency, optimize performance, and reduce disruptions in vPC operations.

```
delay restore 360
auto-recovery
graceful consistency-check
peer-gateway
ip arp synchronize
```

The **auto-recovery** command has a default timer of 240 seconds. This time can be extended by adding the reload-delay variable with time in seconds. The auto-recovery feature for vPC recovery replaces the need for the original peer-config-check-bypass feature.

Step 5: On the first Cisco Nexus 5500UP switch, create a port channel interface to be used as the peer link between the two vPC switches. The peer link is the primary link for communications and for forwarding of data traffic to the peer switch, if required.

```
interface port-channel 10
  switchport mode trunk
  vpc peer-link
  spanning-tree port type network
```

Step 6: On the first Cisco Nexus 5500UP switch, configure the physical interfaces that connect the two Cisco Nexus 5500 switches together to the port channel. A minimum of two physical interfaces is recommended for link resiliency. The channel-group number must match the port-channel number used in the previous step. Different 10-Gigabit Ethernet ports (as required by your specific implementation) may replace the interfaces shown in the example.

```
interface Ethernet1/17
  description vpc peer link
  switchport mode trunk
  channel-group 10 mode active

interface Ethernet1/18
  description vpc peer link
  switchport mode trunk
  channel-group 10 mode active
```

Step 7: On the second Cisco Nexus 5500UP switch, repeat Step 4 through Step 6.

Step 8: Ensure that the vPC peer relationship has formed successfully by using the **show vpc** command.

```
dc5548ax# show vpc
Legend:
         (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                   : 10
Peer status                     : peer adjacency formed ok
vPC keep-alive status           : peer is alive
Configuration consistency status: success
Per-vlan consistency status     : success
Type-2 consistency status       : success
vPC role                        : primary
Number of vPCs configured       : 55
Peer Gateway                    : Enabled
Dual-active excluded VLANs      : -
Graceful Consistency Check      : Enabled

vPC Peer-link status
-------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ ------------------------------------------
1    Po10   up     1
```

Step 9: Verify successful configuration by looking for the peer status of “peer adjacency formed ok” and the keepalive status of “peer is alive”. If the status does not indicate success, double-check the IP addressing assigned for the keepalive destination and source addresses, as well as the physical connections.

Tech Tip

Do not be concerned about the “(*) - local vPC is down, forwarding via vPC peer-link” statement at the top of the command output at this time. After you have defined vPC port channels and if one of its member links is down or not yet configured, this information becomes a legend that shows the meaning of an asterisk next to your port channel in the listing.

Procedure 5 Configure data center core global settings

The data center core requires basic core operational configuration beyond the setup script.

The IP Subnet and VLAN assignment build on the assignments in the *Server Room Deployment Guide*. As you review configuration guidance, you will notice a difference in the second octet of the IP address. The second octet's address is assigned based on which design the configuration was developed in:

- IP address prefix 10.8 is used in the Cisco SBA server room design
- IP address prefix 10.4 is used in the Cisco SBA data center design

In this deployment guide, we have used the third octet of the IP address and added 100 to determine the VLAN number for easier reference. Adding 100 prevents a VLAN number from being one or zero, which can be a problem on some devices, while still making the VLAN ID easy to remember.

Table 3 - Cisco SBA data center VLANs

VLAN	VLAN name	IP address	Comments
148	Servers_1	10.4.48.0/24	General network server use
149	Servers_2	10.4.49.0/24	Used in the "Application Resiliency" chapter for the server load balancing VLAN
150	Servers_3	10.4.50.0/24	General server use
153	FW_Outside	10.4.53.0/25	Used for firewall outside interface routing
154	FW_Inside_1	10.4.54.0/24	Used in the "Network Security" chapter for firewall-protected servers
155	FW_Inside_2	10.4.55.0/24	Used in "Network Security" for firewall+IPS protected servers
156	PEERING_VLAN	10.4.56.0/30	Cisco Nexus 5500 intra-data center Layer 3 peering link
161	VMotion	10.4.61.0/24	Reserved for VMware VMotion traffic future use
162	iSCSI	10.4.62.0/24	Reserved for iSCSI storage traffic
163	DC-Management	10.4.63.0/24	Out-of-band data center management VLAN

Step 1: Create the necessary VLANs for data center operation.

```
vlan [vlan number]
 name [vlan name]
```

Step 2: Configure spanning tree.

Rapid Per-VLAN Spanning-Tree (PVST+) provides an instance of RSTP (802.1w) per VLAN. Rapid PVST+ greatly improves the detection of indirect failures or linkup restoration events over classic spanning tree (802.1D). Cisco Nexus 5500UP runs Rapid PVST+ by default.

Although this architecture is built without any Layer 2 loops, it is a good practice to assign spanning-tree root to the core switches. This design assigns spanning-tree root for a range of VLANs that may be contained in the data center.

- Configure the primary Cisco Nexus 5500UP switch as the spanning-tree root.

  ```
  spanning-tree vlan 1-400 root primary
  ```

- Configure the second Cisco Nexus 5500UP switch as the spanning-tree secondary.

  ```
  Spanning-tree vlan 1-400 root secondary
  ```

Step 3: Configure an in-band management interface. This examples uses an IP address out of the data center core addressing with a 32-bit address (host) mask.

```
interface loopback 0
ip address 10.4.56.254/32
ip pim sparse-mode
```

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. Because of this capability, the loopback address is the best way to manage the switch in-band and provides an additional management point to the out-of-band management interface. Layer 3 process and features are also bound to the loopback interface to ensure process resiliency.

The loopback interface for the second Cisco Nexus 5500UP switch will be 10.4.56.253/32.

Step 4: Configure EtherChannel port channels to use Layer 3 IP address and Layer 4 port number for load balance hashing. This optimizes load balancing on EtherChannel links and improves throughput to the Layer 3 routing engine in the Cisco Nexus 5500UP switch.

```
port-channel load-balance ethernet source-dest-port
```

Procedure 6 Configure the IP routing protocol

Step 1: Configure EIGRP as the IP routing protocol.

```
router eigrp 100
  router-id 10.4.56.254
```

The router ID for the second Cisco Nexus 5500UP switch will be 10.4.56.253/32.

EIGRP is the IP routing protocol used in the data center to be compatible with the Cisco SBA foundation LAN core and WAN. This example uses the same routing process ID so that routes can be exchanged with the LAN core.

In this configuration, the only parameter configured under the EIGRP process (router eigrp 1) is the router-ID. The loopback **0** IP address is used for the EIGRP router ID.

Step 2: Configure EIGRP on Layer 3 interfaces.

```
interface loopback 0
ip router eigrp 100
```

Cisco NX-OS routing configuration follows an interface-centric model. Instead of adding networks to be advertised via network statements, EIGRP is enabled on a per-interface basis. Each Layer 3 interface that carries a network that may be advertised via EIGRP requires the **ip router eigrp** statement.

Step 3: Configure the core Layer 3 peering link.

```
Interface Vlan 156
  ip address 10.4.56.1/30
  ip router eigrp 100
  ip pim sparse-mode
  no shutdown
```

To pass EIGRP routing updates between routing peers, EIGRP must be enabled on each end of a Layer 3 link. To avoid unnecessary EIGRP peering between the core data center switches across all data center VLAN-switched virtual interfaces, a single link will be used for active EIGRP peering in the data center core.

The peer Cisco Nexus 5500UP switch will use IP address 10.4.56.2/30.

Procedure 7 Configure IP routing for VLANs

Every VLAN that needs Layer 3 reachability between VLANs or to the rest of the network requires a Layer 3 switched virtual interface (SVI) to route packets to and from the VLAN.

Step 1: Configure the SVI.

```
interface Vlan [vlan number]
```

Step 2: Configure the IP address for the SVI interface.

```
ip address [ip address]/mask
```

Step 3: Disable IP redirect on the SVI. It is recommended that Internet Control Message Protocol (ICMP) IP redirects in vPC environments be disabled on SVIs for correct operation.

```
no ip redirects
```

Step 4: Configure the EIGRP process number on the interface. This advertises the subnet into EIGRP.

```
ip router eigrp 100
```

Step 5: Configure passive mode EIGRP operation. To avoid unnecessary EIGRP peer processing, configure server VLANs as passive.

```
ip passive-interface eigrp 100
```

Step 6: Configure HSRP. The Cisco Nexus 5500UP switches use HSRP to provide a resilient default gateway in a vPC environment. For ease of use, number the HSRP group number the same as the SVI VLAN number. Configure a priority greater than 100 for the primary HSRP peer, and leave the second switch at the default priority of 100.

```
hsrp [group number]
priority [priority]
ip [ip address of hsrp default gateway]
```

Tech Tip

Both data center core Cisco Nexus 5500UP switches can process packets for the assigned ip address of their SVI and for the HSRP address. In a vPC environment, a packet to either switch destined for the default gateway (HSRP) address is locally switched and there is no need to tune aggressive HSRP timers to improve convergence time.

- The following is an example configuration for the first Cisco Nexus 5500UP switch.

```
interface Vlan148
  no ip redirects
  ip address 10.4.48.2/24
  ip router eigrp 100
  ip passive-interface eigrp 100
  ip pim sparse-mode
  hsrp 148
    priority 110
    ip 10.4.48.1
  no shutdown
  description Servers_1
```

- The following is an example configuration for the peer Cisco Nexus 5500UP switch.

```
interface Vlan148
  no ip redirects
  ip address 10.4.48.3/24
  ip router eigrp 100
  ip passive-interface eigrp 100
  ip pim sparse-mode
  hsrp 148
    ip 10.4.48.1
  no shutdown
  description Servers_1
```

Procedure 8 Configure IP Multicast routing

The Cisco SBA Foundation LAN network enables IP Multicast routing for the organization by using **pim sparse-mode** operation. The configuration of IP Multicast for the rest of the network can be found in the *Cisco SBA—Borderless Networks LAN Deployment Guide*.

Step 1: Configure the data center core switches to discover the IP Multicast rendezvous point (RP) from the Cisco SBA LAN core. Every Layer 3 switch and router must be configured to discover the IP Multicast RP. The **ip pim auto-rp forward listen** command allows for discovery of the RP across **ip pim sparse-mode** links.

```
ip pim auto-rp forward listen
```

Step 2: Configure an unused VLAN for IP Multicast replication synchronization between the core Cisco Nexus 5500UP switches.

```
vpc bind-vrf default vlan 900
```

Tech Tip

The VLAN used for the IP Multicast **bind-vrf** cannot appear anyplace else in the configuration of the Cisco Nexus 5500UP switches. It must not be defined in the VLAN database commands and does not get included in the VLAN allowed list for the vPC core. It will automatically program packet replication across the vPC peer link trunk when needed.

Step 3: Configure IP Multicast to only be replicated across the vPC peer link when there is an orphan port of a vPC.

```
no ip igmp snooping mrouter vpc-peer-link
```

Step 4: Configure all Layer 3 interfaces for IP Multicast operation with the **pim sparse-mode** command.

```
ip pim sparse-mode
```

It is not necessary to configure IP Multicast on the management VLAN interface (interface vlan 163).

Procedure 9 Configure connectivity to the SBA LAN core

Virtual Port Channel **does not** support peering to another Layer 3 router across a vPC. This design will use dual-homed point-to-point Layer 3 interfaces between each data center core Cisco Nexus 5500UP switch to each Cisco Catalyst 6500 core LAN switch for data to and from the data center to the rest of the network. If your design has a single resilient Cisco Catalyst 4500 with redundant supervisors and redundant line cards, you will instead connect each data center Cisco Nexus 5500UP switch to each of the redundant line cards.

Table 4 - Example data center to LAN core with standalone Catalyst 6500 switches

Data Center Core			LAN Core		
Switch	Port	IP Address	IP Address	C6500-1	C6500-2
dc5548a	e1/19	10.4.40.50	10.4.40.49	Te 4/7	—
	e1/20	10.4.40.58	10.4.40.57	—	Te4/7
dc5548b	e1/19	10.4.40.54	10.4.40.53	Te4/8	—
	e1/20	10.4.40.62	10.4.40.61	—	Te4/7

Table 5 - Example data center to collapsed LAN core with Catalyst 6500 VSS pair

Data Center Core			LAN Core	
Switch	Port	IP Address	IP Address	C6500VSS
dc5548a	e1/19	10.4.40.50	10.4.40.49	Te1/4/6
	e1/20	10.4.40.58	10.4.40.57	Te2/4/6
dc5548b	e1/19	10.4.40.54	10.4.40.53	Te1/4/8
	e1/20	10.4.40.62	10.4.40.61	Te2/4/8

Table 6 - Example data center to collapsed LAN core with Catalyst 4500

Data Center Core			LAN Core	
Switch	Port	IP Address	IP Address	C4500
dc5548a	e1/19	10.4.40.50	10.4.40.49	Te1/4
	e1/20	10.4.40.58	10.4.40.57	Te2/4
dc5548b	e1/19	10.4.40.54	10.4.40.53	Te1/7
	e1/20	10.4.40.62	10.4.40.61	Te2/7

It is recommended you have at least two physical interfaces from each switch connected to the network core, for a total port channel of four resilient physical 10-Gigabit Ethernet links and 40 Gbps of throughput.

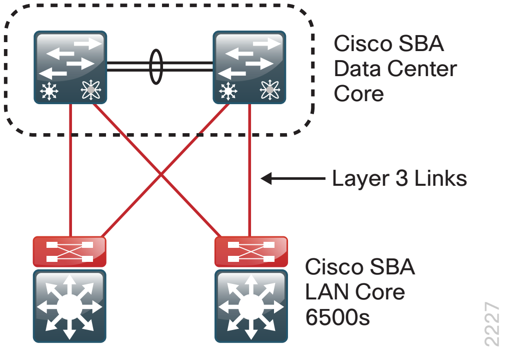

Step 1: On the first data center core Cisco Nexus 5500UP switch, configure two point-to-point Layer 3 interfaces.

```
interface Ethernet1/19
  description Core-1 Ten4/7
  no switchport
  ip address 10.4.40.50/30
  ip router eigrp 100
  ip pim sparse-mode

interface Ethernet1/20
  description Core-2 Ten4/7
  no switchport
  ip address 10.4.40.58/30
  ip router eigrp 100
  ip pim sparse-mode
```

Step 2: On the second data center core Cisco Nexus 5500UP switch, configure two point-to-point Layer 3 interfaces.

```
interface Ethernet1/19
  description Core-1 Ten4/8
  no switchport
  ip address 10.4.40.54/30
  ip router eigrp 100
  ip pim sparse-mode

interface Ethernet1/20
  description Core-2 Ten4/8
  no switchport
  ip address 10.4.40.62/30
  ip router eigrp 100
  ip pim sparse-mode
```

Step 3: On the Cisco SBA LAN Core 6500 switches, configure the four corresponding point-to-point Layer 3 links.

- On the first Cisco Catalyst LAN core switch, configure two links.

```
interface TenGigabitEthernet4/7
 description DC5548a Eth1/19
 no switchport
 ip address 10.4.40.49 255.255.255.252
 ip pim sparse-mode
 macro apply EgressQoS

interface TenGigabitEthernet4/8
 description DC5548b Eth1/19
 no switchport
 ip address 10.4.40.53 255.255.255.252
 ip pim sparse-mode
 macro apply EgressQoS
```

- On the second Cisco Catalyst LAN core switch, configure two links.

```
interface TenGigabitEthernet4/7
 description DC5548a Eth1/20
 no switchport
 ip address 10.4.40.57 255.255.255.252
 ip pim sparse-mode
 macro apply EgressQoS

interface TenGigabitEthernet4/8
 description DC5548b Eth1/20
 no switchport
 ip address 10.4.40.61 255.255.255.252
 ip pim sparse-mode
 macro apply EgressQoS
```

At this point, you should be able to see the IP routes from the rest of the network on the core Cisco Nexus 5500UP switches.

Procedure 10 Configure management switch connection

The first process of this “Ethernet Infrastructure” chapter covered deploying an out-of-band Ethernet management switch. In that process, you configured the switch for Layer 2 operation and uplinks to the data center core as the option of providing Layer 3 access to the management VLAN to provide access beyond the data center. If you have selected this option to provide Layer 3 access to the out-of-band Ethernet VLAN, follow this procedure to program the uplinks and the Layer 3 SVI on the Cisco Nexus 5500UP switches.

For resiliency, the Ethernet out-of-band management switch will be dual-homed to each of the data center core switches by using a vPC port channel.

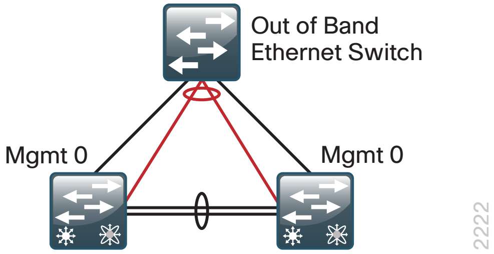

Step 1: Configure the Ethernet out-of-band management VLAN. You will configure the same values on each data center core Cisco Nexus 5500UP switch.

```
vlan 163
  name DC_Management
```

Step 2: Configure vPC port channel to the Ethernet management switch. You will configure the same values on each data center core Cisco Nexus 5500UP switch.

```
interface port-channel21
  description Link to Management Switch for VL163
  switchport mode trunk
  switchport trunk allowed vlan 163
  speed 1000
  vpc 21
```

Step 3: Configure the physical ports to belong to the port channel. You will configure the same values on each data center core Cisco Nexus 5500UP switch.

```
interface Ethernet1/21
  description Link to Management Switch for VL163
  switchport mode trunk
  switchport trunk allowed vlan 163
  speed 1000
  channel-group 21 mode active
```

Step 4: Configure an SVI interface for VLAN 163.

- Configure the first data center core Cisco Nexus 5500UP switch.

```
 interface Vlan163
  description DC-Management
  no ip redirects
  ip address 10.4.63.2/24
  ip router eigrp 100
  ip passive-interface eigrp 100
  hsrp 163
    priority 110
    ip 10.4.63.1
no shutdown
```

- Configure the second data center core Cisco Nexus 5500UP switch.

```
interface Vlan163
  description DC-Management
  no ip redirects
  ip address 10.4.63.3/24
  ip router eigrp 100
  ip passive-interface eigrp 100
  hsrp 163
    ip 10.4.63.1
no shutdown
```

Procedure 11 Configure Fabric Extender connections

Cisco Fabric Extender (FEX) ports are designed to support end host connectivity. There are some design rules to be aware of when connecting devices to Cisco FEX ports:

- Cisco FEX ports **do not** support connectivity to LAN switches that generate spanning-tree BPDU packets. If a Cisco FEX port receives a BPDU packet, it will shut down with an Error Disable status.
- Cisco FEX ports **do not** support connectivity to Layer 3 routed ports where routing protocols are exchanged with the Layer 3 core; they are only for Layer 2–connected end hosts or appliances.
- The Cisco Nexus 5500UP switch running Layer 3 routing supports a maximum of sixteen connected Cisco FEX on a switch.
- Cisco Fabric Extender connections are also configured as port channel connections on Cisco Nexus 5500 Series for uplink resiliency and load sharing.
- If the Cisco FEX is to be single-homed to only one member of the switch pair, it is configured as a standard port channel.
- If the Cisco FEX is to be dual-homed to both members of the vPC switch pair to support single-homed servers or for increased resiliency, it is configured as a vPC on the port channel. Every end node or server connected to a dual-homed FEX is logically dual homed to each of the Cisco Nexus 5500 core switches and will have a vPC automatically generated by the system for the Ethernet FEX edge port.

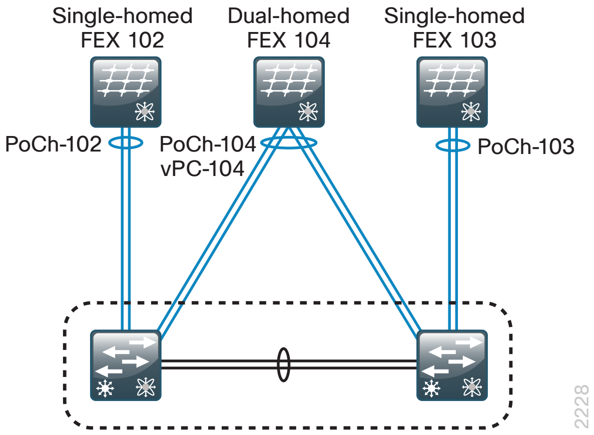

When assigning Cisco FEX numbering, you have the flexibility to use a numbering scheme (different from the example) that maps to some other identifier, such as a rack number that is specific to your environment.

Option 1. Configure single-homed FEX

A single-homed FEX requires configuration for the FEX and the uplinks on the Cisco Nexus 5500 switch it is connected to.

Step 1: Assign the physical interfaces on the connected Cisco Nexus 5500 switch to the port channels that are the supporting Cisco FEX attachment. These Ethernet interfaces form the uplink port channel to the connected FEX.

```
interface Ethernet1/13
  channel-group 102
!
interface Ethernet1/14
  channel-group 102
```

Step 2: Configure port channel interfaces to support the single-homed FEX attachment. The **switchport mode fex-fabric** command informs the Cisco Nexus 5500UP Series switch that a fabric extender should be at the other end of this link.

```
interface port-channel102
  description single-homed 2248
  switchport mode fex-fabric
  fex associate 102
```

Step 3: Configure the second single-homed FEX to the second Cisco Nexus 5500 switch.

```
interface Ethernet1/13
  channel-group 103
!
interface Ethernet1/14
  channel-group 103
!
interface port-channel103
  description single-homed 2248
  switchport mode fex-fabric
  fex associate 103
```

Option 2. Configure dual-homed FEX

A dual-homed FEX requires configuration for the FEX and the uplinks on both of the Cisco Nexus 5500 switches it is connected to.

Step 1: Assign the physical interfaces on the first connected Cisco Nexus 5500 switch to the port channels that are the supporting Cisco FEX attachment. These Ethernet interfaces form the uplink port channel to the connected FEX.

```
interface Ethernet1/25
  channel-group 104
!
interface Ethernet1/26
  channel-group 104
```

Step 2: Configure port channel interfaces on the first connected Cisco Nexus 5500 switch to support the dual-homed Cisco FEX attachment. The **switchport mode fex-fabric** command informs the Cisco Nexus 5500UP Series switch that a fabric extender should be at the other end of this link. The **vpc** command creates the dual-homed port-channel for the dual-homed FEX.

```
interface port-channel104
  description dual-homed 2232
  switchport mode fex-fabric
  fex associate 104
  vpc 104
```

Step 3: Repeat the configuration on the second connected Cisco Nexus 5500 switch.

```
interface Ethernet1/25
  channel-group 104
!
interface Ethernet1/26
  channel-group 104
!
interface port-channel104
  description dual-homed 2232
  switchport mode fex-fabric
  fex associate 104
  vpc 104
```

After configuration is completed for either FEX attachment model, you can power up the FEX and verify the status of the fabric extender modules by using the **show fex** command and then looking for the state of “online” for each unit.

```
dc5548ax# show fex
FEX     FEX         FEX                               FEX
Number  Description State   Model                 Serial
---------------------------------------------------------
102     FEX0102     Online  N2K-C2248TP-1GE    SSI14140643
104     FEX0104     Online  N2K-C2232PP-10GE SSI142602QL
```

Tech Tip

It may take a few minutes for the Cisco FEX to come online after it is programmed, because the initial startup of the Cisco FEX downloads operating code from the connected Cisco Nexus switch.

Procedure 12 Configure end node ports

When configuring Cisco Nexus FEX Ethernet ports for server or appliance connectivity, you must configure the port on one or both of the Cisco Nexus 5500UP core switches depending on the FEX connection (single-homed or dual-homed).

Option 1. Single-homed server to dual-homed FEX

Because the server is connected to a dual-homed FEX, this configuration must be done on both Cisco Nexus 5500UP data center core switches. The spanning-tree mode, VLAN list, and other characteristics for the Ethernet port should be identically programmed on each Cisco Nexus 5500UP switch.

Step 1: When connecting a single-homed server to a dual-homed Cisco FEX, assign physical interfaces to support servers or devices that belong in a single VLAN as access ports. Setting the spanning-tree port type to **edge** allows the port to provide immediate connectivity on the connection of a new device. Enable QoS classification for the connected server or end node as defined in Procedure 3”Configure QoS policies”.

Example

```
interface Ethernet103/1/1
  switchport access vlan 163
  spanning-tree port type edge
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-
QOS
```

Tech Tip

You must assign the Ethernet interface configuration on both data center core Cisco Nexus 5500UP switches as the host is dual homed because it is on a dual-homed Cisco FEX. Failure to configure the port on both Nexus 5500 switches with matching VLAN assignment will prevent the Ethernet interface from being activated.

Step 2: When connecting a single-homed server to a dual-homed Cisco FEX, assign physical interfaces to support servers or devices that require a VLAN trunk interface to communicate with multiple VLANs. Most virtualized servers will require trunk access to support management access plus user data for multiple virtual machines. Setting the spanning-tree port type to **edge** allows the port to provide immediate connectivity on the connection of a new device. Enable QoS classification for the connected server or end node as defined in Procedure 3, "Configure QoS policies".

Because the server is connected to a dual-homed FEX, this configuration must be done on both Cisco Nexus 5500UP data center core switches.

Example

```
interface Ethernet103/1/2
  switchport mode trunk
  switchport trunk allowed vlan 148-163
  spanning-tree port type edge trunk
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
```

Option 2. Dual-homed server using EtherChannel to two single-homed FEX

Because the server is dual-homed using vPC EtherChannel, this configuration must be done on both Cisco Nexus 5500UP data center core switches.

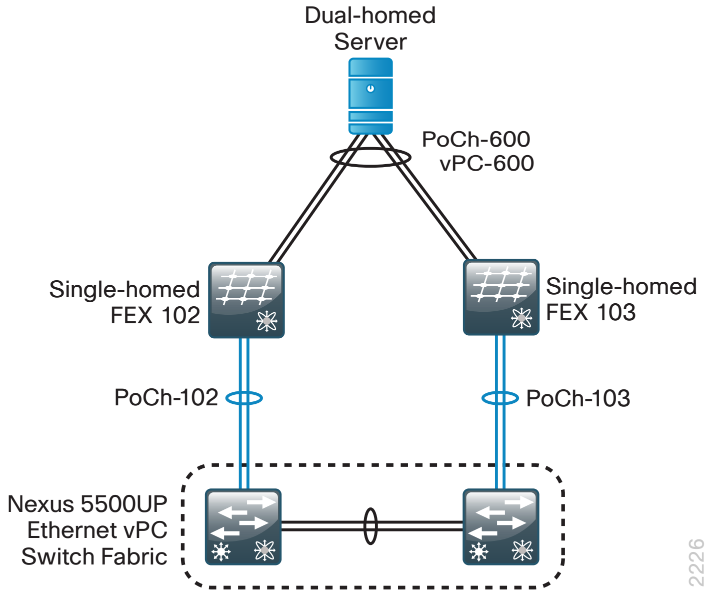

When connecting a dual-homed server that is using IEEE 802.3ad EtherChannel from the server to a pair of single-homed Cisco FEX, you must configure the Cisco FEX Ethernet interface as a port channel and assign a vPC interface to the port channel to talk EtherChannel to the attached server.

Example

Step 1: On the first Cisco Nexus 5500 switch.

```
interface ethernet102/1/1
  switchport mode trunk
  switchport trunk allowed vlan 148-163
  spanning-tree port type edge trunk
  channel-group 600
  no shutdown
interface port-channel 600
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 600
  no shutdown
```

Step 2: On the second Cisco Nexus 5500 switch.

```
interface ethernet103/1/1
  switchport mode trunk
  switchport trunk allowed vlan 148-163
  spanning-tree port type edge trunk
  channel-group 600
  no shutdown
interface port-channel 600
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 600
  no shutdown
```

Tech Tip

When connecting ports via vPC, Cisco NX-OS does consistency checks to make sure that the VLAN list, spanning-tree mode, and other characteristics match between the ports configured on each switch that make up a vPC. If the configuration for each port is not identical with the other, the port will not come up.

Option 3. Dual-homed server using EtherChannel to two dual-homed FEX

This connectivity option, referred to as enhanced vPC, requires Cisco NX-OS release 5.1(3)N1(1) or later for the Cisco Nexus 5500 switches. Dual-homing a server with EtherChannel to a dual-homed FEX is not supported on the older Cisco Nexus 5000 switches.

When connecting a dual-homed server that is using IEEE 802.3ad EtherChannel from the server to a pair of dual-homed Cisco FEX, you must configure the Ethernet interface on each of the Cisco FEX interfaces as a port channel but not as a vPC. The Cisco Nexus 5500 switches will automatically create a vPC to track the dual-homed port channel.

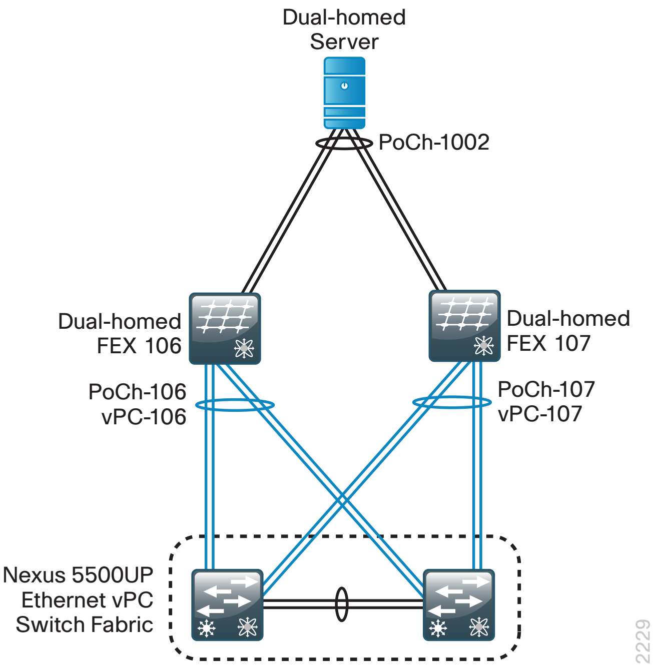

In this configuration option, you use FEX numbers 106 and 107. Both FEX would have to be configured as dual-homed to the Cisco Nexus 5500 data center core switches as defined in Option 2: "Configure dual-homed FEX".

Step 1: Configure the Ethernet interfaces of the first dual-homed FEX on the first core Cisco Nexus 5500 switch for a port channel to the server.

```
interface ethernet106/1/3-4
  channel-group 1002
```

Step 2: Configure the Ethernet interfaces of the second dual-homed FEX on the first core Cisco Nexus 5500 switch for a port channel to the server.

```
interface ethernet107/1/3-4
  channel-group 1002
```

Step 3: Configure the port-channel for the VLANs to be supported.

```
interface port-channel 1002
  switchport mode trunk
  switchport trunk allowed vlan 148-163
  spanning-tree port type edge trunk
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-
QOS
```

Step 4: Repeat the commands on the second core Cisco Nexus 5500 switch with the same settings, because the server and the FEX are dual-homed.

```
interface ethernet106/1/3-4
  channel-group 1002
!
interface ethernet107/1/3-4
  channel-group 1002
!
interface port-channel 1002
  switchport mode trunk
  switchport trunk allowed vlan 148-163
  spanning-tree port type edge trunk
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-
QOS
```

Storage Infrastructure

Business Overview

There is a constant demand for more storage in organizations today. Storage for servers can be physically attached directly to the server or connected over a network. Direct attached storage (DAS) is physically attached to a single server and is difficult to use efficiently because it can be used only by the host attached to it. Storage area networks (SANs) allow multiple servers to share a pool of storage over a Fibre Channel or Ethernet network. This capability allows storage administrators to easily expand capacity for servers supporting data-intensive applications.

Technology Overview

IP-based Storage Options

Many storage systems provide the option for access using IP over the Ethernet network. This approach allows a growing organization to gain the advantages of centralized storage without needing to deploy and administer a separate Fibre Channel network. Options for IP-based storage connectivity include Internet Small Computer System Interface (iSCSI) and network attached storage (NAS).

iSCSI is a protocol that enables servers to connect to storage over an IP connection and is a lower-cost alternative to Fibre Channel. iSCSI services on the server must contend for CPU and bandwidth along with other network applications, so you need to ensure that the processing requirements and performance are suitable for a specific application. iSCSI has become a storage technology that is supported by most server, storage, and application vendors. iSCSI provides block-level storage access to raw disk resources, similar to Fibre Channel. NICs also can provide support to offload iSCSI to a separate processor to increase performance.

Network attached storage (NAS) is a general term used to refer to a group of common file access protocols, the most common implementations use Common Internet File System (CIFS) or network file server (NFS). CIFS originated in the Microsoft network environment and is a common desktop file-sharing protocol. NFS is a multi-platform protocol that originated in the UNIX environment and can be used for shared hypervisor storage. Both NAS protocols provide file-level access to shared storage resources.

Most organizations will have applications for multiple storage access technologies—for example, Fibre Channel for the high performance database and production servers, and NAS for desktop storage access.

Fibre Channel Storage

Fibre Channel allows servers to connect to storage across a fiber-optic network, across a data center, or even across a WAN by using Fibre Channel over IP. Multiple servers can share a single storage array.

This Cisco SBA data center design uses the Cisco Nexus 5500UP series switches as the core that provides Fibre Channel and Fibre Channel over Ethernet (FCoE) SAN switching. The Cisco Nexus 5500UP offers the density required for collapsed Fibre Channel connectivity requirements by supporting both Fibre Channel and FCoE servers and storage arrays. The Cisco MDS 9148 Multilayer Fabric Switch is ideal for a larger SAN fabric with up to 48 Fibre Channel ports, providing 48 line-rate 8-Gbps Fibre Channel ports and cost-effective scalability. The Cisco MDS family of Multilayer SAN Fabric Switches also offers options like hardware-based encryption services, tape acceleration, and Fibre Channel over IP for longer distance SAN extension.

In a SAN, a fabric consists of servers and storage connected to a Fibre Channel switch (Figure 13). It is standard practice in SANs to create two completely separate physical fabrics, providing two distinct paths to the storage. Fibre Channel fabric services operate independently on each fabric so when a server needs resilient connections to a storage array, it connects to two separate fabrics. This design prevents failures or misconfigurations in one fabric from affecting the other fabric.

Figure 13 - Dual fabric SAN with a single disk array

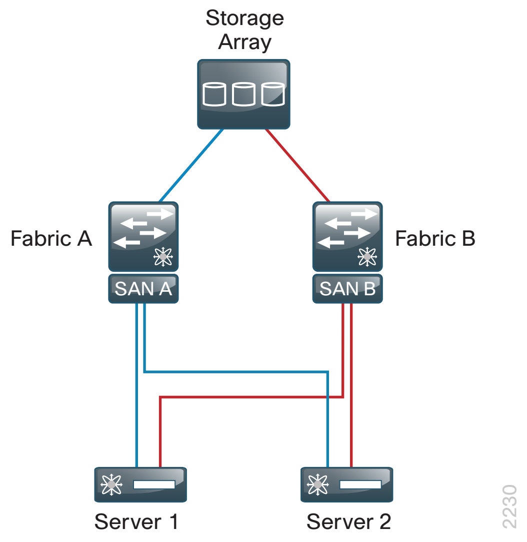

Each server or host on a SAN connects to the Fibre Channel switch with a multi-mode fiber cable from a host bus adapter (HBA). For resilient connectivity, each host connects a port to each of the fabrics.

Each port has a port worldwide name (pWWN), which is the port's address that uniquely identifies it on the network. An example of a pWWN is: 10:00:00:00:c9:87:be:1c. In data networking this would compare to a MAC address for an Ethernet adapter.

VSANs

The virtual storage area network (VSAN) is a technology created by Cisco that is modeled after the virtual local area network (VLAN) concept in Ethernet networks. VSANs provide the ability to create many logical SAN fabrics on a single Cisco MDS 9100 Family switch. Each VSAN has its own set of services and address space, which prevents an issue in one VSAN from affecting other VSANs. In the past, it was a common practice to build physically separate fabrics for production, backup, lab, and departmental environments. VSANs allow all of these fabrics to be created on a single physical switch with the same amount of protection provided by separate switches.

Zoning

The terms *target* and *initiator* will be used throughout this section. *Targets* are disk or tape devices. *Initiators* are servers or devices that initiate access to disk or tape.

Zoning provides a means of restricting visibility and connectivity among devices connected to a SAN. The use of zones allows an administrator to control which initiators can see which targets. It is a service that is common throughout the fabric, and any changes to a zoning configuration are disruptive to the entire connected fabric.

Initiator-based zoning allows for zoning to be port-independent by using the world wide name (WWN) of the end host. If a host's cable is moved to a different port, it will still work if the port is a member of the same VSAN.

Device Aliases

When configuring features such as zoning, quality of service (QoS), and port security on a Cisco MDS 9000 Family switch, WWNs must be specified. The WWN naming format is cumbersome, and manually typing WWNs is error prone. Device aliases provide a user-friendly naming format for WWNs in the SAN fabric (for example: "p3-c210-1-hba0-a" instead of "10:00:00:00:c9:87:be:1c").

Use a naming convention that makes initiator and target identification easy. For example, p3-c210-1-hba0-a in this setup identifies:

- Rack location: p3
- Host type: c210
- Host number: 1
- HBA number: hba0
- Port on HBA: a

Storage Array Tested

The storage arrays used in the testing and validation of this deployment guide are the EMC CX4-120 and the NetApp FAS3200. The specific storage array configuration may vary. Please consult the installation instructions from the specific storage vendor. The Cisco interoperability support matrix for Fibre Channel host bus adapters and storage arrays can be found at: http://www.cisco.com/en/US/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx.html

Tech Tip

Specific interfaces, addresses, and device aliases are examples from the lab. Your WWN addresses, interfaces, and device aliases will likely be different.

Deployment Details

Deployment examples documented in this section include:

- Configuration of Cisco Nexus 5500UP–based SAN network to support Fibre Channel–based storage.
- Configuration of a Cisco MDS SAN switch for higher-density Fibre Channel environments.
- FCoE access to storage from Cisco UCS C-Series servers using Cisco Nexus 5500.

Process

Configuring Fibre Channel SAN on Cisco Nexus 5500UP

1. Configure Fibre Channel operation
2. Configure VSANs
3. Configure Fibre Channel ports
4. Configure device aliases
5. Configure zoning
6. Verify the configuration

Complete each of the following procedures to configure the Fibre Channel SAN on the data center core Cisco Nexus 5500UP switches.

Procedure 1 Configure Fibre Channel operation

The Cisco Nexus 5500UP switch has universal ports that are capable of running Ethernet+FCoE or Fibre Channel on a per-port basis. By default, all switch ports are enabled for Ethernet operation. Fibre Channel ports must be enabled in a contiguous range and be the high numbered ports of the switch baseboard and/or the high numbered ports of a universal port expansion module.

Reader Tip

The first part of this procedure was outlined in Procedure 2 in the “Configuring the Data Center Core” process in the “Ethernet Infrastructure” chapter of this deployment guide. If you have already configured ports for Fibre Channel operation, you can skip Step 1 through Step 3 of this procedure.

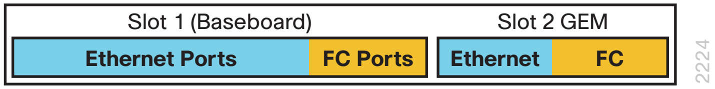

In this design, you enable ports 28 through 32 on the Cisco Nexus 5548UP switch as Fibre Channel ports.

Step 1: Configure universal port mode for Fibre Channel.

```
slot 1
  port 28-32 type fc
```

Tech Tip

Changing port type to **fc** requires a reboot in Cisco Nexus 5500UP version 5.1(3)N1(1a) software to recognize the new port operation. This is subject to change in later releases of software. Ports will not show up in the configuration as fc ports if you did not previously enable the FCoE feature.

Step 2: If you are changing the port type at this time, save your configuration and reboot the switch so that the switch recognizes the new **fc** port type operation. If you have already done this, there is no need to reboot.

Step 3: If you have not done so, enable FCOE operation, which enables both native Fibre Channel and FCoE operation.

```
feature fcoe
```

Step 4: Enable SAN port-channel trunking operation and Fibre Channel N-Port ID Virtualization for connecting to Cisco UCS fabric interconnects.

```
feature npiv
feature fport-channel-trunk
```

Reader Tip

More detail for connecting to a Cisco UCS B-Series fabric interconnect for Fibre Channel operation can be found in the *Cisco SBA—Data Center Unified Computing System Deployment Guide.*

Procedure 2 Configure VSANs

Cisco Data Center Network Manager (DCNM) for SAN Essentials Edition is a no-cost application to configure and manage Cisco MDS and Cisco Nexus SAN switches, available for download from http://www.cisco.com. DCNM for SAN Essentials includes Cisco MDS Device Manager and Cisco SAN Fabric Manager. Managing more than one switch at the same time requires a licensed version.

To manage a switch with Cisco DCNM Device Manager, connect to the switch's management IP address. The CLI can also be used to configure Fibre Channel operation.

Java runtime environment (JRE) is required to run Cisco DCNM Fabric Manager and Device Manager; and should be installed on your desktop before accessing either application.

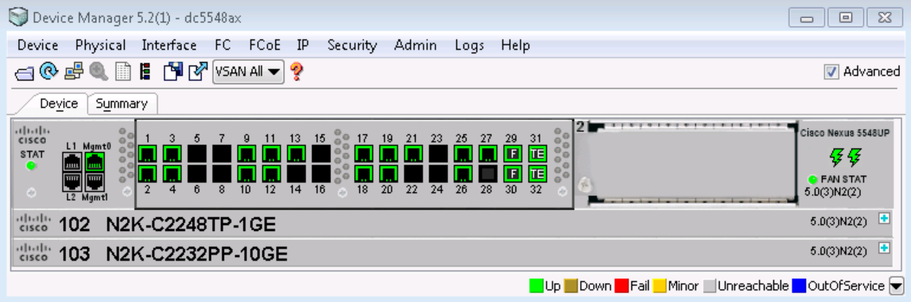

By default, all ports are assigned to VSAN 1 at initialization of the switch. It is a best practice to create a separate VSAN for production and to leave VSAN 1 for unused ports. By not using VSAN 1, you can avoid future problems with merging of VSANs when combining other existing switches that may be set to VSAN 1.

Fibre Channel operates in a SAN-A and SAN-B approach, where you create two separate SAN fabrics. Fibre Channel hosts and targets connect to both fabrics for redundancy. The SAN fabrics operate in parallel.

The example below describes creating two VSANs, one on each data center core Cisco Nexus 5500UP switch.

You can use the CLI or Device Manager to create a VSAN.

Step 1: Install Cisco DCNM for SAN Essentials.

Step 2: Using DCNM Device Manager, connect to the first Cisco Nexus data center core switch IP address (10.4.63.10).

Step 3: Using Device Manager, click **FC > VSANS**. The Create VSAN General window appears.

Step 4: In the **VSAN id** list, choose **4**, and in the name box, type General-Storage.

Step 5: Next to the Interface Members box, click the ellipsis (**...**) button.

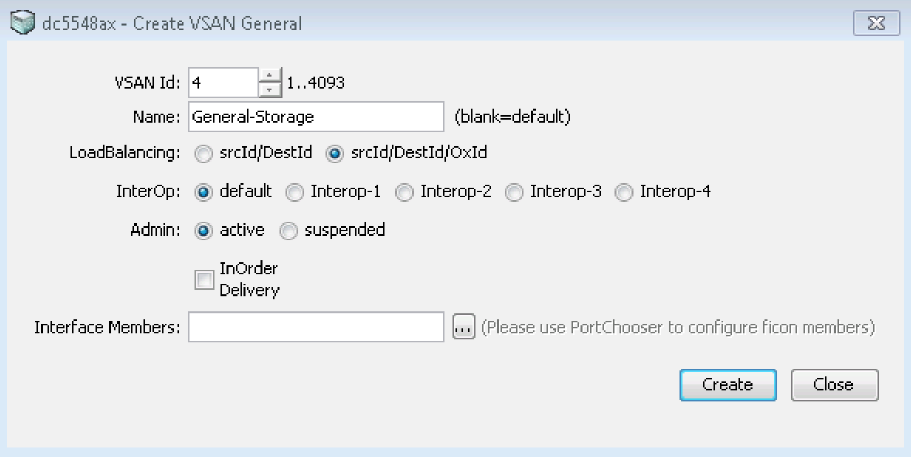

dc5548ax - Create VSAN General

VSAN Id: 4 1..4093

Name: General-Storage (blank=default)

LoadBalancing: ○ srcId/DestId ● srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

□ InOrder Delivery

Interface Members: ... (Please use PortChooser to configure ficon members)

Step 6: Select the interface members by clicking the port numbers you want.

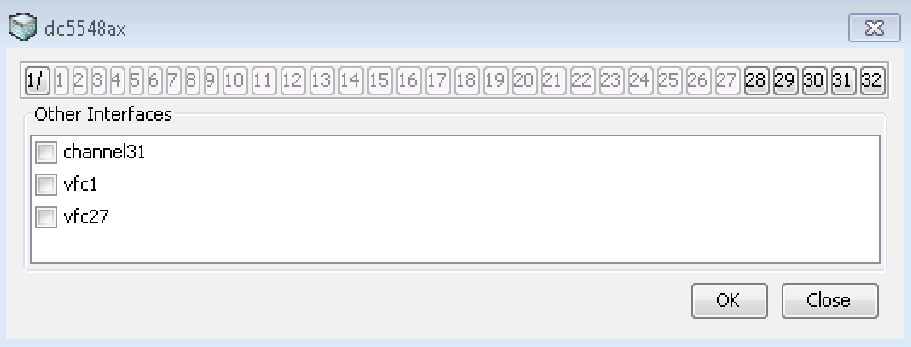

dc5548ax

1/ 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32

Other Interfaces

□ channel31

□ vfc1

□ vfc27

Step 7: Click **Create.** The VSAN is created. You can add additional VSAN members in the Membership tab of the main VSAN window.

The preceding steps apply this configuration in CLI.

```
vsan database
vsan 4 name "General-Storage"
vsan 4 interface fc1/28
```

Step 8: Repeat the steps in this procedure to create a VSAN 5 on the second Cisco Nexus 5500UP switch. Use the same VSAN name.

Procedure 3 Configure Fibre Channel ports

By default, the ports are configured for port mode **Auto**, and this setting should not be changed for most devices that are connected to the fabric. However, you will have to assign a VSAN to the port.

Step 1: If you want to change the port mode by using Device Manager, right-click the port you want to configure.

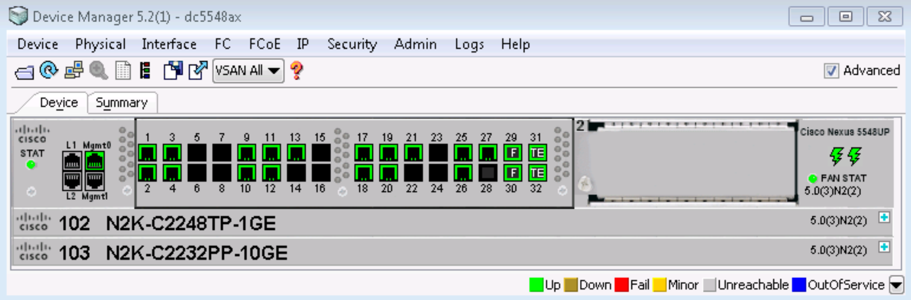
Device Manager 5.2(1) - dc5548ax
Device Physical Interface FC FCoE IP Security Admin Logs Help
VSAN All
☑ Advanced
Device Summary
cisco STAT
L1 Mgmt0 L2 Mgmt1
1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31
2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32
F TE F TE
2
Cisco Nexus 5548UP
FAN STAT
5.0(3)N2(2)
102 N2K-C2248TP-1GE 5.0(3)N2(2)
103 N2K-C2232PP-10GE 5.0(3)N2(2)
🟩Up 🟫Down 🟥Fail 🟨Minor ⬜Unreachable 🟦OutOfService

The General tab appears.

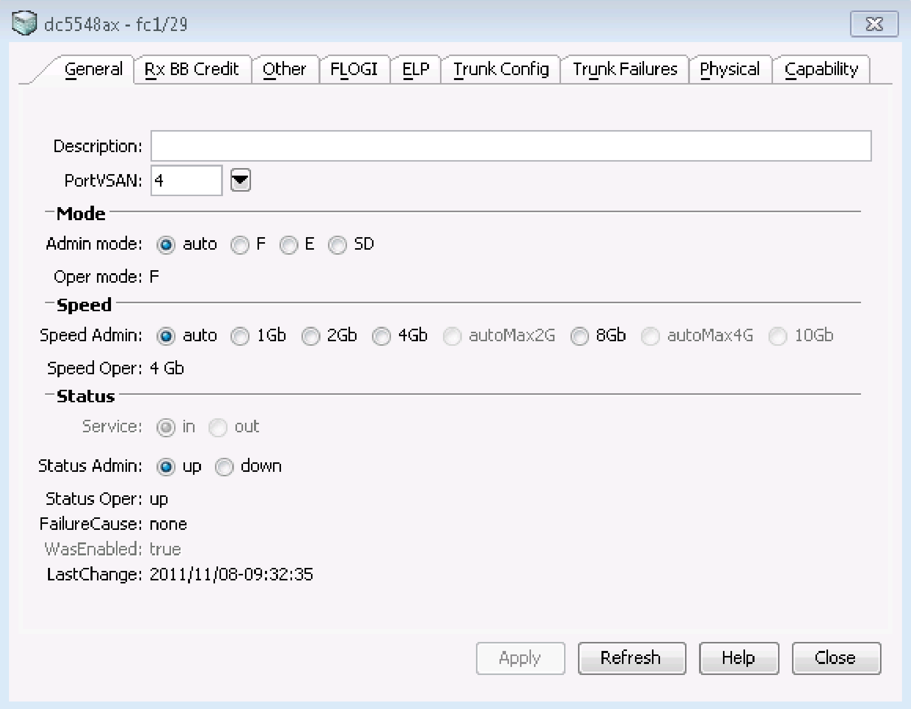
dc5548ax - fc1/29

General | Rx BB Credit | Other | FLOGI | ELP | Trunk Config | Trunk Failures | Physical | Capability

Description:

PortVSAN: 4

Mode

Admin mode: ● auto ○ F ○ E ○ SD

Oper mode: F

Speed

Speed Admin: ● auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb

Speed Oper: 4 Gb

Status

Service: ● in ○ out

Status Admin: ● up ○ down

Status Oper: up

FailureCause: none

WasEnabled: true

LastChange: 2011/11/08-09:32:35

You can see in this figure that the PortVSAN assignment is listed in the top left of the General tab.

Step 2: Next to Status Admin, select **up**. This enables the port.

Step 3: In the PortVSAN list, choose **4** or **5**, depending on which switch you are working on, and then click **Apply**. This changes the VSAN and activates the ports.

The preceding steps apply this configuration in CLI.

```
vsan database
   vsan 4 interface fc1/28
```

This step assigns ports to a VSAN similar to Step 5 in the previous procedure of "Configure VSANs". If you have already created VSANs, you can use this as another way to assign a port to a VSAN.

Step 4: Connect Fibre Channel devices to ports.

Reader Tip

For more information about preparing Cisco UCS B-Series and C-Series servers for connecting to the Fibre Channel network see the *Cisco SBA—Data Center Unified Computing System Deployment Guide.*

Step 5: Display fabric login (FLOGI) by entering the **show flogi database** on the switch CLI.

Tech Tip

When the initiator or target is plugged in or starts up, it automatically logs into the fabric. Upon login, the initiator or target WWN is made known to the fabric. Until you have storage arrays or servers with active HBAs plugged into the switch on Fibre Channel ports, you will not see entries in the FLOGI database.

Example

```
dc5548ax# show flogi database
-----------------------------------------------------------------------
INTERFACE VSAN FCID         PORT NAME           NODE NAME
-----------------------------------------------------------------------
fc1/29    4   0xbc0002  20:41:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
fc1/29    4   0xbc0005  20:00:00:25:b5:77:77:9f  20:00:00:25:b5:00:77:9f
fc1/30    4   0xbc0004  20:42:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
vfc1      4   0xbc0000  20:00:58:8d:09:0e:e0:d2  10:00:58:8d:09:0e:e0:d2
vfc27     4   0xbc0006  50:0a:09:81:89:3b:63:be  50:0a:09:80:89:3b:63:be

Total number of flogi = 5.
```

Procedure 4 Configure device aliases

Device aliases map the long WWNs for easier zoning and identification of initiators and targets. An incorrect device name may cause unexpected results. Device aliases can be used for zoning, port-security, QoS, and **show** commands.

You can configure device aliases via Device Manager or CLI.

Option 1. Configure device aliases by using Device Manager

Step 1: In Device Manager, access the Device Alias window by navigating to **FC > Advanced > Device Alias**.

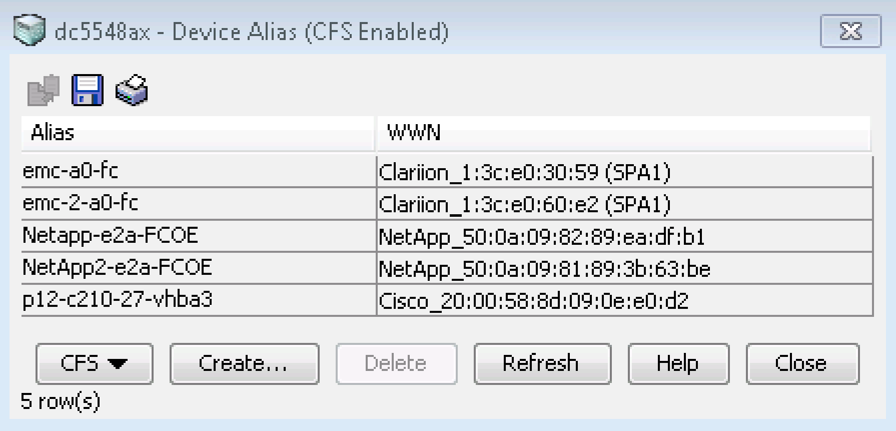

Alias	WWN
emc-a0-fc	Clariion_1:3c:e0:30:59 (SPA1)
emc-2-a0-fc	Clariion_1:3c:e0:60:e2 (SPA1)
Netapp-e2a-FCOE	NetApp_50:0a:09:82:89:ea:df:b1
NetApp2-e2a-FCOE	NetApp_50:0a:09:81:89:3b:63:be
p12-c210-27-vhba3	Cisco_20:00:58:8d:09:0e:e0:d2

5 row(s)

Step 2: Click **Create**.

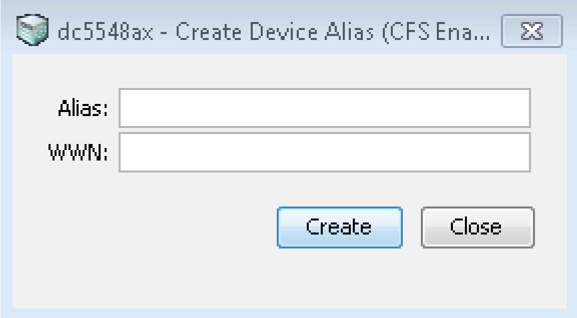

Step 3: In the Alias box, enter a name, and in the **WWN** box, paste in or type the WWN of the host, and then click **Create**.

Step 4: After you have created your devices aliases, click **CFS > Commit**. The changes are written to the database.

Option 2. Configure device aliases by using CLI

Step 1: Enter device alias database configuration mode.

```
device-alias database
```

Step 2: Enter device alias names mapping to a PWWN from the FLOGI database above. As an example:

```
 device-alias name emc-a0-fc pwwn 50:06:01:61:3c:e0:30:59
 device-alias name emc-2-a0-fc pwwn 50:06:01:61:3c:e0:60:e2
 device-alias name Netapp-e2a-FCOE pwwn
50:0a:09:82:89:ea:df:b1
 device-alias name NetApp2-e2a-FCOE pwwn
50:0a:09:81:89:3b:63:be
 device-alias name p12-c210-27-vhba3 pwwn
20:00:58:8d:09:0e:e0:d2
```

Step 3: Exit device alias configuration mode.

```
exit
```

Step 4: Commit the changes.

```
device-alias commit
```

Step 5: Enter the **show flogi database** command. Aliases are now visible.

```
dc5548ax# show flogi database
------------------------------------------------------------------------
INTERFACE  VSAN  FCID  PORT NAME                NODE NAME
------------------------------------------------------------------------
fc1/29    4  0xbc0002  20:41:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
fc1/29    4  0xbc0005  20:00:00:25:b5:77:77:9f  20:00:00:25:b5:00:77:9f
fc1/30    4  0xbc0004  20:42:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
vfc1      4  0xbc0000  20:00:58:8d:09:0e:e0:d2  10:00:58:8d:09:0e:e0:d2
                       [p12-c210-27-vhba3]
vfc27     4  0xbc0006  50:0a:09:81:89:3b:63:be  50:0a:09:80:89:3b:63:be
                       [NetApp2-e2a-FCOE]
```

Procedure 5 Configure zoning

Leading practices for zoning:

- Configure zoning between a single initiator and a single target per zone.
- You can also configure a single initiator to multiple targets in the same zone.
- Zone naming should follow a simple naming convention of *initiator_x_target_x*:
  - p12-ucs-b-fc0-vhba1_emc-2
- Limit zoning to a single initiator with a single target or multiple targets to help prevent disk corruption and data loss.

Zoning can be configured from the CLI and from Cisco DCNM for SAN Fabric Manager.

Option 1. Configure a zone by using CLI

Step 1: In configuration mode, enter the zone name and VSAN number.

```
zone name p12-ucs-b-fc0-vhba1_emc-2 vsan 4
```

Step 2: Specify device members by WWN or device alias.

```
member device-alias emc-2-a0-fc
member pwwn 20:00:00:25:b5:77:77:9f
```

Step 3: Create and activate a zoneset.

```
zoneset name FCOE_4 vsan 4
```

Tech Tip

A *zoneset* is a collection of zones. Zones are members of a zoneset. After you add all the zones as members, you must activate the zoneset. There can only be one active zoneset per VSAN.

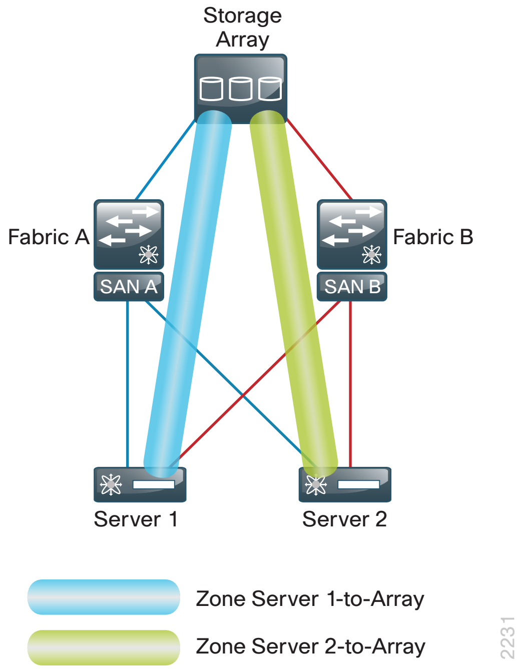

Step 4: Add members to the zoneset.

```
member p12-ucs-b-fc0-vhba1_emc-2
member p12-c210-27-vhba3_netapp-2-e2a
```

Step 5: After all the zones for VSAN 4 are created and added to the zoneset, activate the configuration.

```
zoneset activate name FCOE_4 vsan 4
```

Step 6: Distribute the zone database to other switches in the SAN. This prepares for expanding your Fibre Channel SAN beyond a single switch.

```
zoneset distribute full vsan 4
```

Option 2. Configure a zone by using Cisco DCNM

Step 1: Launch the Cisco DCNM for SAN Fabric Manager installed in Step 1 in the previous procedure of "Configure VSANs".

Step 2: Log in to DCNM-SAN manager. The default username is admin and the password is password.

Step 3: Choose a seed switch by entering the IP address of the first Cisco Nexus 5500UP switch (for example, 10.4.63.10), and then choosing Cisco Nexus 5500UP from the list.

Step 4: From the DCNM-SAN menu, choose **Zone**, and then click **Edit Local Full Zone Database**.

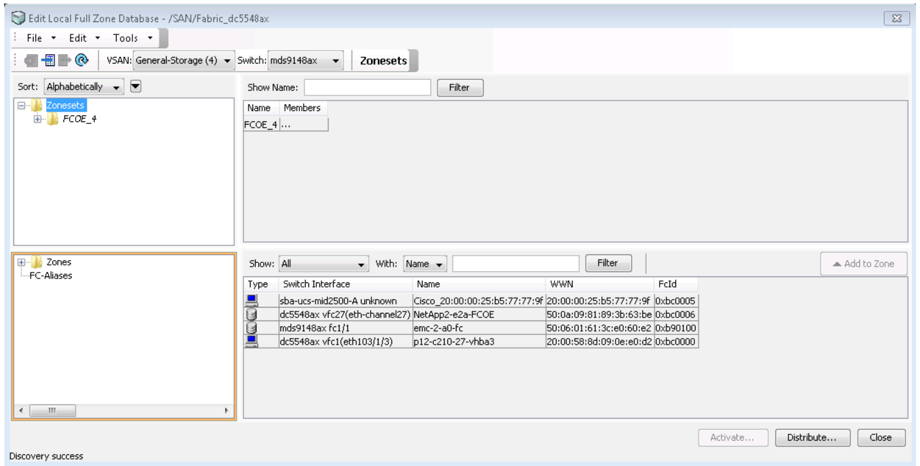

VSAN: General-Storage (4) Switch: mds9148ax

Sort: Alphabetically

Zonesets
- FCOE_4

Show Name:

Name	Members
FCOE_4	...

Zones
FC-Aliases

Show: All With: Name

Type	Switch Interface	Name	WWN	FcId
	sba-ucs-mid2500-A unknown	Cisco_20:00:00:25:b5:77:77:9f	20:00:00:25:b5:77:77:9f	0xbc0005
	dc5548ax vfc27(eth-channel27)	NetApp2-e2a-FCOE	50:0a:09:81:89:3b:63:be	0xbc0006
	mds9148ax fc1/1	emc-2-a0-fc	50:06:01:61:3c:e0:60:e2	0xb90100
	dc5548ax vfc1(eth103/1/3)	p12-c210-27-vhba3	20:00:58:8d:09:0e:e0:d2	0xbc0000

Discovery success

Step 5: In the Zone Database window, in the left pane, right-click **Zones**, and then click **Insert**. This creates a new zone.

Step 6: In the **Zone Name** box, enter the name of the new zone, and then click **OK**.

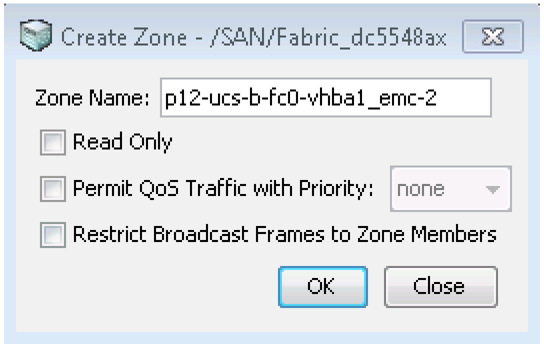

Zone Name: p12-ucs-b-fc0-vhba1_emc-2

□ Read Only

□ Permit QoS Traffic with Priority: none

□ Restrict Broadcast Frames to Zone Members

Step 7: Select the new zone, and then, from the bottom of the right-hand side of the database window, choose initiator or targets you want to add to the zone, and then click **Add to Zone**.

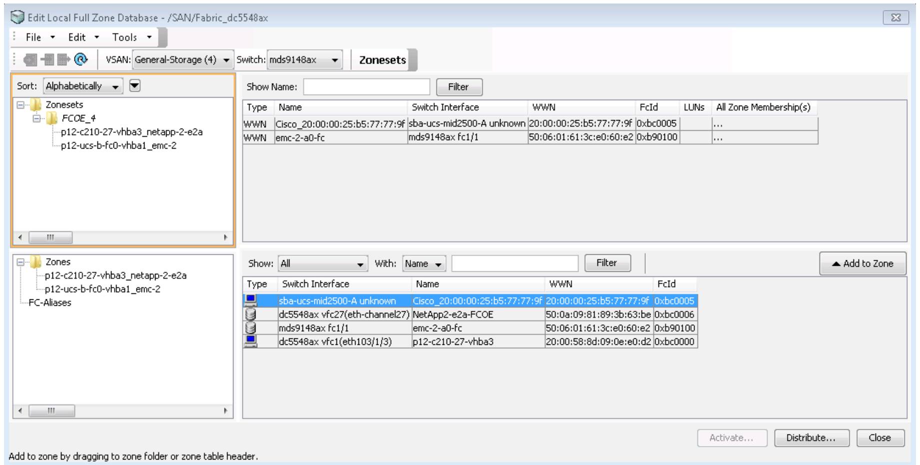

VSAN: General-Storage (4) Switch: mds9148ax

Sort: Alphabetically

Zonesets
- FCOE_4
  - p12-c210-27-vhba3_netapp-2-e2a
  - p12-ucs-b-fc0-vhba1_emc-2

Show Name:

Type	Name	Switch Interface	WWN	FcId	LUNs	All Zone Membership(s)
WWN	Cisco_20:00:00:25:b5:77:77:9f	sba-ucs-mid2500-A unknown	20:00:00:25:b5:77:77:9f	0xbc0005		...
WWN	emc-2-a0-fc	mds9148ax fc1/1	50:06:01:61:3c:e0:60:e2	0xb90100		...

Zones
- p12-c210-27-vhba3_netapp-2-e2a
- p12-ucs-b-fc0-vhba1_emc-2

FC-Aliases

Show: All With: Name

Type	Switch Interface	Name	WWN	FcId
	sba-ucs-mid2500-A unknown	Cisco_20:00:00:25:b5:77:77:9f	20:00:00:25:b5:77:77:9f	0xbc0005
	dc5548ax vfc27(eth-channel27)	NetApp2-e2a-FCOE	50:0a:09:81:89:3b:63:be	0xbc0006
	mds9148ax fc1/1	emc-2-a0-fc	50:06:01:61:3c:e0:60:e2	0xb90100
	dc5548ax vfc1(eth103/1/3)	p12-c210-27-vhba3	20:00:58:8d:09:0e:e0:d2	0xbc0000

Add to zone by dragging to zone folder or zone table header.

Step 8: Right-click **Zoneset** to insert a new zoneset.

Step 9: Drag the zones you just created from the zone box to the zoneset folder that you created.

Step 10: Click **Activate**. This activates the configured zoneset

Step 11: In the Save Configuration dialog box, select **Save Running to Startup Configuration**, and then click **Continue Activation**.

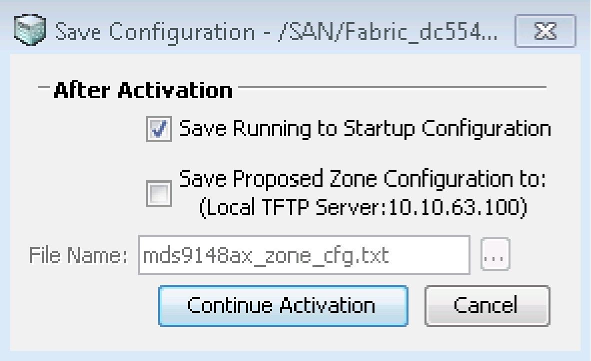

Save Configuration - /SAN/Fabric_dc554...

After Activation

☑ Save Running to Startup Configuration

☐ Save Proposed Zone Configuration to: (Local TFTP Server:10.10.63.100)

File Name: mds9148ax_zone_cfg.txt

Step 12: Configure SAN B the same way by using the procedures in this process to create VSAN 5 on the second data center core Cisco Nexus 5500UP switch.

Procedure 6 Verify the configuration

Step 1: Verify the Fibre Channel login.

In a Fibre Channel fabric, each host or disk requires a Fibre Channel ID (FC ID). When a fabric login (FLOGI) is received from the device, this ID is assigned by the fabric. If the required device is displayed in the FLOGI table, the fabric login is successful.

```
dc5548ax# show flogi database
-----------------------------------------------------------------
INTERFACE VSAN FCID  PORT NAME               NODE NAME
-----------------------------------------------------------------
fc1/29    4   0xbc0002  20:41:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
fc1/29    4   0xbc0005  20:00:00:25:b5:77:77:9f  20:00:00:25:b5:00:77:9f
fc1/30    4   0xbc0004  20:42:00:05:73:a2:b2:40  20:04:00:05:73:a2:b2:41
vfc1      4   0xbc0000  20:00:58:8d:09:0e:e0:d2  10:00:58:8d:09:0e:e0:d2
                        [p12-c210-27-vhba3]
vfc27     4   0xbc0006  50:0a:09:81:89:3b:63:be  50:0a:09:80:89:3b:63:be
                        [NetApp2-e2a-FCOE]

Total number of flogi = 5.
```

Step 2: Verify Fibre Channel Name Server (FCNS) attributes.

The FCNS database shows the same PWWN login along with vendor specific attributes and features. Check that your initiators and targets have logged in and show **FC4-TYPE:FEATURE** attributes as highlighted below. If the feature attributes do not show, you may have a part of the configuration on the end host or storage device misconfigured or a device driver issue.

```
dc5548ax# show fcns database

VSAN 4:
-----------------------------------------------------------------
FCID     TYPE  PWWN                     (VENDOR)    FC4-TYPE:FEATURE
-----------------------------------------------------------------
0xb90100 N     50:06:01:61:3c:e0:60:e2  (Clariion)  scsi-fcp:target
               [emc-2-a0-fc]
0xbc0000 N     20:00:58:8d:09:0e:e0:d2             scsi-fcp:init fc-gs
               [p12-c210-27-vhba3]
0xbc0002 N     20:41:00:05:73:a2:b2:40  (Cisco)     npv
0xbc0004 N     20:42:00:05:73:a2:b2:40  (Cisco)     npv
0xbc0005 N     20:00:00:25:b5:77:77:9f             scsi-fcp:init fc-gs
0xbc0006 N     50:0a:09:81:89:3b:63:be  (NetApp)    scsi-fcp:target
               [NetApp2-e2a-FCOE]

Total number of entries = 6
```

Step 3: Verify active zoneset.

Check the fabric configuration for proper zoning by using the **show zoneset active** command, which displays the active zoneset. Each zone that is a member of the active zoneset is displayed with an asterisk (*) to the left of the member. If there is not an asterisk to the left, the host is either down and not logged into the fabric or there is a misconfiguration of the port VSAN or zoning. Use the **show zone** command to display all configured zones on the Cisco Fibre Channel switches.

```
dc5548ax# show zoneset active
zoneset name FCOE_4 vsan 4
  zone name p12-ucs-b-fc0-vhba1_emc-2 vsan 4
  * fcid 0xb90100 [pwwn 50:06:01:61:3c:e0:60:e2] [emc-2-a0-fc]
  * fcid 0xbc0005 [pwwn 20:00:00:25:b5:77:77:9f]

  zone name p12-c210-27-vhba3_netapp-2-e2a vsan 4
  * fcid 0xbc0006 [pwwn 50:0a:09:81:89:3b:63:be] [NetApp2-e2a-
FCOE]
  * fcid 0xbc0000 [pwwn 20:00:58:8d:09:0e:e0:d2] [p12-c210-27-
vhba3]
```

Step 4: Test Fibre Channel reachability by using the **fcping** command, and then trace the routes to the host by using the **fctrace** command. Cisco created these commands to provide storage networking troubleshooting tools that are familiar to individuals who use ping and traceroute.

Process

Configuring Cisco MDS 9148 Switch SAN Expansion

1. Perform initial setup for Cisco MDS
2. Configure VSANs
3. Configure the trunk for SAN interconnect

If your Fibre Channel SAN environment requires a higher density of Fibre Channel port connectivity, you may choose to use Cisco MDS 9100 series SAN switches.

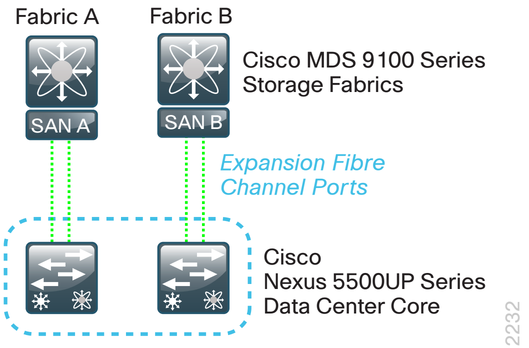
Fabric A
Fabric B
SAN A
SAN B
Cisco MDS 9100 Series Storage Fabrics
Expansion Fibre Channel Ports
Cisco Nexus 5500UP Series Data Center Core
2232

The following procedures describe how to deploy a Cisco MDS 9124 or 9148 SAN switch to connect to the data center core Cisco Nexus 5500UP switches.

Procedure 1 Perform initial setup for Cisco MDS

The following is required to complete this procedure:

- Setting a management IP address
- Configuring console access
- Configuring a secure password

When initially powered on, a new Cisco MDS 9148 switch starts a setup script when accessed from the console.

Step 1: Follow the prompts in the setup script to configure login, out-of-band management, SSH, NTP, switch port modes, and default zone policies.

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: y
  Enter the password for "admin":
  Confirm the password for "admin":
         ---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic
configuration of the system. Setup configures only enough
connectivity for management of the system.
*Note: setup is mainly used for configuring the system
initially, when no configuration is present. So setup
always assumes system defaults and not the current system
configuration values.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
  Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
  Enter the switch name : mds9148ax
  Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: y
    Mgmt0 IPv4 address : 10.4.63.12
    Mgmt0 IPv4 netmask : 255.255.255.0
  Configure the default gateway? (yes/no) [y]: y
    IPv4 address of the default gateway : 10.4.63.1
  Configure advanced IP options? (yes/no) [n]:
  Enable the ssh service? (yes/no) [y]: y
    Type of ssh key you would like to generate (dsa/rsa)
[rsa]: rsa
    Number of rsa key bits <768-2048> [1024]: 2048
  Enable the telnet service? (yes/no) [n]: n
  Enable the http-server? (yes/no) [y]:
 Configure clock? (yes/no) [n]:
 Configure timezone? (yes/no) [n]:
 Configure summertime? (yes/no) [n]:
```

```
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : 10.4.48.17
  Configure default switchport interface state (shut/noshut)
[shut]: noshut
  Configure default switchport trunk mode (on/off/auto) [on]:
  Configure default switchport port mode F (yes/no) [n]: n
  Configure default zone policy (permit/deny) [deny]:
  Enable full zoneset distribution? (yes/no) [n]: y
  Configure default zone mode (basic/enhanced) [basic]:
The following configuration will be applied:
  password strength-check
  switchname mds9148ax
  interface mgmt0
    ip address 10.4.63.12 255.255.255.0
    no shutdown
  ip default-gateway 10.4.63.1
  ssh key rsa 2048 force
  feature ssh
  no feature telnet
  feature http-server
  ntp server 10.4.48.17
  no system default switchport shutdown
  system default switchport trunk mode on
  no system default zone default-zone permit
  system default zone distribute full
  no system default zone mode enhanced

Would you like to edit the configuration? (yes/no) [n]: n
Use this configuration and save it? (yes/no) [y]: y
[########################################] 100%
```

Tech Tip

NTP is critical to troubleshooting and should not be overlooked.

Step 2: Run the setup script for the second Cisco MDS 9100 switch using a unique switch name and 10.4.63.13 for the Mgmt0 IPv4 address.

Step 3: If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the AAA server.

As networks scale in the number of devices to maintain, the operational burden to maintain local user accounts on every device also scales. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root-cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined in the setup script on each Nexus 5500 switch to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

```
feature tacacs+
tacacs-server host 10.4.48.15 key SecretKey
aaa group server tacacs+ tacacs
  server 10.4.48.15
aaa authentication login default group tacacs
```

Reader Tip

The AAA server used in this architecture is the Cisco Access Control System (ACS). For details about Cisco ACS configuration, see the *Cisco SBA—Borderless Networks Device Management Using ACS Deployment Guide.*

Step 4: Set the SNMP strings in order to enable managing Cisco MDS switches with Device Manager. Set both the read-only (**network-operator**) and read/write (**network-admin**) SNMP strings:

```
snmp-server community cisco group network-operator
snmp-server community cisco123 group network-admin
```

Step 5: Configure the clock. In the setup mode, you configured the NTP server address. In this step, configuring the clock enables the clock to use the NTP time for reference and makes the switch output match the local time zone.

```
clock timezone PST -8 0
clock summer-time PDT 2 Sunday march 02:00 1 Sunday nov 02:00 60
```

Procedure 2 Configure VSANs

To configure the Cisco MDS switches to expand the Fibre Channel SAN that you built on the Cisco Nexus 5500UP switches, use the same VSAN numbers for SAN A and SAN B, respectively. The CLI and GUI tools work the same way for Cisco MDS as they do with Cisco Nexus 5500UP.

Step 1: In Device Manager, log in to the first Cisco MDS SAN switch, and then click **FC > VSANS**.

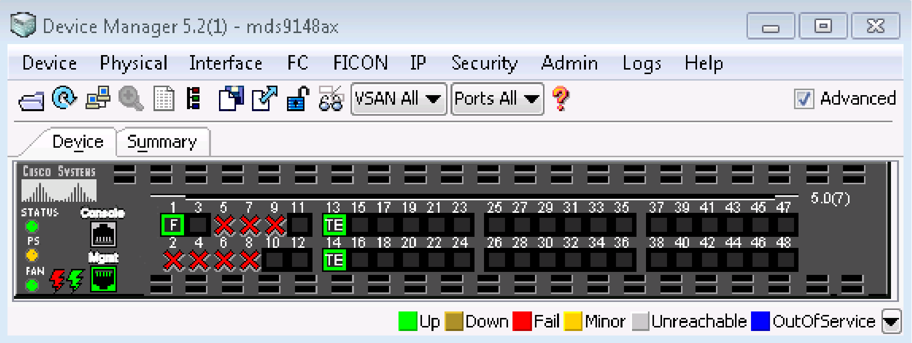

The Create VSAN General window appears.

Step 2: In the VSAN id list, choose **4**, and in the **Name** box, enter **General-Storage**.

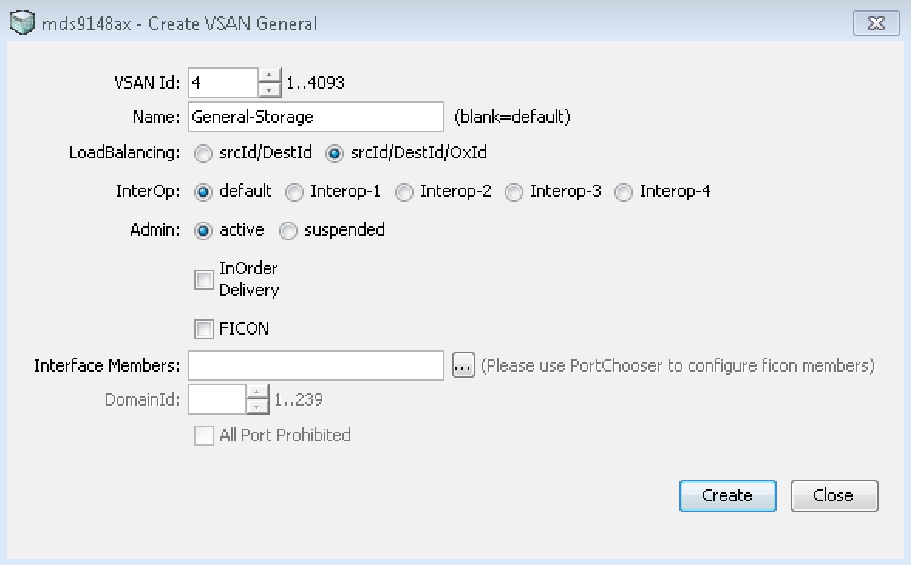

VSAN Id: 4 1..4093

Name: General-Storage (blank=default)

LoadBalancing: ○ srcId/DestId ● srcId/DestId/OxId

InterOp: ● default ○ Interop-1 ○ Interop-2 ○ Interop-3 ○ Interop-4

Admin: ● active ○ suspended

□ InOrder Delivery

□ FICON

Interface Members: (Please use PortChooser to configure ficon members)

DomainId: 1..239

□ All Port Prohibited

Step 3: Click **Create**.

The preceding steps apply this configuration in CLI.

```
vsan database
vsan 4 name "General-Storage"
```

Step 4: Configure the second Cisco MDS SAN switch for VSAN **5** and VSAN name **General-Storage** using Step 1 through Step 3 in this procedure.

Procedure 3 Configure the trunk for SAN interconnect

Connect the Cisco MDS switch to the existing Cisco Nexus 5500UP core Fibre Channel SAN.

Step 1: In Device Manager, navigate to the Cisco MDS switch.

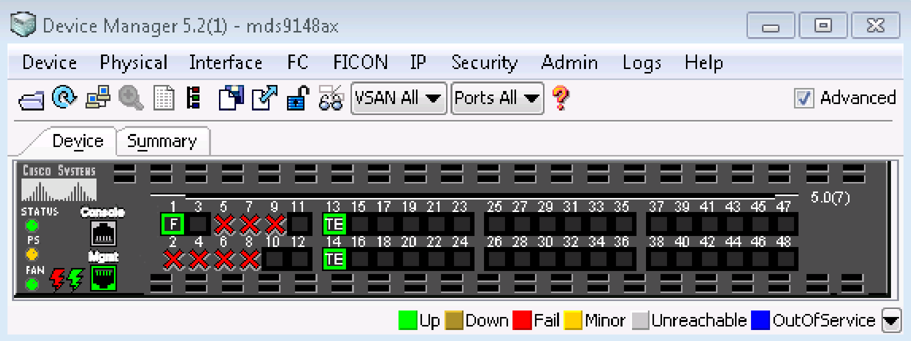

Step 2: In the Device Manager screen, click **Interfaces > Port Channels**, and then click **Create**. Next, you configure the trunk ports on Cisco MDS.

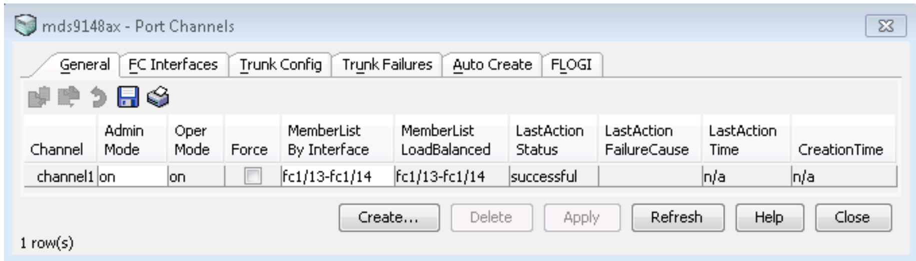

Channel	Admin Mode	Oper Mode	Force	MemberList By Interface	MemberList LoadBalanced	LastAction Status	LastAction FailureCause	LastAction Time	CreationTime
channel1	on	on	□	fc1/13-fc1/14	fc1/13-fc1/14	successful		n/a	n/a

Step 3: Choose the port channel Id number, select **trunk**, select Mode **E**, and then select **Force**.

Step 4: In the Allowed VSANs box, enter **1,4**. For the Cisco MDS switch for SAN Fabric B, the VSANs to enter would be **1** and **5**.

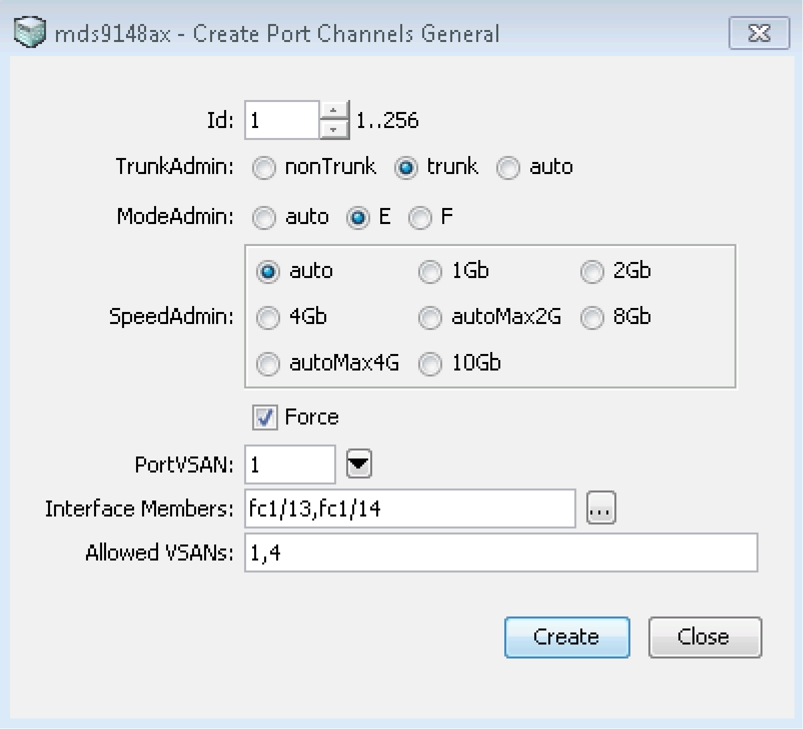

mds9148ax - Create Port Channels General

Id: 1 1..256

TrunkAdmin: ○ nonTrunk ● trunk ○ auto

ModeAdmin: ○ auto ● E ○ F

SpeedAdmin: ● auto ○ 1Gb ○ 2Gb ○ 4Gb ○ autoMax2G ○ 8Gb ○ autoMax4G ○ 10Gb

☑ Force

PortVSAN: 1

Interface Members: fc1/13,fc1/14

Allowed VSANs: 1,4

Step 5: To the right of the Interface Members box, click the ellipsis button (...), and then select the interface members that will belong to this port channel.

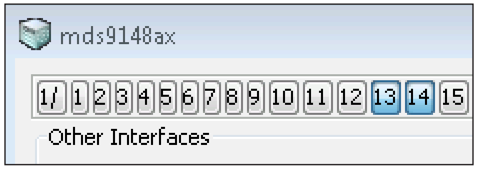

Step 6: Click **Create.** The new port channel is created.

Step 7: Right click the Fibre Channel ports used for the port channel, and then select **enable**.

The preceding steps apply this Cisco MDS 9100 configuration to the MDS SAN-A switch.

```
interface port-channel 1
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 4
  switchport rate-mode dedicated

interface fc1/13
  switchport mode E
  channel-group 1 force
  switchport rate-mode dedicated
  no shutdown
interface fc1/14
  switchport mode E
  channel-group 1 force
  switchport rate-mode dedicated
  no shutdown
```

The preceding steps apply this Cisco MDS 9100 configuration to the MDS SAN-B switch.

```
interface port-channel 1
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 5
  switchport rate-mode dedicated

interface fc1/13
  switchport mode E
  channel-group 1 force
  switchport rate-mode dedicated
  no shutdown
interface fc1/14
  switchport mode E
  channel-group 1 force
  switchport rate-mode dedicated
  no shutdown
```

Step 8: Create the corresponding SAN port channel to connect to the Cisco MDS switch for Cisco Nexus 5500UP by following the preceding steps in this procedure (Procedure 3).

The resulting Cisco Nexus 5500UP CLI for this SAN port channel is the following for the SAN-A switch.

```
interface san-port-channel 31
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 4

interface fc1/31
  switchport description Link to dcmds9148ax port fc-1/13
  switchport mode E
  channel-group 31 force
  no shutdown

interface fc1/32
  switchport description Link to dcmds9148ax port fc1/14
  switchport mode E
  channel-group 31 force
  no shutdown
```

The resulting Cisco Nexus 5500UP CLI for this SAN port channel is the following for the SAN-B switch.

```
interface san-port-channel 31
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 5

interface fc1/31
  switchport description Link to dcmds9148bx port fc-1/13
  switchport mode E
  channel-group 31 force
  no shutdown

interface fc1/32
  switchport description Link to dcmds9148bx port fc1/14
  switchport mode E
  channel-group 31 force
  no shutdown
```

Step 9: Distribute the zone database created on the Cisco Nexus 5500UP switch to the new Cisco MDS 9100 switch.

Configure the Cisco Nexus 5500UP CLI for SAN-A to distribute the zone database to the new Cisco MDS 9100 switch.

```
zoneset distribute full vsan 4
```

Configure the Cisco Nexus 5500UP CLI for SAN-B to distribute the zone database to the new Cisco MDS 9100 switch.

```
zoneset distribute full vsan 5
```

Process

Configuring FCoE Host Connectivity

1. Configure FCoE QoS
2. Configure host-facing FCoE ports
3. Verify FCoE connectivity

Cisco UCS C-Series rack-mount servers ship with onboard 10/100/1000 Ethernet adapters and Cisco Integrated Management Controller (CIMC), which uses a 10/100 Ethernet port. To get the most out of the rack servers and minimize cabling in the Cisco SBA Unified Computing architecture, the Cisco UCS C-Series rack-mount server is connected to a unified fabric. The Cisco Nexus 5500UP Series switch that connects the Cisco UCS 5100 Series Blade Server Chassis to the network can also be used to extend Fibre Channel traffic over 10-Gigabit Ethernet. The Cisco Nexus 5500UP Series switch consolidates I/O onto one set of 10-Gigabit Ethernet cables, eliminating redundant adapters, cables, and ports. A single converged network adapter (CNA) card and set of cables connects servers to the Ethernet and Fibre Channel networks by using FCoE. FCoE and CNA also allows the use of a single cabling infrastructure within server racks.

In the Cisco SBA data center, the Cisco UCS C-Series rack-mount server is configured with a dual-port CNA. Cabling the Cisco UCS C-Series server with a CNA limits the cables to three: one for each port on the CNA and one for the CIMC connection.

Tech Tip

A server connecting to Cisco Nexus 5500UP that is running FCoE consumes a Fibre Channel port license. If you are connecting the FCoE attached servers to a Cisco FEX model 2232PP, only the 5500UP ports connected to the Cisco FEX require a Fibre Channel port license for each port connecting to the Cisco FEX. This way, you could connect up to 32 FCoE servers to a Cisco FEX 2232PP and only use Fibre Channel port licenses for the Cisco FEX uplinks.

A standard server without a CNA could have a few Ethernet connections or multiple Ethernet and Fibre Channel connections. The following figure shows a topology with mixed unified fabric, standard Ethernet and Fibre Channel connections, and optional Cisco MDS 9100 Series for Fibre Channel expansion.

Notes

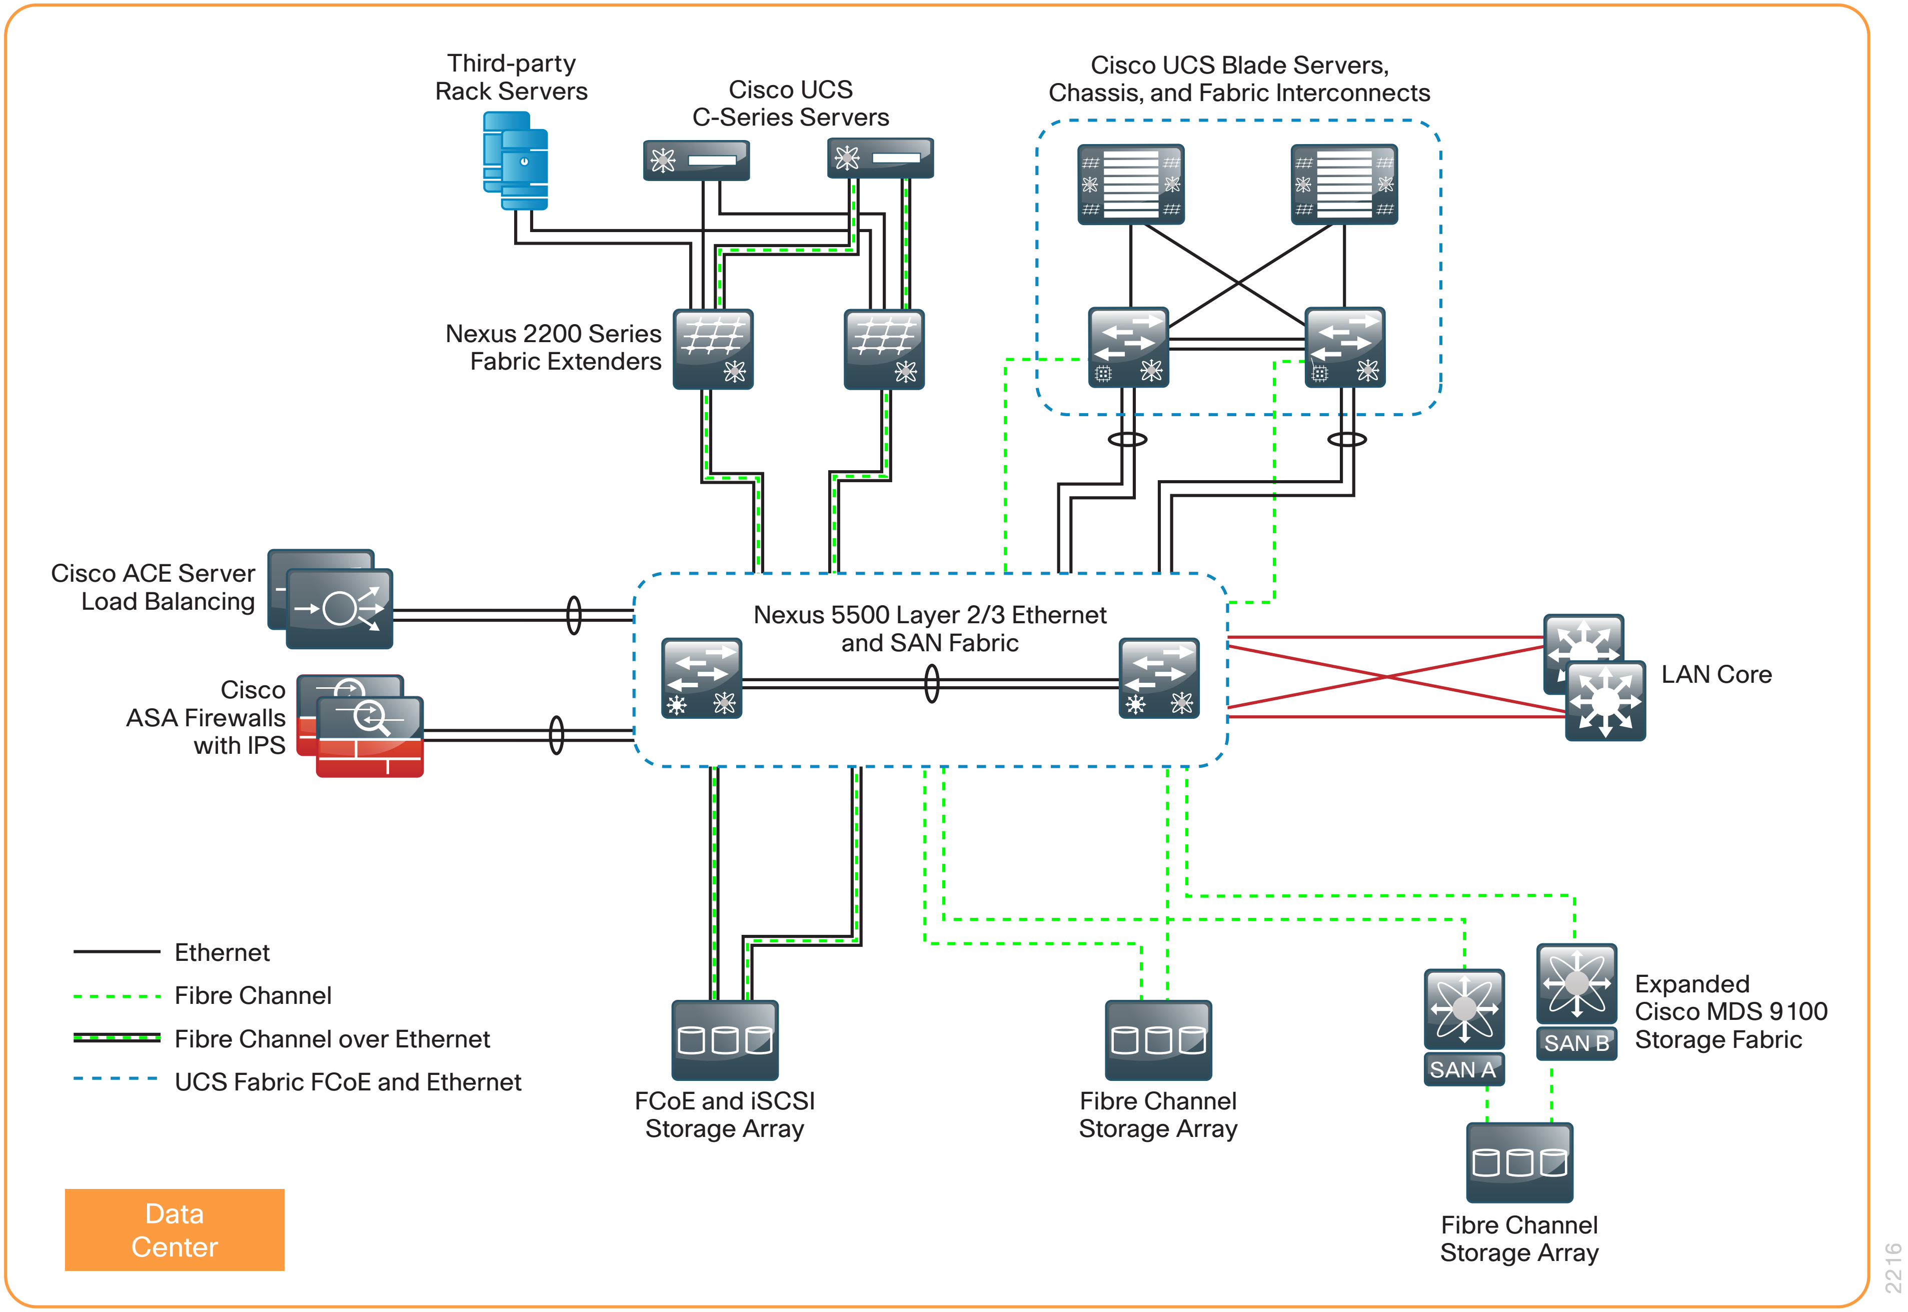
Third-party Rack Servers
Cisco UCS C-Series Servers
Cisco UCS Blade Servers, Chassis, and Fabric Interconnects
Nexus 2200 Series Fabric Extenders
Cisco ACE Server Load Balancing
Cisco ASA Firewalls with IPS
Nexus 5500 Layer 2/3 Ethernet and SAN Fabric
LAN Core
Ethernet
Fibre Channel
Fibre Channel over Ethernet
UCS Fabric FCoE and Ethernet
FCoE and iSCSI Storage Array
Fibre Channel Storage Array
SAN A
SAN B
Expanded Cisco MDS 9100 Storage Fabric
Fibre Channel Storage Array
Data Center
2216

The Cisco UCS C-Series server is connected to both Cisco Nexus 5500UP Series switches from the CNA with twinax or fiber optic cabling. The Cisco UCS server running FCoE can also attach to a single-homed Cisco FEX model 2232PP.

Tech Tip

At this time, FCoE-connected hosts can only connect over 10-Gigabit Ethernet and must use a fiber optic or twinax connection.

The recommended approach is to connect the CIMC 10/100 management port(s) to an Ethernet port on the out-of-band management switch. Alternatively, you can connect the CIMC management port(s) to a Cisco Nexus 2248 fabric extender port in the management VLAN (163).

Cisco Nexus 5500UP Configuration for FCoE

In previous processes, you enabled Cisco Nexus 5500UP Series FCoE functionality. In this process, you perform the following tasks to allow a Cisco C-Series server to connect using FCoE:

- Create a virtual Fibre Channel interface
- Assign the VSAN to a virtual Fibre Channel interface
- Configure the Ethernet port and trunking

Procedure 1 Configure FCoE QoS

Configuration is the same across both of the Cisco Nexus 5500UP Series switches with the exception of the VSAN configured for SAN fabric A and for SAN fabric B.

The Cisco Nexus 5500UP, unlike Cisco Nexus 5010, does not preconfigure QoS for FCoE traffic.

Step 1: Ensure that the Cisco Nexus 5500UP data center core switches have been programmed with a QoS policy to support lossless FCoE transport. The QoS policy for the data center core Nexus 5500UP switches was defined in Procedure 3 “Configure QoS policies.”

Tech Tip

You must have a QoS policy on the Cisco Nexus 5500UP switches that classifies FCoE for lossless operation.

Procedure 2 Configure host-facing FCoE ports

On the Cisco Nexus 5500UP switch, configure the Ethernet ports connected to the CNA on the dual-homed host.

Step 1: Create a VLAN that will carry FCoE traffic to the host.

- In the following, VLAN 304 is mapped to VSAN 4. VLAN 304 carries all VSAN 4 traffic to the CNA over the trunk for the first Cisco Nexus 5500UP switch.

```
vlan 304
fcoe vsan 4
exit
```

- On the second Cisco Nexus 5500UP switch, VLAN 305 is mapped to VSAN 5.

```
vlan 305
fcoe vsan 5
exit
```

Step 2: Create a virtual Fibre Channel (vfc) interface for Fibre Channel traffic, and then bind it to the corresponding host Ethernet interface. You need to do this in order to be able to map an FCoE interface to Fibre Channel.

This example shows binding to a Cisco FEX 2232PP Ethernet interface. This command will be the same on both Cisco Nexus 5500UP switches.

```
interface  vfc1
bind interface Ethernet 103/1/3
no shutdown
exit
```

Step 3: Add the vfc interface to the VSAN database.

- On the first Cisco Nexus 5500UP switch, the vfc is mapped to VSAN 4.

```
vsan database
vsan 4 interface vfc 1
exit
```

- On the second Cisco Nexus 5500UP switch, the vfc is mapped to VSAN 5.

```
vsan database
vsan 5 interface vfc 1
exit
```

Step 4: Configure the Ethernet interface to operate in trunk mode, configure the interface with the FCoE VSAN and any data VLANs required by the host, and configure the spanning-tree port type as **trunk edge**.

- This example shows the configuration of the first Cisco Nexus 5500UP switch.

```
interface Ethernet 103/1/3
switchport mode trunk
switchport trunk allowed vlan 148-162,304
spanning-tree port type edge trunk
no shut
```

- This example shows the configuration of the second Cisco Nexus 5500UP switch.

```
interface Ethernet 103/1/3
switchport mode trunk
switchport trunk allowed vlan 148-162,305
spanning-tree port type edge trunk
no shut
```

Step 5: Configure VSAN on a Cisco UCS C-Series server.

Tech Tip

The Cisco UCS C-Series server using the Cisco P81E CNA must have the FCoE VSANs configured for virtual host bus adapter (vHBA) operation to connect to the Fibre Channel fabric. For more information on configuring the C-Series server for FCoE connectivity, please see the *Cisco SBA—Data Center Unified Computing System Deployment Guide*.

Procedure 3 Verify FCoE connectivity

Step 1: On the Cisco Nexus 5500UP switches, use the **show interface** command to verify the status of the virtual Fibre Channel interface. The interface should now be up as seen below if the host is properly configured to support the CNA.

Reader Tip

Host configuration is beyond the scope of this guide. Please see CNA documentation for specific host drivers and configurations.

```
dc5548ax# show interface vfc1
vfc1 is trunking (Not all VSANs UP on the trunk)
    Bound interface is Ethernet103/1/3
    Hardware is Virtual Fibre Channel
    Port WWN is 20:00:54:7f:ee:17:cf:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 4
    Trunk vsans (admin allowed and active) (1,4)
    Trunk vsans (up)                       (4)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1)
    1 minute input rate 1672 bits/sec, 209 bytes/sec, 0
frames/sec
    1 minute output rate 320 bits/sec, 40 bytes/sec, 0 frames/
sec
      117038 frames input, 39607100 bytes
        0 discards, 0 errors
      128950 frames output, 33264140 bytes
        0 discards, 0 errors
    last clearing of "show interface" counters never
    Interface last changed at Tue Nov  8 11:11:29 2011
```

Step 2: On the Cisco Nexus 5500UP switches, display the FCoE addresses.

```
    dc5548ax# show fcoe database
-----------------------------------------------------------------
INTERFACE  FCID        PORT NAME                MAC ADDRESS
-----------------------------------------------------------------
vfc1       0xbc0000    20:00:58:8d:09:0e:e0:d2  58:8d:09:0e:e0:d2
```

Step 3: Show the FLOGI database for FCoE login. The vfc1 addresses appear in the current FLOGI database on the Cisco Nexus 5500 switch.

```
dc5548ax# show flogi database
-----------------------------------------------------------------
INTERFACE VSAN FCID     PORT NAME               NODE NAME
-----------------------------------------------------------------
vfc1    4  0xbc0000  20:00:58:8d:09:0e:e0:d2  10:00:58:8d:09:0e:e0:d2
                     [p12-c210-27-vhba3]
```

Step 4: Show the FCNS database for FCoE login. The FCNS database shows the FCoE host logged in and the FC-4 TYPE:FEATURE information.

```
dc5548ax# show fcns database
VSAN 4:
-----------------------------------------------------------------
FCID         TYPE  PWWN              (VENDOR)     FC4-TYPE:FEATURE
-----------------------------------------------------------------
0xbc0000     N     20:00:58:8d:09:0e:e0:d2        scsi-fcp:init fc-gs
                   [p12-c210-27-vhba3]
```

Now you can configure zoning and device aliases per the procedures in the "Configuring Fibre Channel SAN on Cisco Nexus 5500UP" process, earlier in this chapter.

Tech Tip

Much of the configuration of the Cisco Nexus 5500UP Series switch can also be done from within Device Manager; however, Device Manager for SAN Essentials cannot be used to configure VLANs or Ethernet trunks on the Cisco Nexus 5500UP Series switches.

Compute Connectivity

Business Overview

As an organization grows, the number and type of servers required to handle the information processing tasks of the organization grows as well. This imposes several challenges:

- Increased data center square footage and rack space
- More power and cooling, particularly in light of the fact that every new CPU generation increases wattage dissipation as core speeds increase
- Increased complexity of the data-networking cable plant to provide adequate capacity and capability for increasing server counts
- More hardware capital expense to buy server platforms and spares, and greater operational expense to administer and maintain diverse hardware and OS platforms
- Migration from existing servers and applications to newer platforms and connection methods, which requires a flexible architecture that accommodates both legacy and new servers and applications
- Increased resiliency and migration-path challenges, because appliance-centric or server-centric application platforms tend to be platform-centric and may not lend themselves well to being load-balanced or moved to disparate platforms

Organizations frequently need to optimize their investment in server resources, so that the organization can add new applications while controlling costs as they move from a small server room environment into a data center.

Scaling a data center with conventional servers, networking equipment, and storage resources can pose a significant challenge to a growing organization. Multiple hardware platforms and technologies must be integrated to deliver the expected levels of performance and availability to application end users. These components in the data center also need to be managed and maintained, typically with a diverse set of management tools with different interfaces and approaches.

Technology Overview

Server virtualization offers the capability to run multiple application servers on a common hardware platform, allowing an organization to focus on maximizing the application capability of the data center while minimizing costs. Increased capability and reduced costs are realized through multiple aspects:

- Multiple applications can be combined in a single hardware chassis, reducing the number of boxes that must be accommodated in data-center space
- Simplified cable management, due to fewer required cable runs and greater flexibility in allocating network connectivity to hosts on an as-needed basis
- Improved resiliency and application portability as hypervisors allow workload resiliency and load-sharing across multiple platforms, even in geographically dispersed locations
- Applications that are deployed on standardized hardware platforms, which reduces platform-management consoles and minimizes hardware spare stock challenges
- Minimized box count reduces power and cooling requirements, because there are fewer lightly loaded boxes idling away expensive wattage

The ability to virtualize server platforms to handle multiple operating systems and applications with hypervisor technologies building virtual machines (VMs) allows the organization to lower capital and operating costs by collapsing more applications onto fewer physical servers. The hypervisor technology also provides the ability to cluster many virtual machines into a domain where workloads can be orchestrated to move around the data center to provide resiliency and load balancing, and to allow new applications to be deployed in hours versus days or weeks.

The ability to move VMs or application loads from one server to the next, whether the server is a blade server in a chassis-based system or a standalone rack-mount server, requires the network to be flexible and scalable, allowing any VLAN to appear anywhere in the data center. Cisco Virtual Port Channel (vPC) and Fabric Extender (FEX) technologies are used extensively in the Cisco SBA data center to provide flexible Ethernet connectivity to VLANs distributed across the data center in a scalable and resilient manner.

Streamlining the management of server hardware and its interaction with networking and storage equipment is another important component of using this investment in an efficient manner. Cisco offers a simplified reference model for managing a small server room as it grows into a full-fledged data center. This model benefits from the ease of use offered by Cisco UCS. Cisco UCS provides a single graphical management tool for the provisioning and management of servers, network interfaces, storage interfaces, and the network components directly attached to them. Cisco UCS treats all of these components as a cohesive system, which simplifies these complex interactions and allows an organization to deploy the same efficient technologies as larger enterprises do, without a dramatic learning curve.

The primary computing platforms targeted for the Cisco SBA Unified Computing reference architecture are Cisco UCS B-Series Blade Servers and Cisco UCS C-Series Rack-Mount Servers. The Cisco UCS Manager graphical interface provides ease of use that is consistent with the goals of Cisco SBA. When deployed in conjunction with the Cisco SBA data center network foundation, the environment provides the flexibility to support the concurrent use of the Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack-Mount Servers, and third-party servers connected to 1- and 10-Gigabit Ethernet connections.

The following sections describe features that enhance connectivity options in the data center.

Cisco Nexus Virtual Port Channel

As described in the “Ethernet Infrastructure” chapter, Virtual Port Channel (vPC) allows links that are physically connected to two different Cisco Nexus switches to appear to a third downstream device to be coming from a single device and as part of a single Ethernet port channel. The third device can be a server, switch, or any other device or appliance that supports IEEE 802.3ad port channels. For Cisco EtherChannel technology, the term “multichassis EtherChannel” (MCEC) refers to this technology. MCEC links from a device connected to the data center core that provides spanning-tree loop–free topologies, allowing VLANs to be extended across the data center while maintaining a resilient architecture.

A *vPC domain* consists of two vPC peer switches connected by a peer link. Of the vPC peers, one is primary and one is secondary. The system formed by the switches is referred to as a vPC domain. The vPC peer link between the two Cisco Nexus switches is the most important connectivity element in the system. This link is used to create the illusion of a single control plane between the two switches, and carries critical control plane packets as well as data packets when devices are single-homed due to design or EtherChannel link failure. For a VLAN to be forwarded on a vPC, that VLAN must exist on the peer link and both vPC peer switches.

The *vPC peer-keepalive link* is used to resolve dual-active scenarios in which the peer link connectivity is lost. If the vPC peer link connectivity is lost, the secondary vPC peer will shut down all vPC member links and the primary vPC switch will continue forwarding packets, providing a resilient architecture.

A *vPC port* is a port that is assigned to a vPC channel group. The ports that form the vPC are split between the vPC peers, must be defined identically on both vPC switches, and are referred to as vPC member ports. A non-vPC port, also known as an orphaned port, is a port that belongs to a VLAN that is part of a vPC, but is not programmed as a vPC member. The following figure illustrates vPC ports and orphan ports. The Host with Active-Standby teaming interfaces would be considered as vPC orphan ports.

Figure 14 - vPC member and non-member ports

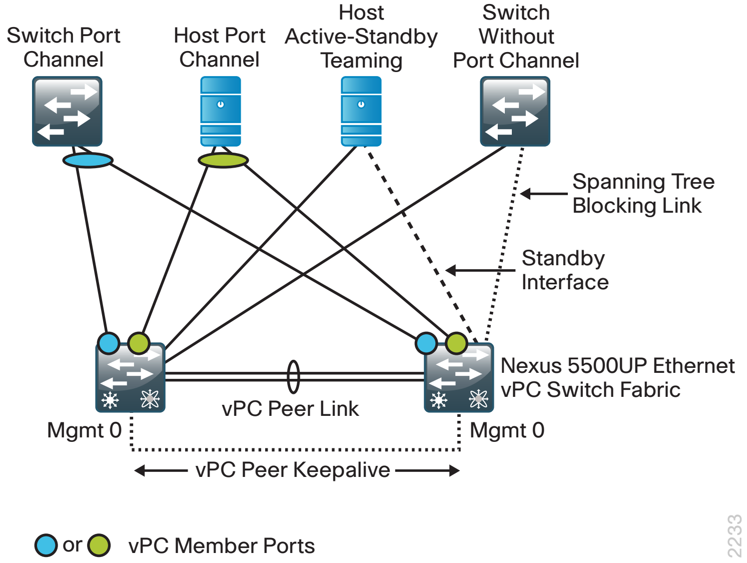

The important point to remember about vPC orphan ports is that if the vPC peer link is lost and the secondary vPC shuts down vPC ports, it will not shut down vPC orphan ports unless programmed to do so with the **vpc orphan-port suspend** command on the switch interface.

Example

```
interface Ethernet103/1/2
  description to_teamed_adapter
  switchport mode access
  switchport access vlan 50
  vpc orphan-port suspend

interface Ethernet104/1/2
  description to_teamed_adapter
  switchport mode access
  switchport access vlan 50
  vpc orphan-port suspend
```

Reader Tip

The fundamental concepts of vPC are described in detail in the whitepaper titled “Cisco NX-OS Virtual PortChannel: Fundamental Design Concepts with NXOS 5.0” located on www.cisco.com

The complete vPC domain programming for the Cisco Nexus 5500UP switches is detailed in the Procedure 4, ”Configure virtual port channel,” earlier in this guide.

Cisco Nexus Fabric Extender

As described earlier in the “Ethernet Infrastructure” chapter , the Cisco Fabric Extender (FEX) acts as a remote line card to the attached Cisco Nexus 5500UP switch that it is connected to. This allows for central configuration of all switch ports on the data center core switches, and provides fan out to higher-density Fast Ethernet, 1-Gigabit Ethernet, and 10-Gigabit Ethernet for top-of-rack server connectivity. Because the Cisco FEX acts as a line card on the Cisco Nexus 5500UP switch, extending VLANs to server ports on different Cisco FEXs does not create spanning-tree loops across the data center.

The Cisco FEX can be single-homed to a data center core switch (also called *straight-through mode*) or dual-homed using vPC (also called *active/active mode*).

Figure 15 - Cisco Nexus FEX connectivity to data center core

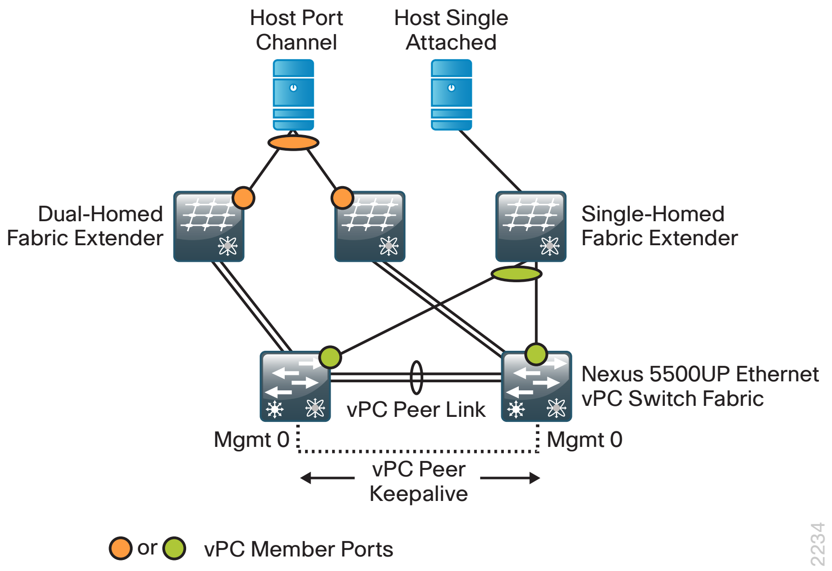

The dual-homed (active/active) Cisco FEX uses vPC to provide resilient connectivity to both data center core switches for single attached host servers. Each host is considered to be vPC connected through the associated connectivity to a vPC dual-homed Cisco FEX. The Cisco FEX–to-core connectivity ranges from 4 to 8 uplinks, depending on the Cisco FEX type in use, and the Cisco FEX uplinks can be configured as a port channel as well.

The host connected to a pair of single-homed Cisco FEXs can be configured for port channel operation to provide resilient connectivity to both data center core switches through the connection to each Cisco FEX. The Cisco FEX–to-core connectivity ranges from 4 to 8 uplinks, depending on the Cisco FEX type in use, and the Cisco FEX uplinks are typically configured as a port channel as well.

Tech Tip

Devices such as LAN switches that generate spanning-tree bridge protocol data units (BPDUs) should not be connected to Cisco FEXs. The Cisco FEX is designed for host connectivity and will error disable a port that receives a BPDU packet.

The complete Cisco FEX connectivity programming to the Cisco Nexus 5500UP data center core switches and Ethernet port configuration for server connection is detailed in the "Ethernet Infrastructure" chapter, earlier in this guide.

Cisco UCS System Network Connectivity

Both Cisco UCS B-Series Blade Servers and C-Series Rack Mount Servers integrate cleanly into the Cisco SBA data center design. The Cisco Nexus 5500UP data center core provides 1-Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel SAN connectivity in a single platform.

Cisco UCS B-Series Blade Chassis System Components

The Cisco UCS Blade Chassis system has a unique architecture that integrates compute, data network access, and storage network access into a common set of components under a single-pane-of-glass management interface. The primary components included within this architecture are as follows:

- Cisco UCS 6200 Series Fabric Interconnects—Provide both network connectivity and management capabilities to the other components in the system.
- Cisco UCS 2200 Series Fabric Extenders—Logically extend the fabric from the fabric interconnects into each of the enclosures for Ethernet, FCoE, and management purposes.
- Cisco UCS 5100 Series Blade Server Chassis—Provides an enclosure to house up to eight half-width or four full-width blade servers, their associated fabric extenders, and four power supplies for system resiliency.
- Cisco UCS B-Series Blade Servers—Available in half-width or full-width form factors, with a variety of high-performance processors and memory architectures to allow customers to easily customize their compute resources to the specific needs of their most critical applications.
- Cisco UCS B-Series Network Adapters—A variety of mezzanine adapter cards that allow the switching fabric to provide multiple interfaces to a server.

The following figure shows an example of the physical connections required within a Cisco UCS Blade Chassis system to establish the connection between the fabric interconnects and a single blade chassis. The links between the blade chassis and the fabric interconnects carry all server data traffic, centralized storage traffic, and management traffic generated by Cisco UCS Manager.

Figure 16 - Cisco UCS Blade Chassis System component connections

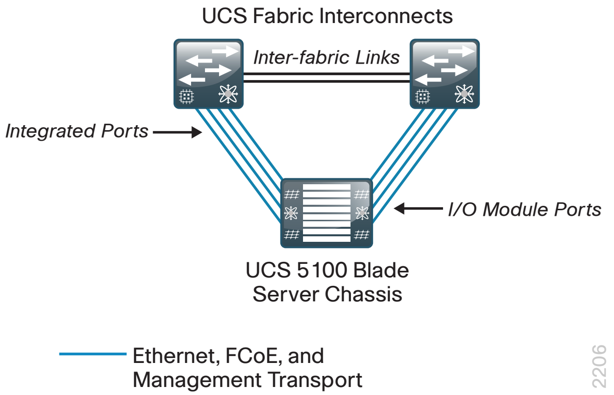

Cisco UCS Manager

Cisco UCS Manager is embedded software resident on the fabric interconnects, providing complete configuration and management capabilities for all of the components in Cisco UCS. This configuration information is replicated between the two fabric interconnects, providing a highly available solution for this critical function. The most common way to access Cisco UCS Manager for simple tasks is to use a web browser to open the Java-based GUI. For command-line or programmatic operations against the system, a CLI and an XML API are also included with the system.

Cisco UCS B-Series System Network Connectivity

Cisco UCS 6200 Series Fabric Interconnects provide connectivity for Cisco UCS Blade Server systems. The following figure shows a detailed example of the connections between the fabric interconnects and the Cisco Nexus 5500UP Series data center core.

The default and recommended configuration for the fabric interconnects is end-host mode, which means they do not operate as full LAN switches but rather rely on the upstream data center switching fabric. In this way, Cisco UCS appears to the network as a virtualized compute cluster with multiple physical connections. Individual server traffic is pinned to specific interfaces, with failover capability in the event of loss of the primary link. The Ethernet traffic from the fabric interconnects shown in Figure 17 uses vPC links to the data center core for resiliency and traffic load sharing. The Fibre Channel links to the core use SAN port channels for load sharing and resiliency as well.

Figure 17 - Cisco UCS fabric interconnect to core

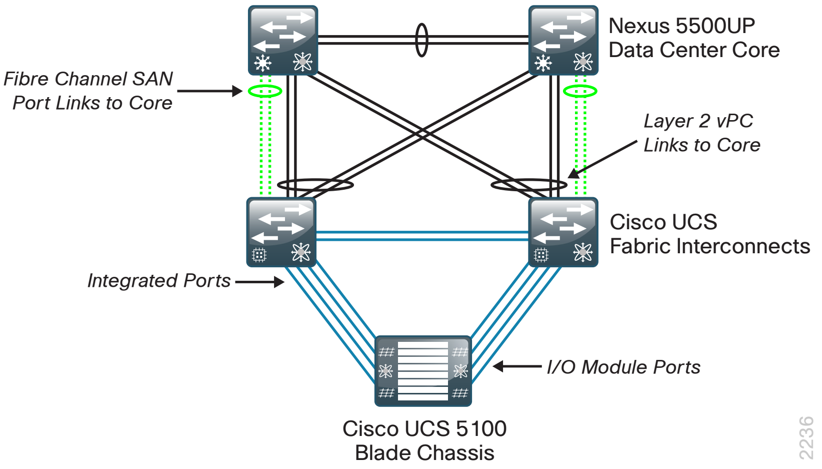

Detailed configuration for Cisco UCS B-Series deployment can be found in the *Cisco SBA—Data Center Unified Computing System Deployment Guide*.

Cisco UCS C-Series Network Connectivity

Cisco UCS C-Series Rack-Mount Servers balance simplicity, performance, and density for production-level virtualization, web infrastructure, and data center workloads. Cisco UCS C-Series servers extend Unified Computing innovations and benefits to rack-mount servers.

The Cisco Nexus switching fabric provides connectivity for 10-Gigabit or 1-Gigabit Ethernet attachment for Cisco UCS C-Series servers, depending on the throughput requirements of the applications or virtual machines in use and the number of network interface cards installed per server. Figure 18 shows some examples of dual-homed connections from Cisco UCS C-Series servers to single-homed Cisco FEXs, providing 1-Gigabit and 10-Gigabit Ethernet connections. Ten-Gigabit Ethernet connections capable of supporting Ethernet and FCoE are available either through the Cisco Nexus 2232PP Fabric Extender or by using 10-Gigabit ports directly on the Cisco Nexus 5500UP Series switch pair. Connections for Fast Ethernet or 1-Gigabit Ethernet can also use the Cisco Nexus 2248TP Fabric Extender.

Figure 18 - Example Cisco UCS C-Series FEX Connections

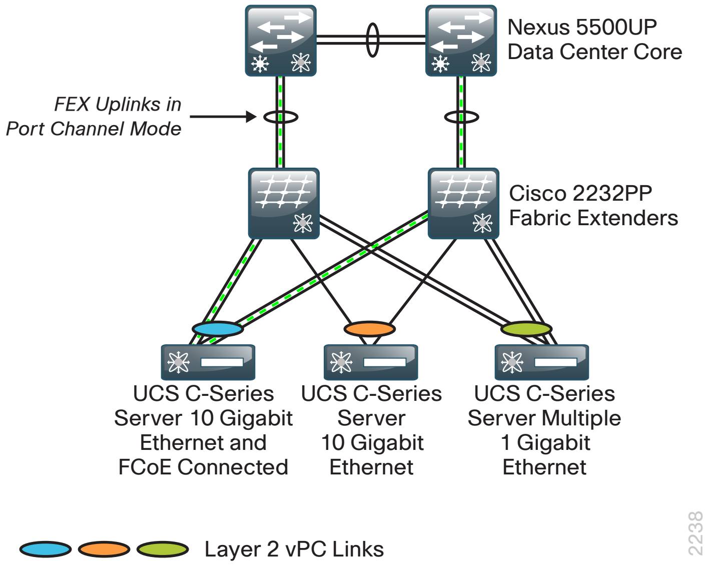

The Cisco UCS C-Series Server connectivity to Cisco FEX options in Figure 18 above all make use of vPC connections by using IEEE 802.3ad EtherChannel from the host to single-homed Cisco Nexus 2232PP FEXs. When using vPC for server connections, each server interface must be identically configured on each data center core Cisco Nexus 5500UP switch. The Cisco FEX–to–data center core uplinks use a port channel to load balance server connections over multiple links and provide added resiliency.

The Cisco UCS C-Series Server with 10-Gigabit Ethernet and FCoE connectivity uses a converged network adapter (CNA) in the server and must connect to either a Cisco Nexus 2232PP FEX or directly to the Cisco Nexus 5500UP switch. This is because FCoE uplinks must use a fiber optic or twinax connection to maintain bit error rate (BER) thresholds for Fibre Channel transport. Cisco supports FCoE on 10-Gigabit Ethernet only at this time. If used with vPC, the Ethernet traffic is load balanced across the server links with EtherChannel and Fibre Channel runs up each link to the core, with SAN-A traffic on one link to the connected Cisco FEX and data center core switch, and SAN-B traffic on the other link to the connected Cisco FEX and data center core switch, as is typical of Fibre Channel SAN traffic.

Example

- This example shows the configuration of the FEX interface on the first Cisco Nexus 5500UP switch.

```
interface Ethernet 103/1/3
  description Dual-homed server FCoE link to SAN-A VSAN 304
  switchport mode trunk
  switchport trunk allowed vlan 148-163,304
  spanning-tree port type edge trunk
  no shut
```

- This example shows the configuration of the FEX interface on the second Cisco Nexus 5500UP switch.

```
interface Ethernet 104/1/3
  description Dual-homed server FCoE link to SAN-B VSAN 305
  switchport mode trunk
  switchport trunk allowed vlan 148-163,305
  spanning-tree port type edge trunk
  no shut
```

The Cisco UCS C-Series Server with 10-Gigabit Ethernet without FCoE can connect to a Cisco Nexus 2232 FEX or directly to the Cisco Nexus 5500UP switch. These server connections can be fiber optic, copper, or twinax, depending on the Cisco FEX and server combination used. If used with vPC, the Ethernet traffic is load balanced across the server links with EtherChannel.

The Cisco UCS C-Series Server with multiple 1-Gigabit Ethernet uses vPC to load balance traffic over multiple links using EtherChannel. The use of vPC is not a requirement. In a non-vPC server connection where you want independent server interfaces, you may prefer connecting to a dual-homed Cisco FEX for resiliency unless the server operating system provides resilient connectivity.

Configuration for the Cisco Nexus FEX to Cisco Nexus 5500UP switch connections is detailed in the "Ethernet Infrastructure" chapter earlier in this guide. Detailed configuration for Cisco UCS C-Series deployment can be found in the *Cisco SBA—Data Center Unified Computing System Deployment Guide*.

Single-Homed Server Connectivity

As an organization grows, it may need to provide connectivity in the data center for many legacy servers and appliances with a single Fast Ethernet or Gigabit Ethernet. To provide added resiliency for these servers, a dual-homed Cisco FEX using vPC for the Cisco FEX connection to the data center is recommended as shown in the figure below.

Figure 19 - Single-homed server to dual-homed Cisco FEX

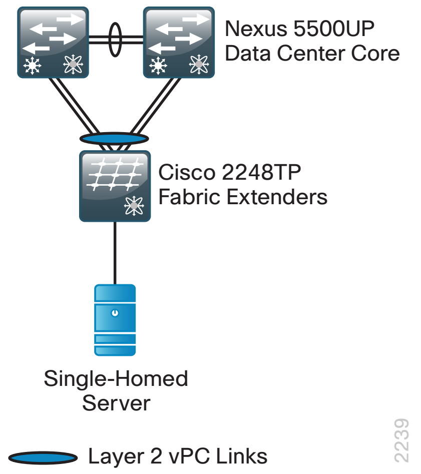

The vPC connection from the Cisco Nexus 2248TP FEX provides both control plane and data plane redundancy for servers connected to the same Cisco FEX. This topology provides resiliency for the attached servers in the event of a fabric uplink or Cisco Nexus 5500UP core switch failure; however there is no resiliency in the event of a Cisco Nexus 2248TP failure. All servers connected to the vPC dual-homed Cisco FEX are vPC connections and must be configured on each data center core Cisco Nexus 5500UP switch. Although this approach does provide added resiliency, single-homed servers hosting important applications should be migrated to dual-homed connectivity to provide sufficient resiliency.

Server with Teamed Interface Connectivity

Server NIC teaming comes in many options and features. NIC adapters and operating systems capable of using IEEE 802.3ad EtherChannel from servers to a Cisco FEX would use the vPC option covered in the "Cisco UCS C-Series Connectivity" section, earlier in this chapter. NIC adapters and operating systems using an active/standby method for connecting to the Cisco FEX are best served by a dual-homed Cisco FEX as shown in the figure below.

Figure 20 - Server with active/standby NIC–to–Cisco FEX connection

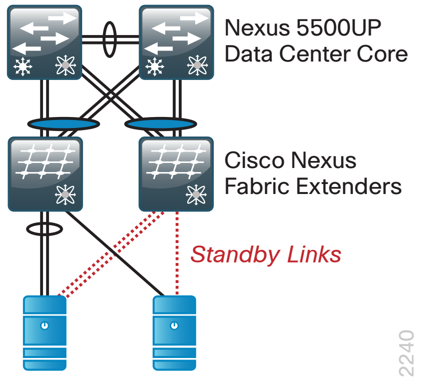

The vPC connection from the Cisco Nexus 2248TP FEX provides both control plane and data plane redundancy for servers connected to each Cisco FEX. This topology provides resiliency for the attached servers in the event of a fabric uplink or Cisco Nexus 5500UP core switch failure. In the event of a Cisco FEX failure, the NIC teaming switches to the standby interfaces.

Enhanced Fabric Extender and Server Connectivity

The dual-homed Cisco Nexus fabric extender enhances system reliability by connecting the FEX to both core switches. With a dual-homed FEX, one of the data center core Cisco Nexus 5500UP switches can be taken out of service and traffic will continue to flow over the FEX uplinks to the remaining active data center core switch. Until recently, the dual-homed FEX connection was unable to support a server that was connected to two dual-homed FEX with a single EtherChannel connected server. This condition meant that you may have needed a mix of single-homed FEX and dual-homed FEX in the data center to support different server connectivity requirements.

As of Cisco NX-OS release 5.1(3)N1(1) for the Cisco Nexus 5500 Series switches, the Cisco Nexus 5500 switch can now support a port channel connected to two dual-homed FEXes as shown in Figure 21. This new capability is referred to as Enhanced vPC. The Cisco 5000 switch does not support Enhanced vPC.

Figure 21 - Enhanced vPC

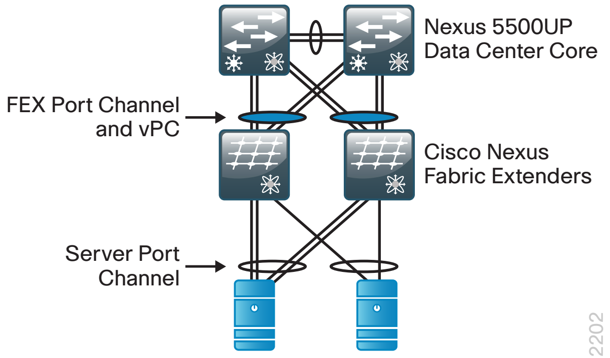

With Enhanced vPC, the dual-homed FEX uplinks are programmed with a port channel and vPC that connects it to both data core switches, and the Ethernet interfaces on the FEX connected to the server interfaces are programmed with a different port channel for the server port channel. The Cisco Nexus 5500 switches then automatically create a vPC to enable the server port channel that is connected to the dual-homed FEX pair. The result is a more resilient and simplified FEX deployment in the data center that can support single- and dual-homed servers with or without EtherChannel from the server.

Enhanced vPC also supports a dual-homed server with EtherChannel running FCoE. However, this may not be suitable for a high-bandwidth FCoE environment, because the FCoE traffic can only use a subset of the FEX uplinks to the data center core as shown in Figure 22. The FCoE traffic can only use the FEX-to–Cisco Nexus 5500 uplinks on the left side or right side, respectively, because SAN traffic must maintain SAN-A and SAN-B isolation and therefore cannot connect to both data center core switches. Non-FCoE

Ethernet traffic (for example, IP connectivity) from the dual-homed FEX can utilize all FEX-to–data center core uplinks, maximizing traffic load balancing and bandwidth.

Figure 22 - Enhanced vPC with FCoE traffic

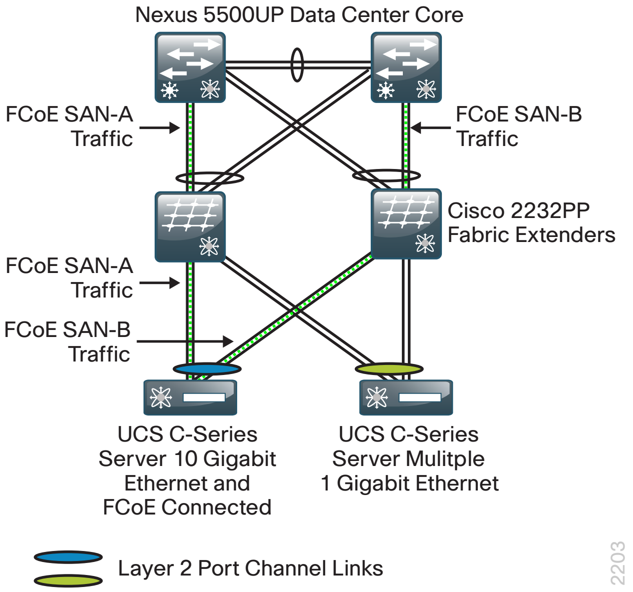
Nexus 5500UP Data Center Core
FCoE SAN-A Traffic
FCoE SAN-B Traffic
Cisco 2232PP Fabric Extenders
FCoE SAN-A Traffic
FCoE SAN-B Traffic
UCS C-Series Server 10 Gigabit Ethernet and FCoE Connected
UCS C-Series Server Mulitple 1 Gigabit Ethernet
Layer 2 Port Channel Links
2203

Third-Party Blade Server System Connectivity

Blade server systems are available from manufacturers other than Cisco. In the event you have a non–Cisco blade server system to connect to your data center, you have multiple options for connecting to your Cisco SBA data center design.

The first option is using a blade server system with a pass-through module that extends server interfaces directly out of the blade server chassis without using an internal switch fabric in the blade server system. When using pass-through modules, the server NIC connections can use the Cisco Nexus FEX for high-density port fan out and resilient connections, as shown in the figure below.

Figure 23 - Third-party blade server system with pass-through module

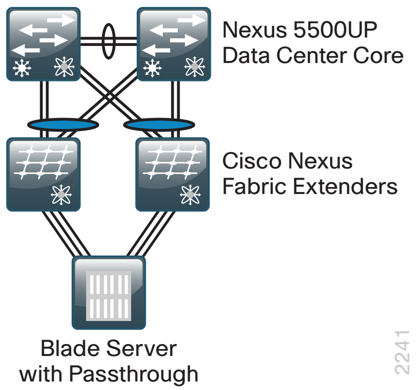
Nexus 5500UP Data Center Core
Cisco Nexus Fabric Extenders
Blade Server with Passthrough
2241

A second option for connecting a non-Cisco blade server system to the Cisco SBA data center involves a blade server system that has an integrated Ethernet switch. In this scenario, the integrated switch in the blade server chassis generates spanning-tree BPDUs and therefore cannot be connected to fabric extenders. Another consideration is that a blade server with an integrated switch generally uses a few high-speed 10-Gigabit Ethernet uplinks where direct connection to the Cisco Nexus 5500UP switch core, as shown in Figure 24, is recommended.

Figure 24 - Third-party blade server system with integrated switch

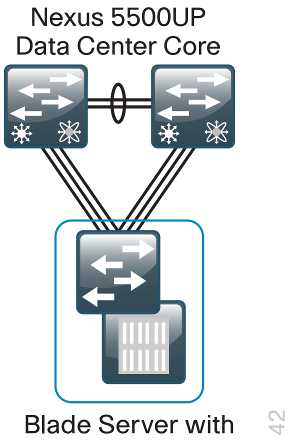

A third option is imbedding Cisco Nexus fabric extenders directly into the non–Cisco blade server system to connect to the Cisco SBA data center core, as shown in Figure 25. Although this option has not been tested and documented in the Cisco SBA data center deployment guide, it has proven to be a desirable connectivity option for many organizations.

Figure 25 - Non–Cisco blade server system with embedded Cisco Nexus fabric extenders

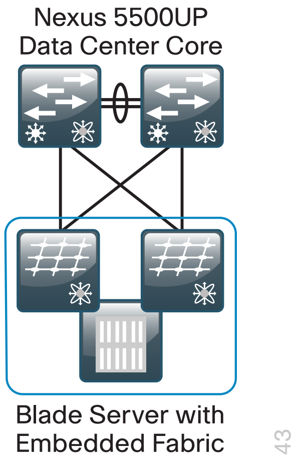

Summary

The compute connectivity options outlined in this chapter show how the Cisco SBA data center foundation design integrates with Cisco UCS to build a flexible and scalable compute connectivity. The data center architecture also provides support for resilient, non–Cisco server and blade system connectivity. For further detail on deploying Cisco UCS Server systems, please refer to the *Cisco SBA—Data Center Unified Computing System Deployment Guide*.

Notes

Network Security

Business Overview

In today's business environment, the data center contains some of the organization's most valuable assets. Customer and personnel records, financial data, email stores, and intellectual property must be maintained in a secure environment to assure confidentiality and availability. Additionally, portions of networks in specific business sectors may be subject to industry or government regulations that mandate specific security controls to protect customer or client information.

To protect the valuable electronic assets located in the data center, network security helps ensure the facility is protected from automated or human-operated snooping and tampering, and it helps prevent compromise of hosts by resource-consuming worms, viruses, or botnets.

Although worms, viruses, and botnets pose a substantial threat to centralized data—particularly from the perspective of host performance and availability—servers must also be protected from employee snooping and unauthorized access. Statistics have consistently shown that the majority of data loss and network disruptions have occurred as the result of human-initiated activity (intentional or accidental) carried out within the boundaries of the business's network.

Technology Overview

To minimize the impact of unwanted network intrusions, firewalls and intrusion prevention systems (IPSs) should be deployed between clients and centralized data resources.

Figure 26 - Deploy firewall inline to protect data resources

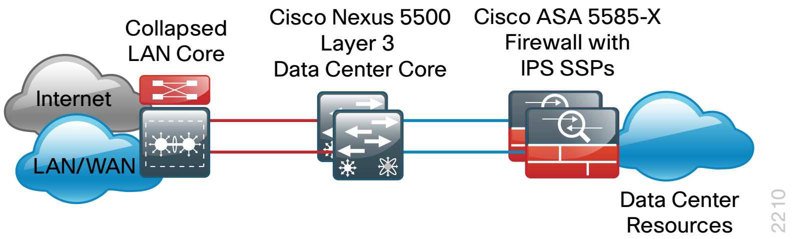
Collapsed LAN Core
Cisco Nexus 5500 Layer 3 Data Center Core
Cisco ASA 5585-X Firewall with IPS SSPs
Internet
LAN/WAN
Data Center Resources
2210

Because everything else outside the protected VLANs hosting the data center resources can be a threat, the security policy associated with protecting those resources has to include the following potential threat vectors.

Data center threat landscape:

- Internet
- Remote access and teleworker VPN hosts
- Remote office/branch networks
- Business partner connections
- Campus networks
- Unprotected data center networks
- Other protected data center networks

The data center security design employs a pair of Cisco Adaptive Security Appliance (ASA) 5585-X with SSP-20 firewall modules and matching IPS Security Service Processors (SSP) installed. This configuration provides up to 10 Gbps of firewall throughput. The IPS and firewall SSPs deliver 3 Gbps of concurrent throughput. There is a range of Cisco ASA 5585-X with IPS firewalls to meet your processing requirements.

All of the ports on modules installed in the Cisco ASA chassis are available to the firewall SSP, which offers a very flexible configuration. The Cisco ASA firewalls are dual-homed to the data center core Cisco Nexus 5500UP switches using two 10-Gigabit Ethernet links for resiliency. The pair of links on each Cisco ASA is configured as an EtherChannel, which provides load balancing as well as rapid and transparent failure recovery. The Cisco NX-OS Virtual Port Channel (vPC) feature on the Cisco Nexus 5500UP data core switches allow the firewall EtherChannel to span the two data center core switches (multichassis EtherChannel) but appear to be connected to a single upstream switch. This EtherChannel link is configured as a VLAN trunk in order to support access to multiple secure VLANs in the data center. One VLAN on the data center core acts as the outside VLAN for the firewall, and any hosts or servers that reside in that VLAN are outside the firewall and therefore receive no protection from Cisco ASA for attacks originating from anywhere else in the organization's network. Other VLANs on the EtherChannel trunk will be designated as being firewalled from all the other data center threat vectors or firewalled with additional IPS services.

The pair of Cisco ASAs is configured for firewall active/standby high availability operation to ensure that access to the data center is minimally impacted by outages caused by software maintenance or hardware failure. When Cisco ASA appliances are configured in active/standby mode, the standby appliance does not handle traffic, so the primary device must be sized to provide enough throughput to address connectivity requirements between the core and the data center. Although the IPS modules do not actively exchange state traffic, they participate in the firewall appliances' active/standby status by way of reporting their status to the firewall's status monitor. A firewall failover will occur if either the Cisco ASA itself has an issue or the IPS module becomes unavailable.

The Cisco ASAs are configured in routing mode; as a result, the secure network must be in a separate subnet from the client subnets. IP subnet allocation would be simplified if Cisco ASA were deployed in transparent mode; however, hosts might inadvertently be connected to the wrong VLAN, where they would still be able to communicate with the network, incurring an unwanted security exposure.

The data center IPSs monitor for and mitigate potential malicious activity that is contained within traffic allowed by the security policy defined on the Cisco ASAs. The IPS sensors can be deployed in promiscuous intrusion detection system (IDS) mode so that they only monitor and alert for abnormal traffic. The IPS modules can be deployed inline in IPS mode to fully engage their intrusion prevention capabilities, wherein they will block malicious traffic before it reaches its destination. The choice to have the sensor drop traffic or not is one that is influenced by several factors: risk tolerance for having a security incident, risk aversion for inadvertently dropping valid traffic, and other possibly externally driven reasons like compliance requirements for IPS. The ability to run in IDS mode or IPS is highly configurable to allow the maximum flexibility in meeting a specific security policy.

Security Topology Design

The Cisco SBA secure data center design provides two secure VLANs in the data center. The number of secure VLANs is arbitrary; the design is an example of how to create multiple secured networks to host services that require separation. High-value applications, such as Enterprise Resource Planning and Customer Relationship Management, may need to be separated from other applications in their own VLAN.

Figure 27 - Example design with secure VLANs

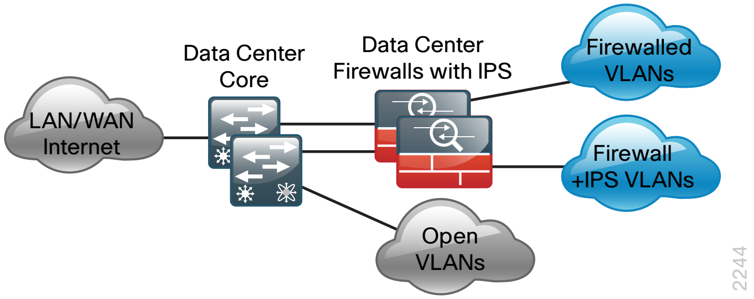

As another example, services that are indirectly exposed to the Internet (via a web server or other application servers in the Internet demilitarized zone) should be separated from other services, if possible, to prevent Internet-borne compromise of some servers from spreading to other services that are not exposed. Traffic between VLANs should be kept to a minimum, unless your security policy dictates service separation. Keeping traffic between servers intra-VLAN will improve performance and reduce the load on network devices.

For this deployment, open VLANs without any security policy applied are configured physically and logically on the data center core switches. For devices that need an access policy, they will be deployed on a VLAN behind the firewalls. Devices that require both an access policy and IPS traffic inspection will be deployed on a different VLAN that exists logically behind the Cisco ASAs. Because the Cisco ASAs are physically attached only to the data center core Nexus switches, these protected VLANs will also exist at Layer 2 on the data center core switches. All protected VLANs are logically connected via Layer 3 to the rest of the network through Cisco ASA and, therefore, are reachable only by traversing the appliance.

Security Policy Development

An organization should have an IT security policy as a starting point in defining its firewall policy. If there is no organization-wide security policy, it will be very difficult to define an effective policy for the organization while maintaining a secure computing environment.

To effectively deploy security between the various functional segments of a business's network, you should seek the highest level of detail possible regarding the expected network behaviors. If you have greater detail of the expectations, you will be better positioned to define a security policy that enables a business's application traffic and performance requirements while optimizing security.

Reader Tip

A detailed examination of regulatory compliance considerations exceeds the scope of this document; you should include industry regulation in your network security design. Non-compliance may result in regulatory penalties such as fines or suspension of business activity.

Network security policies can be broken down into two basic categories: whitelist policies and blacklist policies. A *whitelist policy* offers a higher implicit security posture, blocking all traffic except that which must be allowed (at a sufficiently granular level) to enable applications. Whitelist policies are generally better positioned to meet regulatory requirements because only traffic that must be allowed to conduct business is allowed. Other traffic is blocked and does not need to be monitored to assure that unwanted activity is not occurring. This reduces the volume of data that will be forwarded to an IDS or IPS, and also minimizes the number of log entries that must be reviewed in the event of an intrusion or data loss.

Figure 28 - Whitelist policy example

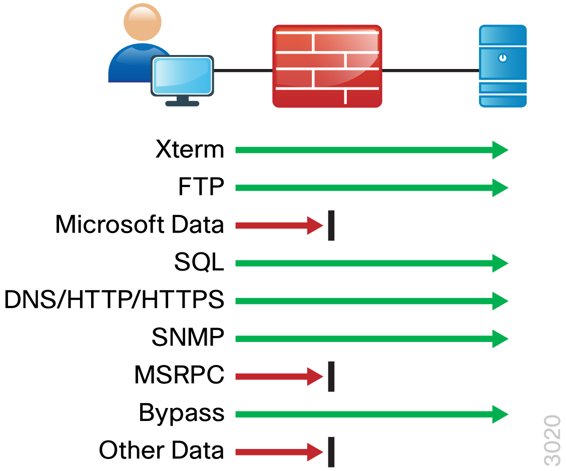

Inversely, a *blacklist policy* only denies traffic that specifically poses the greatest risk to centralized data resources. A blacklist policy is simpler to maintain and less likely to interfere with network applications. A whitelist policy is the best-practice option if you have the opportunity to examine the network's requirements and adjust the policy to avoid interfering with desired network activity.

Figure 29 - Blacklist policy example

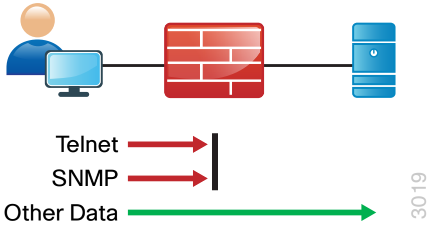

Cisco ASA firewalls implicitly end access lists with a deny-all rule. Blacklist policies include an explicit rule, prior to the implicit deny-all rule, to allow any traffic that is not explicitly allowed or denied.

Whether you choose a whitelist or blacklist policy basis, consider IDS or IPS deployment for controlling malicious activity on otherwise trustworthy application traffic. At a minimum, IDS or IPS can aid with forensics to determine the origin of a data breach. Ideally, IPS can detect and prevent attacks as they occur and provide detailed information to track the malicious activity to its source. IDS or IPS may also be required by the regulatory oversight to which a network is subject (for example, PCI 2.0).

A blacklist policy that blocks high-risk traffic offers a lower-impact—but less secure—option (compared to a whitelist policy) in cases where a detailed study of the network's application activity is impractical, or if the network availability requirements prohibit application troubleshooting. If identifying all of the application requirements is not practical, you can apply a blacklist policy with logging enabled to generate a detailed history of the policy. With details about its network's behavior in hand, an organization can more easily develop an effective whitelist policy.

Deployment Details

Data center security deployment is addressed in five discrete processes:

- “Configuring Cisco ASA Firewall Connectivity,” which configures network connections for the Cisco ASA firewalls on the Cisco Nexus 5500UP data center core.
- “Configuring the Data Center Firewall,” which configures Cisco ASA initial setup and the connections to the data center core.
- “Configuring Firewall High Availability,” which configures high availability active/standby state for the firewall pair.
- “Evaluating and Deploying Firewall Security Policy,” which outlines the process for identifying security policy needs and applying a configuration to meet requirements.
- “Deploying Cisco IPS,” which integrates connectivity and policy configuration in one process.

Process

Configuring Cisco ASA Firewall Connectivity

1. Configure firewall VLANs on Nexus 5500s
2. Configure port channels on core switches

Complete the following procedures to configure connectivity between the Cisco ASA chassis and the core. Note that this design describes a configuration wherein the Cisco ASA firewalls are connected to the Nexus 5500UP data center core switches by using a pair of 10-Gigabit Ethernet interfaces in an EtherChannel. The Cisco ASA firewall connects between the data center core–routed interface and the protected VLANs that also reside on the switches.

Connect the interfaces on the primary Cisco ASA firewall to the first Cisco Nexus 5500 data center core switch, and the secondary Cisco ASA firewall to the second Cisco Nexus 5500 data center core switch. Cisco ASA network ports are connected as follows:

- TenGigabitEthernet 0/8 connects to the Cisco Nexus 5500UP switch ethernet 1/1
- TenGigabitEthernet 0/9 connects to the Cisco Nexus 5500UP switch ethernet 1/2
- GigabitEthernet 0/1 connects via a crossover or straight-through Ethernet cable to the other firewall for the failover link

Table 7 - Data Center Firewall VLANs

VLAN	IP address	Trust state	Use
153	10.4.53.1 /25	Untrusted	Firewall to data center core routing
154	10.4.54.X /24	Trusted	Firewall protected VLAN
155	10.4.55.X /24	Trusted	Firewall + IPS protected VLAN

Procedure 1 Configure firewall VLANs on Nexus 5500s

Step 1: Configure the outside (untrusted) and inside (trusted) VLANs on the first Cisco Nexus 5500UP data center core switch.

```
vlan 153
  name FW_Outside
vlan 154
  name FW_Inside_1
vlan 155
  name FW_Inside_2
```

Step 2: Configure the Layer 3 SVI for VLAN 153 on the first Cisco Nexus 5500UP data center core switch. Set the HSRP address for the default gateway to 10.4.53.1 and the HSRP priority for this switch to 110.

```
interface Vlan153
  no shutdown
  description FW_Outside
  no ip redirects
  ip address 10.4.53.2/25
  ip router eigrp 100
  ip passive-interface eigrp 100
  ip pim sparse-mode
  hsrp 153
    priority 110
    ip 10.4.53.1
```

Step 3: Configure static routes pointing to the trusted subnets behind the Cisco ASA firewall on the first Cisco Nexus 5500UP data center core switch.

```
ip route 10.4.54.0/24 Vlan 153 10.4.53.126
ip route 10.4.55.0/24 Vlan 153 10.4.53.126
```

Step 4: Redistribute the trusted subnets into the existing EIGRP routing process on the first Cisco Nexus 5500UP data center core switch. This design uses route maps to control which static routes will be redistributed.

```
route-map static-to-eigrp permit 10
  match ip address 10.4.54.0/24
route-map static-to-eigrp permit 20
  match ip address 10.4.55.0/24
!
router eigrp 100
  redistribute static route-map static-to-eigrp
```

Step 5: Configure the outside (untrusted) and inside (trusted) VLANs on the second Cisco Nexus 5500UP data center core switch.

```
vlan 153
  name FW_Outside
vlan 154
  name FW_Inside_1
vlan 155
  name FW_Inside_2
```

Step 6: Configure the Layer 3 SVI for VLAN 153 on the second Cisco Nexus 5500UP data center core switch. Set the HSRP address for the default gateway to 10.4.53.1 and leave the HSRP priority for this switch at the default setting.

```
interface Vlan153
  no shutdown
  description FW_Outside
  no ip redirects
  ip address 10.4.53.3/25
  ip router eigrp 100
  ip passive-interface eigrp 100
  ip pim sparse-mode
  hsrp 153
    ip 10.4.53.1
```

Step 7: Configure static routes pointing to the trusted subnets behind the Cisco ASA firewall on the second Cisco Nexus 5500UP data center core switch.

```
ip route 10.4.54.0/24 Vlan 153 10.4.53.126
ip route 10.4.55.0/24 Vlan 153 10.4.53.126
```

Step 8: Redistribute the trusted subnets into the existing EIGRP routing process on the second Cisco Nexus 5500UP data center core switch. This design uses route maps to control which static routes will be redistributed.

```
route-map static-to-eigrp permit 10
  match ip address 10.4.54.0/24
route-map static-to-eigrp permit 20
  match ip address 10.4.55.0/24
!
router eigrp 100
   redistribute static route-map static-to-eigrp
```

Procedure 2 Configure port channels on core switches

The Cisco ASA firewalls protecting applications and servers in the data center will be dual-homed to each of the data center core Cisco Nexus 5500UP switches by using EtherChannel links.

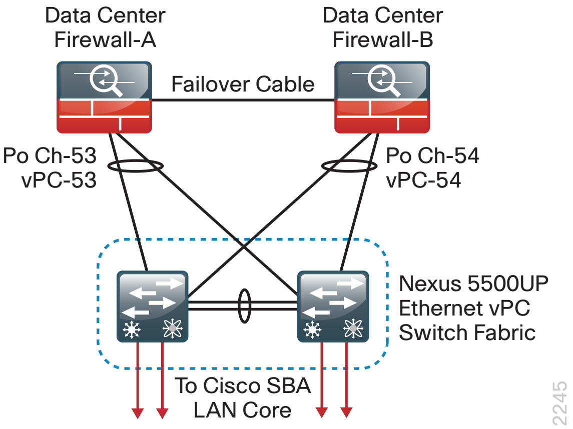

Dual-homed or multichassis EtherChannel connectivity to the Cisco Nexus 5500UP switches uses vPCs, which allow Cisco ASA to connect to both of the data center core switches with a single logical EtherChannel.

Step 1: Configure the physical interfaces that will make up the port channels on the first Cisco Nexus 5500UP data center core switch.

```
interface Ethernet1/1
  description DC5585a Ten0/8
  channel-group 53 mode active
!
interface Ethernet1/2
  description DC5585b Ten0/8
  channel-group 54 mode active
```

When you assign the channel group to a physical interface, it creates the logical EtherChannel (port-channel) interface that will be configured in the next step.

Step 2: Configure the logical port-channel interfaces on the first data center core switch. The physical interfaces tied to the port channel will inherit the settings from the logical port-channel interface. Assign the QoS policy created in Procedure 3, “Configure QoS policies,” to the port channel interfaces.

```
interface port-channel53
  switchport mode trunk
  switchport trunk allowed vlan 153-155
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 53
!
interface port-channel54
  switchport mode trunk
  switchport trunk allowed vlan 153-155
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 54
```

The port channels are created as vPC port channels, because the fabric interfaces are dual-homed EtherChannels to both Nexus 5500UP data center core switches.

Step 3: Apply following configuration to the second Cisco Nexus 5500UP data center core switch.

```
interface Ethernet1/1
  description DC5585a Ten0/9
  channel-group 53 mode active
!
interface Ethernet1/2
  description DC5585b Ten0/9
  channel-group 54 mode active
!
interface port-channel53
  switchport mode trunk
  switchport trunk allowed vlan 153-155
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 53
!
interface port-channel54
  switchport mode trunk
  switchport trunk allowed vlan 153-155
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS
  vpc 54
```

Process

Configuring the Data Center Firewall

1. Configure initial Cisco ASA settings
2. Configure firewall connectivity
3. Configure firewall static route to the core
4. Configure user authentication
5. Configure time synchronization and logging
6. Configure device management protocols

Configuration for this process is applied using CLI through the console port on the Cisco ASA firewall that is the primary unit of the high-availability pair. The standby unit synchronizes the configuration from the primary unit when it is programmed in the next process, “Configuring Firewall High Availability.”

The factory default password for enable mode is <CR>.

Table 8 - Cisco ASA 5500X firewall and IPS module addressing

ASA firewall failover status	Firewall IP address	IPS module IP address
Primary	10.4.53.126 /25	10.4.63.21 /24
Secondary	10.4.53.125 /25	10.4.63.23 /24

Table 9 - Common network services used in the deployment examples

Service	Address
Domain name	cisco.local
Active Directory, DNS, DHCP server	10.4.48.10
Cisco ACS server	10.4.48.15
NTP server	10.4.48.17

Procedure 1 Configure initial Cisco ASA settings

Connect to the console of the Cisco ASA firewall and perform the following global configuration.

Step 1: Select anonymous monitoring preference. When you enter configuration mode for an unconfigured unit, you are prompted for anonymous reporting. You are given a choice to enable anonymous reporting of error and health information to Cisco. Select the choice appropriate for your organization's security policy.

```
*************************** NOTICE ***************************

Help to improve the ASA platform by enabling anonymous
reporting, which allows Cisco to securely receive minimal
error and health information from the device. To learn more
about this feature, please visit: http://www.cisco.com/go/
smartcall

Would you like to enable anonymous error reporting to help
improve the product? [Y]es, [N]o, [A]sk later:N
```

Step 2: Configure the Cisco ASA firewall host name to make it easy to identify.

```
hostname DC5585ax
```

Step 3: Disable the dedicated management port. This design does not use it.

```
interface Management0/0
  shutdown
```

Step 4: Configure local user authentication.

```
Username [username] password [password]
```

Tech Tip

All passwords in this document are examples and should not be used in production configurations. Follow your company's policy, or—if no policy exists—create a password using a minimum of eight characters with a combination of uppercase, lowercase, and numbers.

Step 5: Configure enable password.

```
enable password [password]
```

Procedure 2 Configure firewall connectivity

Two 10-Gigabit Ethernet links connect each Cisco ASA chassis to the two core Cisco Nexus switches. The two interfaces are paired in a port channel group. Subinterfaces are created on the port channel for the outside VLAN 153 and all the protected VLANs inside (154 and 155). Each interface created will be assigned the correct VLAN, an appropriate name, a security level, and an IP address and netmask.

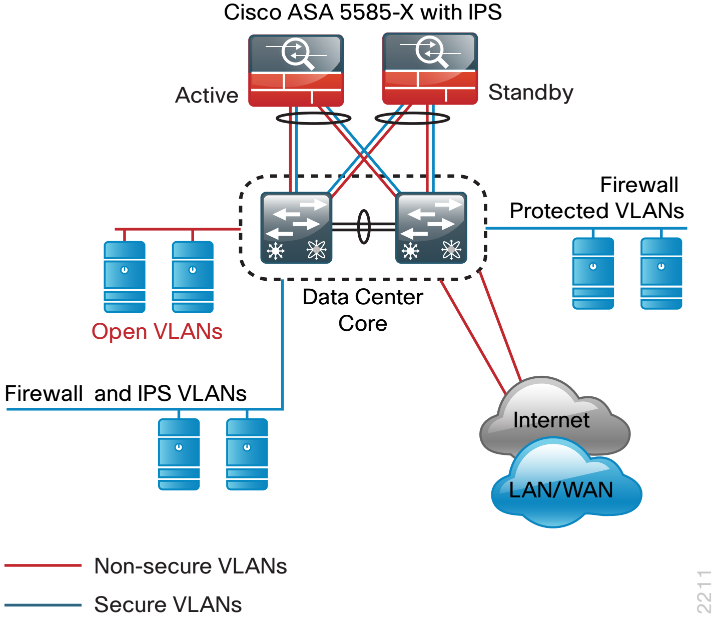

All interfaces on Cisco ASA have a security-level setting. The higher the number, the more trusted the interface, relative to other interfaces. By default, the inside interface is assigned 100, the highest security level. The outside interface is assigned 0. By default, traffic can pass from a high-security interface to a lower-security interface. In other words, traffic from an inside network is permitted to an outside network, but not conversely.

Step 1: Configure the port channel group by using the two 10-Gigabit Ethernet interfaces.

```
interface Port-channel10
 description ECLB Trunk to 5548 Switches
 no shutdown
!
interface TenGigabitEthernet0/8
 description Trunk to DC5548x eth1/1
 channel-group 10 mode passive
 no shutdown
!
interface TenGigabitEthernet0/9
 description Trunk to DC5548x eth1/2
 channel-group 10 mode passive
 no shutdown
```

Step 2: Configure the subinterfaces for the three VLANs: VLAN 153 outside, VLAN 154 inside the firewall, and VLAN 155 inside the firewall with IPS.

```
interface Port-channel10.153
 description DC VLAN Outside the FW
 vlan 153
 nameif outside
 security-level 0
 ip address 10.4.53.126 255.255.255.128 standby 10.4.53.125
 no shutdown
!
interface Port-channel10.154
 description DC VLAN Inside the Firewall
 vlan 154
 nameif DC-InsideFW
 security-level 75
 ip address 10.4.54.1 255.255.255.0 standby 10.4.54.2
 no shutdown
!
interface Port-channel10.155
 description DC VLAN Inside the FW w/ IPS
 vlan 155
 nameif DC-InsideIPS
 security-level 75
 ip address 10.4.55.1 255.255.255.0 standby 10.4.55.2
 no shutdown
```

Procedure 3 Configure firewall static route to the core

Because the Cisco ASAs are the gateway to the secure VLANs in the data center, the Cisco ASA pair is configured to use a static route to the HSRP address of the Cisco Nexus switches on outside VLAN 153.

Step 1: Configure the static route pointing to the data center core HSRP address on the Cisco ASA pair.

```
route outside 0.0.0.0 0.0.0.0 10.4.53.1 1
```

Procedure 4 Configure user authentication

(Optional)

If you want to reduce operational tasks per device, configure centralized user authentication by using the TACACS+ protocol to authenticate management logins on the infrastructure devices to the AAA server.

As networks scale in the number of devices to maintain, the is an operational burden to maintain local user accounts on every device also scales. A centralized AAA service reduces operational tasks per device and provides an audit log of user access for security compliance and root-cause analysis. When AAA is enabled for access control, all management access to the network infrastructure devices (SSH and HTTPS) is controlled by AAA.

Reader Tip

The AAA server used in this architecture is the Cisco Secure Access Control Server (ACS). Configuration of Cisco Secure ACS is discussed in the *Cisco SBA—Borderless Networks Device Management Using ACS Deployment Guide*.

TACACS+ is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database was defined already to provide a fallback authentication source in case the centralized TACACS+ server is unavailable.

Step 1: Configure the TACACS+ server.

```
aaa-server AAA-SERVER protocol tacacs+
aaa-server AAA-SERVER (outside) host 10.4.48.15 SecretKey
```

Step 2: Configure the appliance's management authentication to use the TACACS+ server first, and then the local user database if the TACACS+ server is unavailable.

```
aaa authentication enable console AAA-SERVER LOCAL
aaa authentication ssh console AAA-SERVER LOCAL
aaa authentication http console AAA-SERVER LOCAL
aaa authentication serial console AAA-SERVER LOCAL
```

Step 3: Configure the appliance to use AAA to authorize management users.

```
aaa authorization exec authentication-server
```

Tech Tip

User authorization on the Cisco ASA firewall, unlike Cisco IOS devices, does not automatically present the user with the enable prompt if they have a privilege level of 15.

Procedure 5 Configure time synchronization and logging

Logging and monitoring are critical aspects of network security devices to support troubleshooting and policy-compliance auditing.

NTP is designed to synchronize time across a network of devices. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the organization's network.

Network devices should be programmed to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source.

There is a range of detail that can be logged on the appliance. Informational-level logging provides the ideal balance between detail and log-message volume. Lower log levels produce fewer messages, but they do not produce enough detail to effectively audit network activity. Higher log levels produce a larger volume of messages, but do not add sufficient value to justify the number of messages logged.

Step 1: Configure the NTP server IP address.

```
ntp server 10.4.48.17
```

Step 2: Configure the time zone.

```
clock timezone PST -8 0
clock summer-time PDT recurring
```

Step 3: Configure which logs to store on the appliance.

```
logging enable
logging buffered informational
```

Procedure 6 Configure device management protocols

Cisco Adaptive Security Device Manager (ASDM) requires that the appliance's HTTPS server be available. Be sure that the configuration includes networks where administrative staff has access to the device through Cisco ASDM; the appliance can offer controlled Cisco ASDM access for a single address or management subnet (in this case, 10.4.48.0/24).

HTTPS and SSH are more secure replacements for the HTTP and Telnet protocols. They use SSL and TLS to provide device authentication and data encryption.

Use SSH and HTTPS protocols in order to more securely manage the device. Both protocols are encrypted for privacy, and the unsecure protocols—Telnet and HTTP—are turned off.

SNMP is enabled to allow the network infrastructure devices to be managed by an NMS. SNMPv2c is configured for a read-only community string.

Step 1: Allow internal administrators to remotely manage the appliance over HTTPS and SSH.

```
domain-name cisco.local
http server enable
http 10.4.48.0 255.255.255.0 outside
ssh 10.4.48.0 255.255.255.0 outside
ssh version 2
```

Step 2: Configure the appliance to allow SNMP polling from the NMS.

```
snmp-server host outside 10.4.48.35 community [cisco]
snmp-server community [cisco]
```

Process

Configuring Firewall High Availability

1. Configure the primary appliance for HA
2. Configure the secondary Cisco ASA for HA

Cisco ASAs are set up as a highly available active/standby pair. Active/standby is used, rather than an active/active configuration, because this allows the same appliance to be used for firewall and VPN services if required in the future (VPN functionality is disabled on the appliance in active/active configuration). In the event that the active appliance fails or needs to be taken out of service for maintenance, the secondary appliance assumes all active firewall and IPS functions. In an active/standby configuration, only one device is passing traffic at a time; thus, the Cisco ASAs must be sized so that the entire traffic load can be handled by either device in the pair.

Both units in the failover pair must be the same model, with identical feature licenses and IPS (if the software module is installed). For failover to be enabled, the secondary ASA unit needs to be powered up and cabled to the same networks as the primary unit.

One interface on each appliance is configured as the state-synchronization interface, which the appliances use to share configuration updates, determine which device in the high availability pair is active, and exchange state information for active connections. The failover interface carries the state synchronization information. All session state is replicated from the primary to the secondary unit through this interface. There can be a substantial amount of data, and it is recommended that this be a dedicated interface.

By default, the appliance can take from 2 to 25 seconds to recover from a failure. Tuning the failover poll times can reduce that to 0.5 to 5 seconds. On an appropriately sized appliance, the poll times can be tuned down without performance impact to the appliance, which minimizes the downtime a user experiences during failover. It is recommended that you do not reduce the failover timer intervals below the values in this guide.

Procedure 1 Configure the primary appliance for HA

Step 1: Enable failover on the primary appliance, and then assign it as the primary unit.

```
failover
failover lan unit primary
```

Step 2: Configure the failover interface. Enter a key for the failover that you will later enter on the secondary appliance to match.

```
failover lan interface failover GigabitEthernet0/1
failover key [key]
failover replication http
failover link failover GigabitEthernet0/1
```

Step 3: If you want to speed up failover in the event of a device or link failure, you can tune the failover timers. With the default setting, depending on the failure, Cisco ASA can take from 2 to 25 seconds to fail over to the standby unit. Tuning the failover poll times can reduce that to 0.5 to 5 seconds, depending on the failure.

On an appliance with low to average load, the poll times can be tuned down without performance impact.

```
failover polltime unit msec 200 holdtime msec 800
failover polltime interface msec 500 holdtime 5
```

Step 4: Configure the failover interface IP address.

```
failover interface ip failover 10.4.53.130 255.255.255.252 
standby 10.4.53.129
```

Step 5: Enable the failover interface.

```
interface GigabitEthernet0/1
 no shutdown
```

Step 6: Configure failover to monitor the inside and outside interfaces so that the active firewall will defer to the standby firewall if connectivity is lost on the data center VLANs.

```
monitor-interface outside
monitor-interface DC-InsideFW
monitor-interface DC-InsideIPS
```

Procedure 2 Configure the secondary Cisco ASA for HA

Step 1: On the secondary Cisco ASA, enable failover and assign it as the secondary unit.

```
failover
failover lan unit secondary
```

Step 2: Configure the failover interface.

```
failover lan interface failover GigabitEthernet0/1
failover key [key]
failover replication http
failover link failover GigabitEthernet0/1
```

Step 3: Configure the failover interface IP address.

```
failover interface ip failover 10.4.53.130 255.255.255.252 
standby 10.4.53.129
```

Step 4: Enable the failover interface.

```
interface GigabitEthernet0/1
 no shutdown
```

Step 5: Verify high availability standby synchronization between the Cisco ASA devices. On the CLI of the primary appliance, issue the **show failover state** command.

```
DC5585ax# show failover state

             State          Last Failure Reason   Date/Time
This host -  Primary
             Active         None
Other host - Secondary
             Standby Ready  None        15:18:12 UTC May 25 2012

====Configuration State===
        Sync Done
====Communication State===
        Mac set
```

Process

Evaluating and Deploying Firewall Security Policy

1. Evaluate security policy requirements
2. Deploy the appropriate security policy

This process describes the steps required to evaluate which type of policy fits an organization's data center security requirements and provides the procedures necessary to apply these policies.

Procedure 1 Evaluate security policy requirements

Step 1: Evaluate security policy requirements by answering the following questions:

- What applications will be served from the secure data center?
- Can the applications' traffic be characterized at the protocol level?
- Is a detailed description of application behavior available to facilitate troubleshooting if the security policy interferes with the application?
- What is the network's baseline performance expectation between the controlled and uncontrolled portions of the network?
- What is the peak level of throughput that security controls will be expected to handle, including bandwidth-intensive activity such as workstation backups or data transfers to a secondary data replication site?

Step 2: For each data center VLAN, determine which security policy enables application requirements. Each VLAN that requires a firewall needs either a permissive (blacklist) or restrictive (whitelist) security policy.

Procedure 2 Deploy the appropriate security policy

Network security policy configuration is fairly arbitrary to suit the policy and management requirements of an organization. Thus, examples here should be used as a basis for security policy configuration.

Option 1. Deploy a whitelist security policy

A basic whitelist data-service policy can be applied to allow common business services such as HTTP, HTTPS, DNS, and other services typically seen in Microsoft-based networks.

Step 1: Control access so only specific hosts may be accessed.

```
object network BladeWeb1Secure
 host 10.4.54.100
 object network BladeWeb2Secure
 host 10.4.55.100
!
object-group network Application-Servers
 description HTTP, HTTPS, DNS, MSExchange
 network-object object BladeWeb1Secure
 network-object object BladeWeb2Secure
!
object-group service MS-App-Services
 service-object tcp destination eq domain
 service-object tcp destination eq www
 service-object tcp destination eq https
 service-object tcp destination eq netbios-ssn
 service-object udp destination eq domain
 service-object udp destination eq nameserver
 service-object udp destination eq netbios-dgm
 service-object udp destination eq netbios-ns
!
 access-list global_access extended permit object-group MS-App-Services any object-group Application-Servers
```

Step 2: Specify which resources certain users (for example, IT management staff or network users) can use to access management resources. In this example, management hosts in the IP address range 10.4.48.224–255 are allowed SSH and SNMP access to data center subnets.

```
object network Secure-Subnets
 subnet 10.4.54.0 255.255.255.0
object network SecureIPS-Subnets
 subnet 10.4.55.0 255.255.255.0
!
object network Mgmt-host-range
 range 10.4.48.224 10.4.48.254
object-group network DC_Secure_Subnet_List
 network-object object Secure-Subnets
 network-object object SecureIPS-Subnets
object-group service Mgmt-Traffic
 service-object tcp destination eq ssh
 service-object udp destination eq snmp
access-list global_access extended permit object-group Mgmt-
Traffic object Mgmt-host-range object-group DC_Secure_Subnet_
List
```

Step 3: If you want to allow access to an application for firewall policy troubleshooting, configure a bypass rule. A bypass rule allows wide-open access to hosts that are added to the appropriate network object group. The bypass rule must be carefully defined to avoid opening access to hosts or services that must otherwise be blocked. In a whitelist policy, the bypass rule is typically disabled, and it is only called into use whenever firewall policy troubleshooting is required to allow access to an application.

The following policy defines two hosts and applies them to the bypass rule.

```
object-group network Bypass-Rule
 description Open Policy for Server Access
 network-object object BladeWeb1Secure
 network-object object BladeWeb2Secure
access-list global_access extended permit ip any object-group 
Bypass-Rule
```

This disables the bypass rule:

```
access-list global_access extended permit ip any object-group 
Bypass-Rule  inactive
```

Tech Tip

The bypass rule group is useful for troubleshooting or providing temporary access to services on the host that must be opened for maintenance or service migration. It is typically disabled unless being used for troubleshooting.

Step 4: Save your Cisco ASA firewall configuration.

```
copy running-config startup-config
```

Option 2. Deploy a blacklist security policy

If an organization does not have the desire or resources to maintain a granular, restrictive policy to control access between centralized data and the user community, a simpler, easy-to-deploy policy that limits only the highest-risk traffic may be more attractive. This policy is typically configured such that only specific services' access is blocked; all other traffic is permitted.

Step 1: Allow SNMP queries and SSH requests from a specific address range that will be allocated for IT staff. Network administrative users may need to issue SNMP queries from desktop computers to monitor network activity and SSH to connect to devices.

```
object network Secure-Subnets
 subnet 10.4.54.0 255.255.255.0
object network SecureIPS-Subnets
 subnet 10.4.55.0 255.255.255.0
!
object network Mgmt-host-range
 range 10.4.48.224 10.4.48.254
object-group network DC_Secure_Subnet_List
 network-object object Secure-Subnets
 network-object object SecureIPS-Subnets
object-group service Mgmt-Traffic
 service-object tcp destination eq ssh
 service-object udp destination eq snmp
access-list global_access extended permit object-group Mgmt-Traffic object Mgmt-host-range object-group DC_Secure_Subnet_List
```

Step 2: Block Telnet, SSH, and SNMP to all other hosts.

```
access-list global_access extended deny object-group Mgmt-Traffic any any
```

Step 3: Configure a rule to permit application traffic through to the servers in the secure server subnets, that was not specifically denied by the blacklist rule in Step 2. Note that logging is disabled on this policy to prevent the firewall from having to log all accesses to the server network.

```
access-list global_access extended permit ip any object-group DC_Secure_Subnet_List log disable
```

Step 4: Save your Cisco ASA firewall configuration.

```
copy running-config startup-config
```

Process

Deploying Cisco IPS

1. Apply initial configuration
2. Complete basic configuration
3. Configure signature updates

From a security standpoint, intrusion detection systems (IDS) and intrusion prevention systems (IPS) are complementary to firewalls because firewalls are generally access-control devices that are built to block access to an application or host. In this way, a firewall can be used to remove access to a large number of application ports, reducing the threat to the servers. IDS and IPS sensors look for attacks in network and application traffic that is permitted to go through the firewall. If it detects an attack, the IDS sensor generates an alert to inform the organization about the activity. IPS is similar in that it generates alerts due to malicious activity and, additionally, it can apply an action to block the attack before it reaches the destination.

Promiscuous versus Inline Deployment Modes

There are two primary deployment modes when using IPS sensors: *promiscuous* (IDS) or *inline* (IPS). There are specific reasons for each deployment model based on risk tolerance and fault tolerance.

- In promiscuous mode (IDS), the sensor inspects copies of packets, which prevents it from being able to stop a malicious packet when it sees one.
- An IDS sensor must use another inline enforcement device in order to stop malicious traffic. This means that for activity such as single-packet attacks (for example, slammer worm over User Datagram Protocol[UDP]), an IDS sensor could not prevent the attack from occurring. However, an IDS sensor can offer great value when identifying and cleaning up infected hosts.
- In an inline (IPS) deployment, because the packet flow is sent through the sensor and returned to Cisco ASA, the sensor inspects the actual data packets.
- The advantage IPS mode offers is that when the sensor detects malicious behavior, the sensor can simply drop the malicious packet. This allows the IPS device a much greater capacity to actually prevent attacks.

Deployment Considerations

Use IDS when you do not want to impact the availability of the network or create latency issues. Use IPS when you need higher security than IDS can provide, and when you need the ability to drop malicious data packets.

The secure data center design using a Cisco ASA 5585-X with IPS implements a policy for IPS, which sends all traffic to the IPS module inline.

Your organization may choose an IPS or IDS deployment depending on regulatory and application requirements. It is very easy to initially deploy an IDS, or promiscuous, design and then move to IPS after you understand the traffic and performance profile of your network and you are comfortable that production traffic will not be affected.

Procedure 1 Apply initial configuration

Use the sensor's CLI in order to set up basic networking information, specifically, the IP address, gateway address, and access lists that allow remote access. After these critical pieces of data are entered, the rest of the configuration is accomplished by using IPS Device Manager (IDM), the embedded GUI console. Unlike the Cisco ASA firewalls used in the Cisco SBA design, IPS modules use an out-of-band management connection for configuration and monitoring. The sensor's management port is connected to the data center management VLAN configured in earlier in this guide in Procedure 4, "Configure switch access ports," so that the sensors can route to or directly reach the management station.

Step 1: Connect to the IPS SSP console through the serial console on the IPS SSP module on the front panel of the Cisco ASA 5585-X firewall.

Tech Tip

You can also gain access to the console on the IPS SSP by using the **session 1** command from the CLI of the Cisco ASA's SSP.

Step 2: Log in to the IPS device. The default username and password are both cisco. You will be prompted to change the login password for the "cisco" user.

Step 3: At the IPS module's CLI, launch the System Configuration Dialogue.

```
sensor# setup
```

The IPS module enters the interactive setup.

Step 4: Define the IPS module's host name. Note that unlike Cisco IOS devices where the host name instantly changes the CLI prompt to reflect the new host name, the IPS will display the new host name for the CLI prompt upon the next login to the sensor.

```
--- Basic Setup ---
--- System Configuration Dialog ---
At any point you may enter a question mark ‘?’ for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets ‘[]’.
Current time: Mon Oct 12 23:31:38 2009
Setup Configuration last modified: Mon Oct 12 23:22:27 2009
Enter host name [sensor]: IPS-SSP20-A
```

Step 5: Define the IP address and gateway address for the IPS module's external management port.

```
Enter IP interface [192.168.1.62/24,192.168.1.250]:
10.4.63.21/24,10.4.63.1
```

Step 6: Define the access list, and then press **Enter**. This controls management access to the IPS module. For this network, all addresses in the headquarters subnet (10.4.0.0/16) are allowed. Press **Enter** at a blank Permit prompt to go to the next step.

```
Modify current access list?[no]: yes
Current access list entries:
  No entries
Permit: 10.4.0.0/16
```

Step 7: Accept the default answer (no) for the next three questions.

```
Use DNS server for Global Correlation? [no]:
Use HTTP proxy server for Global Correlation? [no]:
Modify system clock settings?[no]:
```

Note the following:

- Global correlation is disabled until later in the configuration process.
- You will configure time details in the IPS module's GUI console.

Step 8: Accept the default answer (off) for the option to participate in the SensorBase Network.

```
Participation in the SensorBase Network allows Cisco to
collect aggregated statistics about traffic sent to your IPS.
SensorBase Network Participation level? [off]:
```

The IPS SSP displays your configuration and a brief menu with four options.

Step 9: In the System Configuration dialog, save your configuration and exit setup by entering **2**.

```
The following configuration was entered.
[removed for brevity]
exit
[0] Go to the command prompt without saving this
configuration.
[1] Return to setup without saving this configuration.
[2] Save this configuration and exit setup.
[3] Continue to Advanced setup.
Enter your selection [3]: 2
Warning: DNS or HTTP proxy is required for global correlation
inspection and reputation filtering, but no DNS or proxy
servers are defined.
--- Configuration Saved ---
Complete the advanced setup using CLI or IDM.
To use IDM, point your web browser at https://<sensor-ip-
address>.
```

Step 10: Repeat this procedure for the IPS sensor installed in the other Cisco ASA chassis. In Step 4, be sure to use a different host name (IPS-SSP20-B) and in Step 5, be sure to use a different IP address (10.4.63.23) on the other sensor's management interface.

Procedure 2 Complete basic configuration

After the basic setup in the System Configuration dialog is complete, you will use the startup wizard in the integrated management tool, Cisco ASDM, to complete the remaining tasks in order to configure a basic setup:

- Configure time settings
- Configure DNS and NTP servers
- Define a basic IDS configuration
- Configure inspection service rule policy
- Assign interfaces to virtual sensors

Using ASDM to configure the IPS module operation allows you to set up the communications path from the firewall to the IPS module, as well as configure the IPS module settings.

Connect to the sensor by navigating to the firewall outside interface programmed in Step 2 of the "Configure firewall connectivity" procedure, using a secure HTTP session (https://10.4.53.126). Next, click **Run ASDM**, which will run ASDM from a Java Web Start application; alternatively, you can choose **Install ASDM Launcher and Run ASDM**, which will allow you to connect to multiple security appliances.

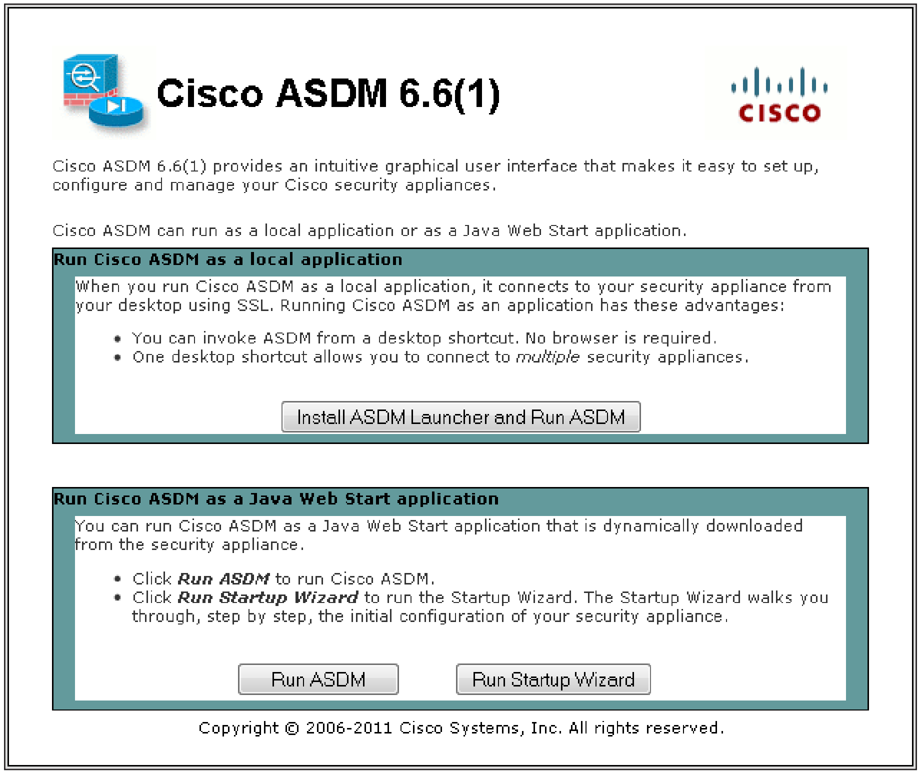

Cisco ASDM 6.6(1)

Cisco ASDM 6.6(1) provides an intuitive graphical user interface that makes it easy to set up, configure and manage your Cisco security appliances.

Cisco ASDM can run as a local application or as a Java Web Start application.

Run Cisco ASDM as a local application

When you run Cisco ASDM as a local application, it connects to your security appliance from your desktop using SSL. Running Cisco ASDM as an application has these advantages:

- You can invoke ASDM from a desktop shortcut. No browser is required.
- One desktop shortcut allows you to connect to *multiple* security appliances.

Run Cisco ASDM as a Java Web Start application

You can run Cisco ASDM as a Java Web Start application that is dynamically downloaded from the security appliance.

- Click ***Run ASDM*** to run Cisco ASDM.
- Click ***Run Startup Wizard*** to run the Startup Wizard. The Startup Wizard walks you through, step by step, the initial configuration of your security appliance.

Step 1: Enter the username and password configured for the Cisco ASA firewall in Step 4 of the “Configure initial Cisco ASA settings” procedure.

Step 2: Navigate to the **Intrusion Prevention** tab in Cisco ASDM, enter the required connection information for IPS-SSP20-A access, and then click **Continue**.

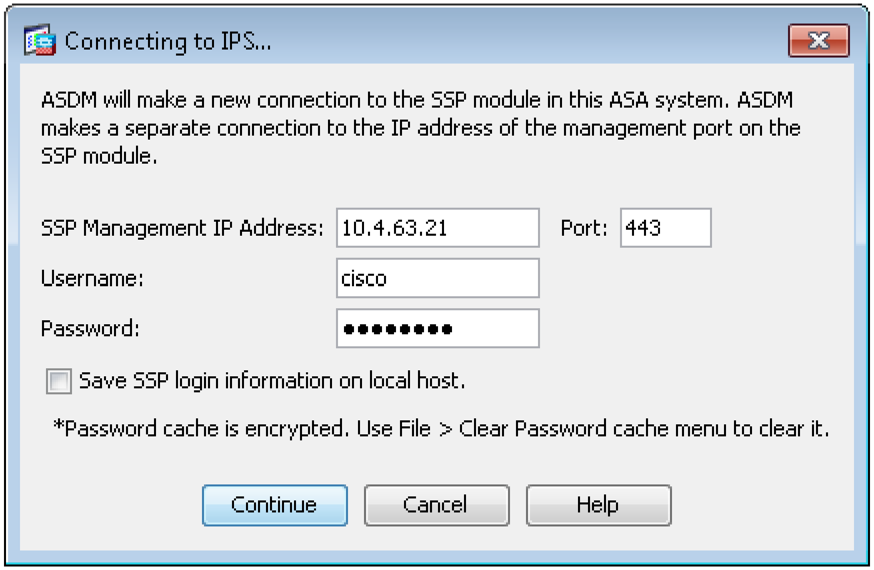

Connecting to IPS...

ASDM will make a new connection to the SSP module in this ASA system. ASDM makes a separate connection to the IP address of the management port on the SSP module.

SSP Management IP Address: 10.4.63.21 Port: 443

Username: cisco

Password: ●●●●●●●●●

□ Save SSP login information on local host.

*Password cache is encrypted. Use File > Clear Password cache menu to clear it.

ASDM will download the IPS information from the appliance for IPS-SSP20-A.

Step 3: Click **Configuration**, navigate to the **IPS** tab, and then click **Launch Startup Wizard**.

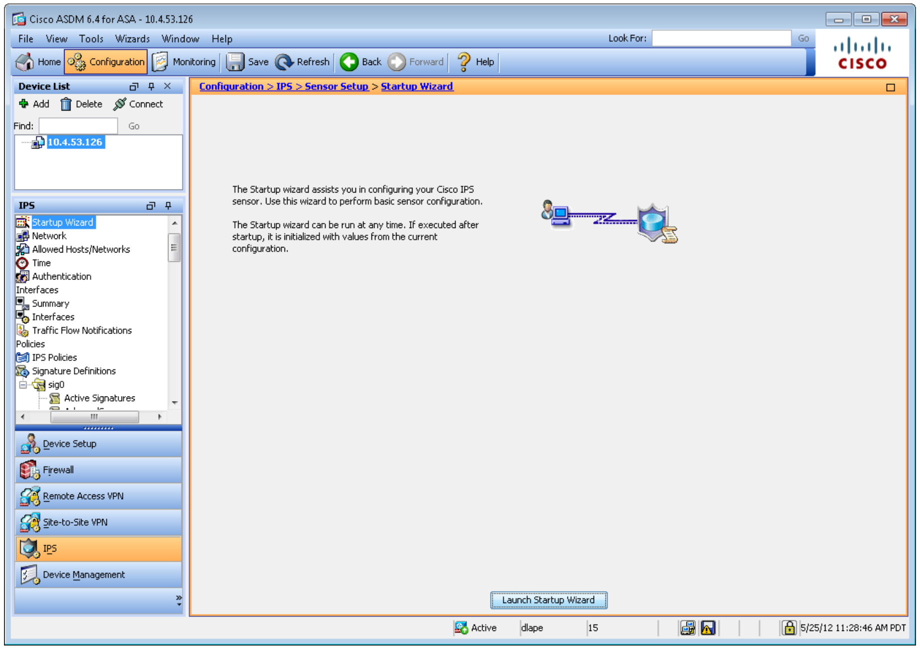

Configuration > IPS > Sensor Setup > Startup Wizard

The Startup wizard assists you in configuring your Cisco IPS sensor. Use this wizard to perform basic sensor configuration.

The Startup wizard can be run at any time. If executed after startup, it is initialized with values from the current configuration.

Step 4: Review the Startup Wizard Introduction, and then click **Next**.

Step 5: On the Sensor Setup page, configure the DNS server address, time zone, and NTP server address.

Tech Tip

NTP is particularly important for security event correlation if you use a Security Event Information Manager product to monitor security activity on your network.

Step 6: If necessary for your time zone, select **Enable Summertime**. Ensure that **Authenticated NTP** is not selected, and then click **Next**.

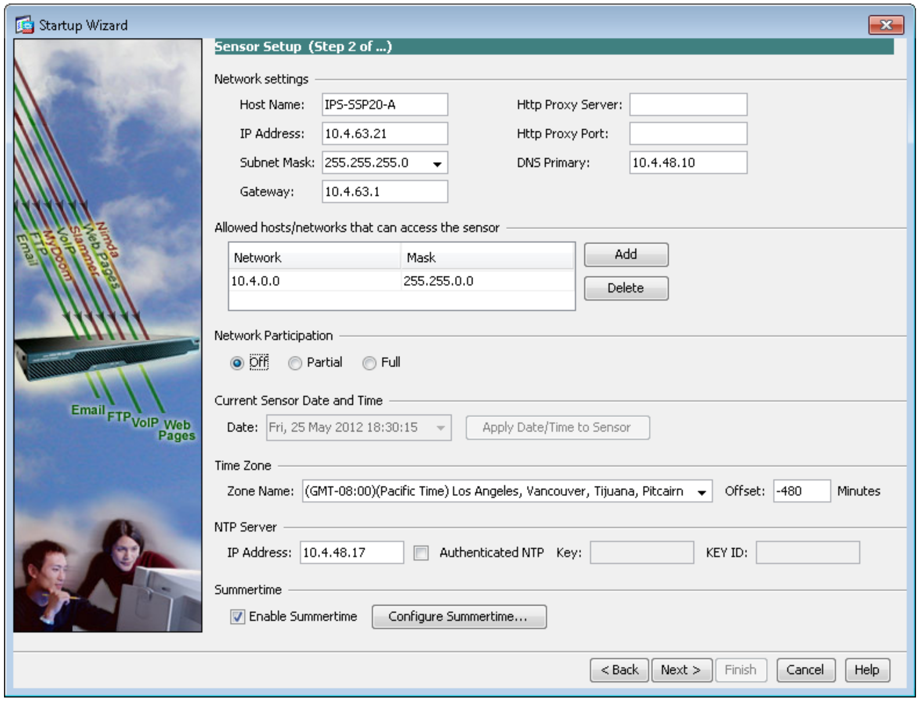

Sensor Setup (Step 2 of ...)

Network settings

Host Name: IPS-SSP20-A

IP Address: 10.4.63.21

Subnet Mask: 255.255.255.0

Gateway: 10.4.63.1

Http Proxy Server:

Http Proxy Port:

DNS Primary: 10.4.48.10

Allowed hosts/networks that can access the sensor

Network	Mask
10.4.0.0	255.255.0.0

Add Delete

Network Participation

● Off ○ Partial ○ Full

Current Sensor Date and Time

Date: Fri, 25 May 2012 18:30:15 Apply Date/Time to Sensor

Time Zone

Zone Name: (GMT-08:00)(Pacific Time) Los Angeles, Vancouver, Tijuana, Pitcairn Offset: -480 Minutes

NTP Server

IP Address: 10.4.48.17 □ Authenticated NTP Key: KEY ID:

Summertime

☑ Enable Summertime Configure Summertime...

You must now decide the sensor mode. In IPS mode, the sensor is inline in the traffic path. In this mode, the sensor inspects—and can drop—traffic that is malicious. Alternatively, in IDS mode, a copy of the traffic is passively sent to the sensor and the sensor inspects—and can send alerts about—traffic that is malicious. IPS mode provides more protection from Internet threats and has a low risk of blocking important traffic at this point in the network, particularly when it is coupled with reputation-based technologies. You can deploy IDS mode as a temporary solution to see what kind of impact IPS would have on the network and what traffic would be stopped. After you understand the impact on your network's performance and after you perform any necessary tuning, you can easily change the sensor to IPS mode.

This procedure assigns IPS mode.

Step 7: On the Startup Wizard: Virtual Sensors page, click **Next**.

Step 8: On the Startup Wizard: Traffic Allocation page, click **Add**.

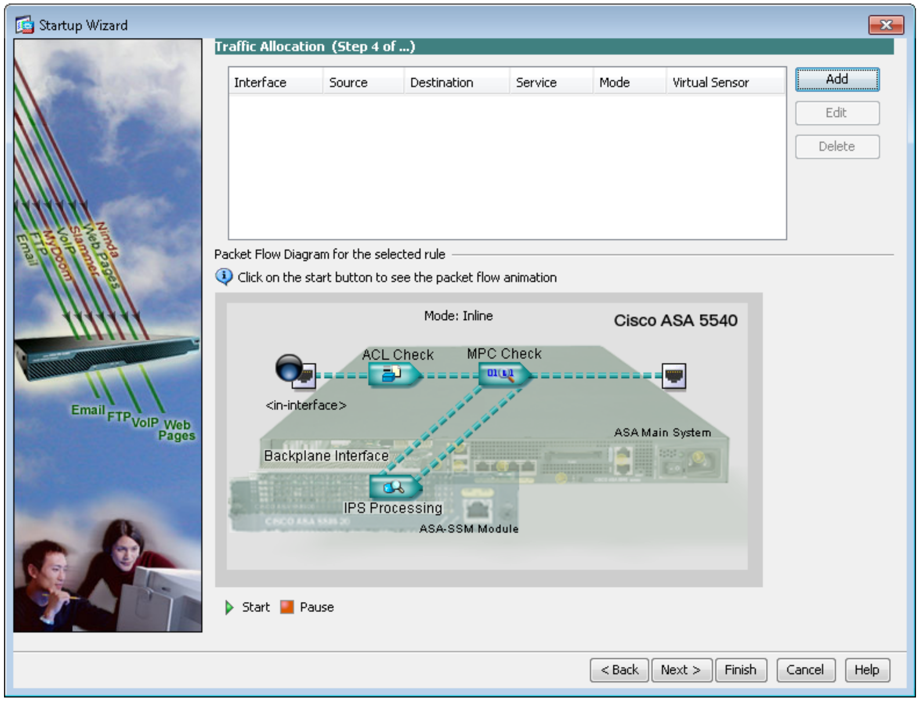

Startup Wizard

Traffic Allocation (Step 4 of ...)

Interface	Source	Destination	Service	Mode	Virtual Sensor

Packet Flow Diagram for the selected rule

Click on the start button to see the packet flow animation

Step 9: In the Specify traffic for IPS Scan window, in the Interface list, choose **DC-InsideIPS**, and next to Traffic Inspection Mode, select **Inline**, and then click **OK.** Note that if the Cisco ASA firewall already has a default Traffic Allocation policy, IDM displays a warning that “The Service Rule Policy you are trying to create already exists.” You can cancel the window and proceed to the next step if you receive this warning.

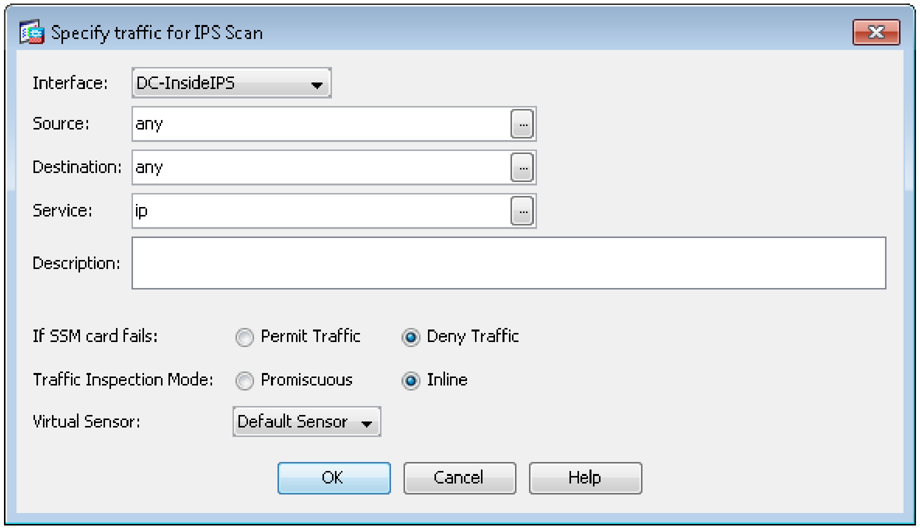

Specify traffic for IPS Scan

Interface: DC-InsideIPS

Source: any

Destination: any

Service: ip

Description:

If SSM card fails: ○ Permit Traffic ● Deny Traffic

Traffic Inspection Mode: ○ Promiscuous ● Inline

Virtual Sensor: Default Sensor

Step 10: On the Traffic Allocation page, in the Packet Flow Diagram for the selected Rule panel, verify the traffic allocation configuration by clicking **Start**. The animation illustrates a packet being copied to the IPS module and the egress interface. The animation may display an incorrect platform compared to the one you are configuring.

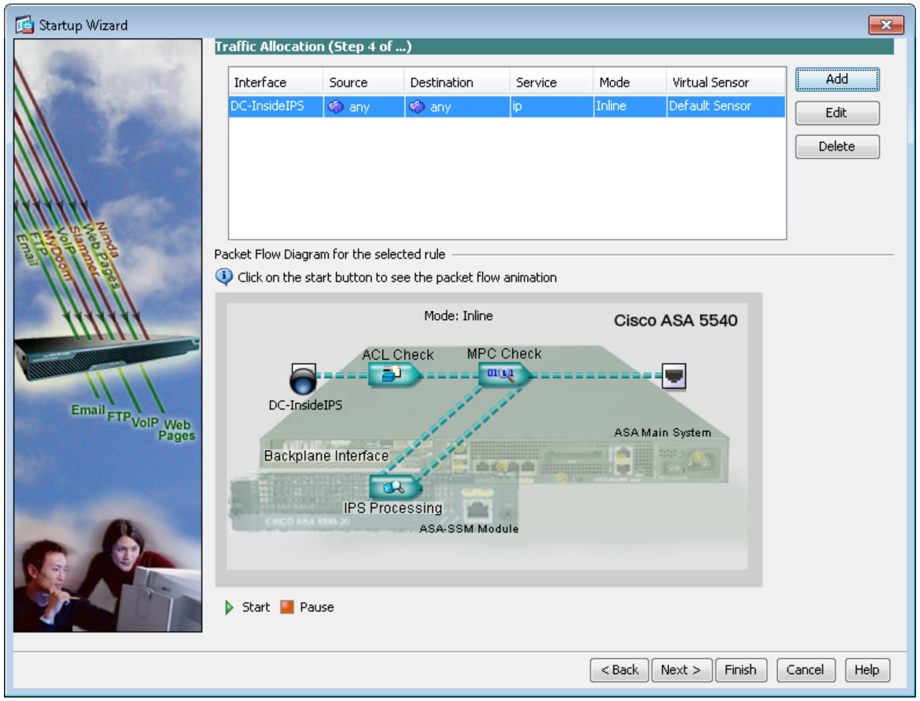

Startup Wizard

Traffic Allocation (Step 4 of ...)

Interface	Source	Destination	Service	Mode	Virtual Sensor
DC-InsideIPS	any	any	ip	Inline	Default Sensor

Packet Flow Diagram for the selected rule

Click on the start button to see the packet flow animation

Step 11: On the Startup Wizard page, click **Finish**, and then click **Yes** when you are prompted if you want to commit your changes to the sensor.

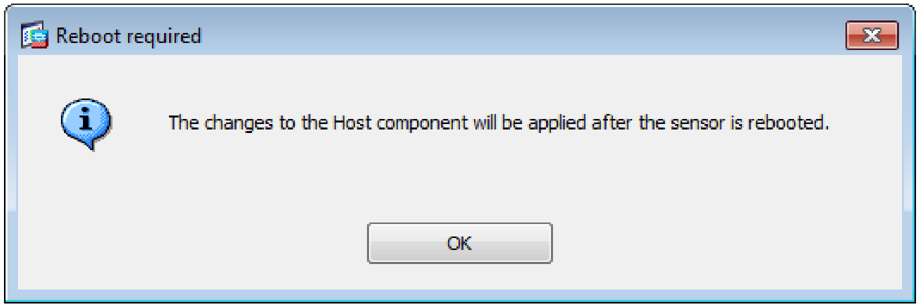

Reboot required

The changes to the Host component will be applied after the sensor is rebooted.

The system notifies you that the IPS sensor requires a reboot to apply the new configuration. Click **OK**, and proceed to the next step, and delay the reboot until the end of this procedure.

Step 12: Navigate to **Policies > IPS Policies**.

On the main panel, note that there is an Event Action Override to Deny Packet Inline for all High Risk events.

Step 13: If you want to see information about what High Risk means, in the main panel, click **Risk Category**.

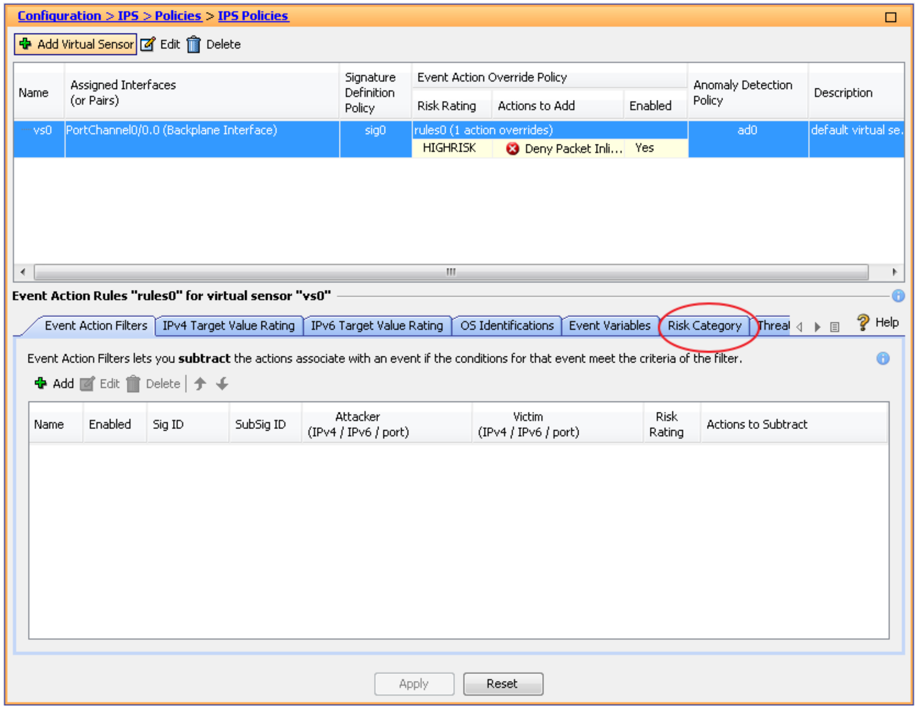

Configuration > IPS > Policies > IPS Policies

Name	Assigned Interfaces (or Pairs)	Signature Definition Policy	Event Action Override Policy			Anomaly Detection Policy	Description
			Risk Rating	Actions to Add	Enabled		
vs0	PortChannel0/0.0 (Backplane Interface)	sig0	rules0 (1 action overrides)			ad0	default virtual se.
			HIGHRISK	Deny Packet Inli...	Yes		

Event Action Rules "rules0" for virtual sensor "vs0"

Event Action Filters lets you **subtract** the actions associate with an event if the conditions for that event meet the criteria of the filter.

Name	Enabled	Sig ID	SubSig ID	Attacker (IPv4 / IPv6 / port)	Victim (IPv4 / IPv6 / port)	Risk Rating	Actions to Subtract

In the default case, High Risk means events that have a Risk Rating from 90 to 100. In this deployment, you reduce the risk of dropping non-malicious traffic by editing the Deny Packet action such that it triggers only when the Risk Rating is 100. This means that the sensor will now use the Deny Packet action only on events with a Risk Rating equal to 100, which only occurs when the most accurate, highest-risk signatures fire.

Step 14: In the Virtual Sensor panel, right-click the **vs0** entry, and then click **Edit**.

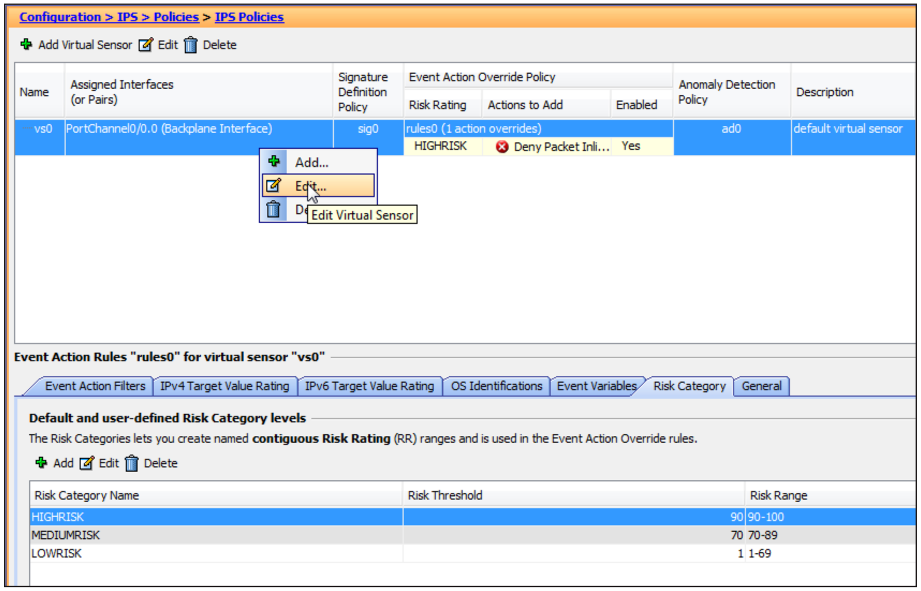

Configuration > IPS > Policies > IPS Policies

Add Virtual Sensor Edit Delete

Name	Assigned Interfaces (or Pairs)	Signature Definition Policy	Event Action Override Policy			Anomaly Detection Policy	Description
			Risk Rating	Actions to Add	Enabled		
vs0	PortChannel0/0.0 (Backplane Interface)	sig0	rules0 (1 action overrides)			ad0	default virtual sensor
			HIGHRISK	Deny Packet Inli...	Yes		

Event Action Rules "rules0" for virtual sensor "vs0"

Default and user-defined Risk Category levels

The Risk Categories lets you create named **contiguous Risk Rating** (RR) ranges and is used in the Event Action Override rules.

Add Edit Delete

Risk Category Name	Risk Threshold	Risk Range
HIGHRISK	90	90-100
MEDIUMRISK	70	70-89
LOWRISK	1	1-69

Step 15: In the Event Action Rule work pane, select **Deny Packet Inline Override**, and then click **Delete**.

Step 16: Click **Add**, highlight and delete the HIGHRISK value in the Risk Rating drop down box and enter a value of **100-100** , select **Deny Packet Inline**, click **OK,** and then click **Apply.**

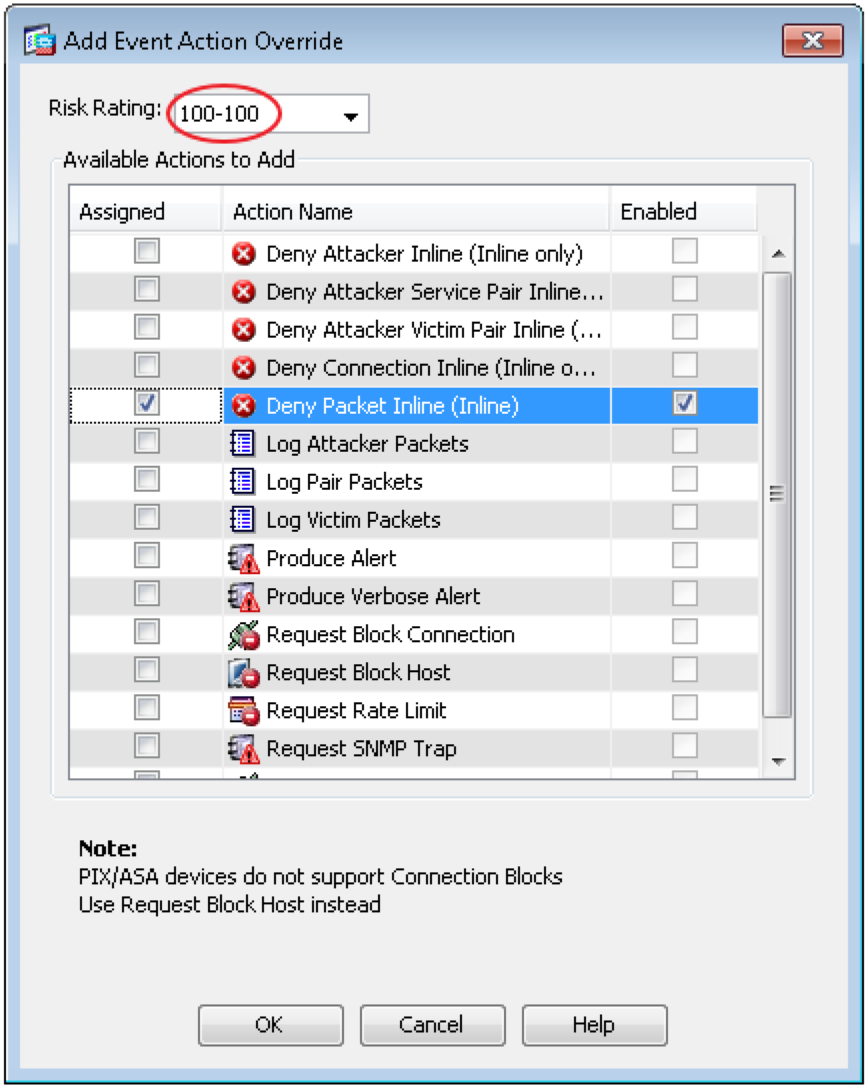

Add Event Action Override

Risk Rating: 100-100

Available Actions to Add

Assigned	Action Name	Enabled
☐	Deny Attacker Inline (Inline only)	☐
☐	Deny Attacker Service Pair Inline...	☐
☐	Deny Attacker Victim Pair Inline (...	☐
☐	Deny Connection Inline (Inline o...	☐
☑	Deny Packet Inline (Inline)	☑
☐	Log Attacker Packets	☐
☐	Log Pair Packets	☐
☐	Log Victim Packets	☐
☐	Produce Alert	☐
☐	Produce Verbose Alert	☐
☐	Request Block Connection	☐
☐	Request Block Host	☐
☐	Request Rate Limit	☐
☐	Request SNMP Trap	☐

Note:
PIX/ASA devices do not support Connection Blocks
Use Request Block Host instead

Step 17: Navigate to **IPS > Reboot Sensor**, click **Reboot Sensor**, and then click **OK** again.

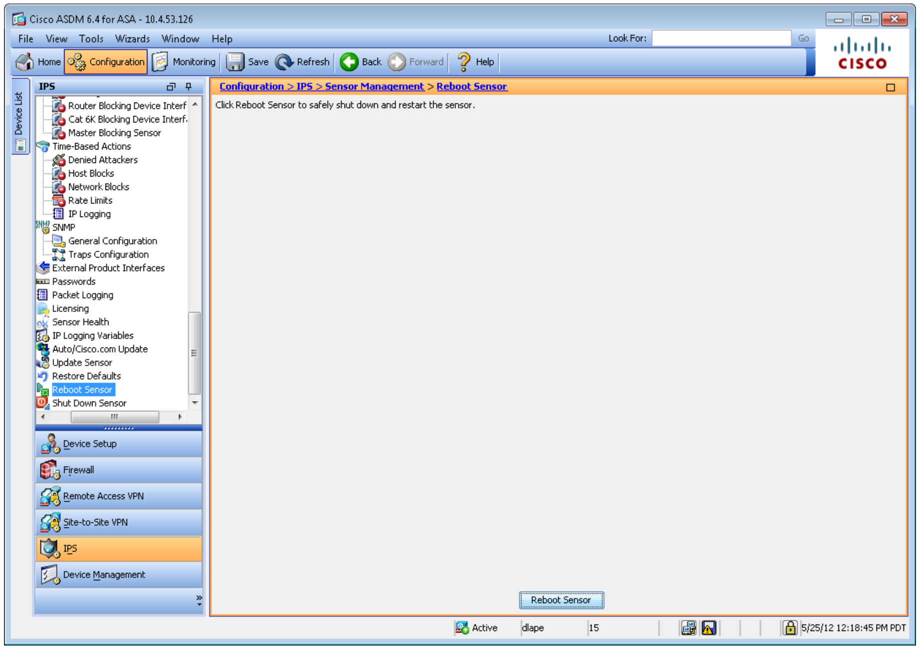

Click Reboot Sensor to safely shut down and restart the sensor.

The GUI console will disconnect from the IPS session and request you log into the ASDM session on the Cisco ASA firewall again. The primary firewall will now switch to standby state because it has lost connectivity to the IPS module in the primary appliance.

Step 18: Log in to the ASDM session on the firewall using the same IP address and credentials you used in Step 1 of this procedure. You are now logging into the secondary active firewall + IPS module pair. Repeat Step 11 through Step 27 for the IPS module installed in the other Cisco ASA chassis, using the **IPS SSP20-B** module name and **10.4.63.23** IP address.

There is no configuration synchronization between the two sensors.

Reader Tip

Cisco IME is a standalone application that can configure and monitor activity for up to 10 sensors (as of IME 7.1.1). Cisco IME is available at no extra cost on Cisco.com in the same web location as Cisco IPS software updates and upgrades.

Procedure 3 Configure signature updates

(Optional)

IDS and IPS devices are generally only as good as their last update, and because of this, it is important that you keep the sensors updated. To this end, the easiest solution is to configure each sensor to retrieve signature updates directly from Cisco.com. Use Cisco ASDM to program the following steps.

Step 1: Click **Configuration**, click **IPS**, navigate to **Configuration> IPS** > **Sensor Management** > **Auto/Cisco.com Update**, select **Enable Signature and Engine Updates from Cisco.com**, and then expand the **Cisco.com Server Settings** panel.

Step 2: Provide a valid cisco.com username and password that holds entitlement to download IPS software updates.

Step 3: Select **Daily**, enter a time between 12:00 a.m. and 4:00 a.m. for the **Start Time**, select each day, and then click **Apply**.

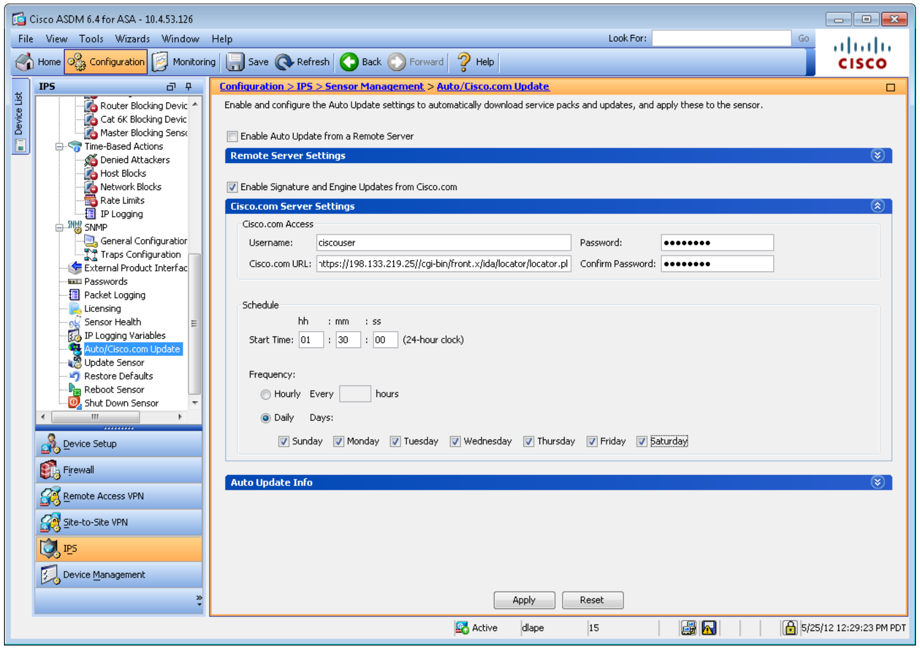

Configuration > IPS > Sensor Management > Auto/Cisco.com Update

Enable and configure the Auto Update settings to automatically download service packs and updates, and apply these to the sensor.

☐ Enable Auto Update from a Remote Server

Remote Server Settings

☑ Enable Signature and Engine Updates from Cisco.com

Cisco.com Server Settings

Cisco.com Access

Username: ciscouser Password: ●●●●●●●●

Cisco.com URL: ıttps://198.133.219.25//cgi-bin/front.x/ida/locator/locator.pl Confirm Password: ●●●●●●●●

Schedule

Start Time: 01 : 30 : 00 (hh : mm : ss) (24-hour clock)

Frequency:

○ Hourly Every hours

● Daily Days:

☑ Sunday ☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday ☑ Friday ☑ Saturday

Auto Update Info

Tech Tip

Using the auto-update feature from Cisco.com will update only the sensor’s engine files and signature files. Major and minor code versions and service packs are not updated with this mechanism.

Notes

Application Resiliency

The network is playing an increasingly important role in the success of a business. Key applications, such as enterprise resource planning, e-commerce, email, and portals, must be available around-the-clock to provide uninterrupted business services. However, the availability of these applications is often threatened by network overloads, as well as server and application failures. Furthermore, resource utilization is often out of balance, resulting in the low-performance resources being overloaded with requests while the high-performance resources remain idle. Application performance, as well as availability, directly affects employee productivity and the bottom-line of a company. As more users work more hours while using key business applications, it becomes even more important to address application availability and performance issues to ensure achievement of business processes and objectives.

Some of the factors that make applications difficult to deploy and deliver effectively over the network include:

- Inflexible application infrastructure—Application infrastructure design has historically been done on an application-by-application basis. This means that the infrastructure used for a particular application is often unique to that application. This type of design tightly couples the application to the infrastructure and offers little flexibility. Because the application and infrastructure are tightly coupled, it is difficult to partition resources and levels of control to match changing business requirements.
- Server availability and load—The mission-critical nature of applications puts a premium on server availability. Despite the benefits of server virtualization technology, the number of physical servers continues to grow based on new application deployments, which in turn increases power and cooling requirements.
- Application security and compliance—Many of the new threats to network security are the result of application- and document-embedded attacks that compromise application performance and availability. Such attacks can also potentially cause the loss of vital application data, while leaving networks and servers unaffected.

One possible solution to improve application performance and availability is to rewrite the application completely to make it network-optimized. However, this requires application developers to have a deep understanding of how different applications respond to bandwidth constraints, delay, jitter, and other network variances. In addition, developers need to accurately predict each end-user's foreseeable access method. This is simply not feasible for every business application, particularly traditional applications that took years to write and customize.

Technology Overview

The idea of improving application performance began in the data center. The Internet boom ushered in the era of the server load balancers (SLBs). SLBs balance the load on groups of servers to improve server response to client requests, and have evolved to take on additional responsibilities, such as application proxies and complete Layer 4 through 7 application switching.

Cisco Application Control Engine (Cisco ACE) is the latest SLB offering from Cisco. Its main role is to provide Layer 4 through 7 switching, but Cisco ACE also provides an array of acceleration and server offload benefits, including TCP-processing offload, SSL-processing offload, and compression. The Cisco ACE appliance sits in the data center in front of the application servers and provides a range of services to maximize server and application availability, security, and asymmetric application acceleration (from server to client browser). As a result, Cisco ACE gives IT departments more control over application and server infrastructure, which enables them to manage and secure application services more easily and improve performance.

Cisco ACE provides the following benefits:

- Scalability—Cisco ACE scales the performance of a server-based program, such as a web server, by distributing its client requests across multiple servers, known as a server farm. As traffic increases, additional servers can be added to the farm. With the advent of server virtualization, application servers can be staged and added dynamically as capacity requirements change.
- High availability—Cisco ACE provides high availability by automatically detecting the failure of a server and repartitioning client traffic among the remaining servers within seconds, while providing users with continuous service.
- Application acceleration—Cisco ACE improves application performance and reduces response time by minimizing latency and compressing data transfers for any HTTP-based application, for any internal or external end-user.
- Server offload—Cisco ACE offloads TCP processing, SSL processing, and compression from the server, which allows the server to handle more requests so more users can be served, and reduces bandwidth requirements by up to 90% without increasing the number of servers. Running SSL on the web application servers is a tremendous drain on server resources. By offloading SSL processing, those resources can be applied to traditional web-application functions. In addition, because persistence information used by the content switches is inside the HTTP header, this information is no longer visible when carried inside SSL sessions. By terminating these sessions before applying content switching decisions, persistence options become available for secure sites.
- Flexible licensing model—Cisco ACE is available in a number of performance options, from 500 Mbps to 4 Gbps of throughput, depending on which license is purchased. You can purchase a 1 Gbps license for your Cisco ACE appliance and then, as your performance requirements increase, upgrade the same hardware to 4 Gbps with a new license.
- Health monitoring—Cisco ACE uses both active and passive techniques to monitor server health. By periodically probing servers and monitoring the return traffic from the real servers, Cisco ACE rapidly detects server failures and quickly reroutes connections to available servers. A variety of health-checking features are supported, including the ability to verify web servers, SSL servers, application servers, databases, FTP servers, and streaming media servers.
- Effective content allocation—Cisco ACE may be used to push requests for cacheable content, such as image files, to a set of caches that can serve them more cost-effectively than the application servers.
- Cisco ACE can be used to partition components of a single web application across several application server clusters. For example, the URLs, www.mycompany.com/quotes/getquote.jsp and www.mycompany.com/trades/order.jsp, could be located on two different server clusters even though the domain name is the same. This partitioning allows the application developer to easily scale the application to several servers without numerous code modifications. Furthermore, it maximizes the cache coherency of the servers by keeping requests for the same pages on the same servers.

There are several ways to integrate Cisco ACE into the data center network. Logically, the Cisco ACE appliance is deployed in front of the application cluster. Requests to the application cluster are directed to a virtual IP address (VIP) configured on the appliance. Cisco ACE receives connections and HTTP-requests, and routes them to the appropriate application server based on configured policies.

Physically, the network topology can take many forms. One-armed mode is the simplest deployment method, in which the Cisco ACE is connected off to the side of the layer 2/layer 3 infrastructure. It is not directly in the path of traffic flow and receives only traffic that is specifically intended for it. Traffic, which should be directed to it, is controlled by careful design of VLANs, virtual server addresses, server default gateway selection, or policy routes on the layer 2/layer 3 switch.

Deployment Details

Cisco ACE 4710 hardware is always deployed in pairs for highest availability, with one primary and one secondary appliance. If the primary Cisco ACE appliance fails, the secondary appliance takes control. Depending on how session-state redundancy is configured, this failover may take place without disrupting the client-to-server connection.

Each Cisco ACE has a port channel that is connected to the switch to scale performance. In this design, the appliance uses two links for 2Gbps of available throughput, but two additional gigabit ports are available. By using four ports, the Cisco ACE appliance can scale the solution to 4Gbps. Cisco ACE operates in an active standby mode, and to maintain performance in a failure scenario, all of the links from each Cisco ACE appliance connect to only a single switch. This prevents the scenario in which Cisco ACE is connected to both switches, and a switch failure cuts the available bandwidth in half.

Process

Configuring Connectivity to the Data Center Core Switches

1. Configure port channels on core switches

Procedure 1 Configure port channels on core switches

The Cisco ACE server load balancers serving applications and servers in the data center will each connect to one of the data center core Cisco Nexus 5500UP switches by using EtherChannel links.

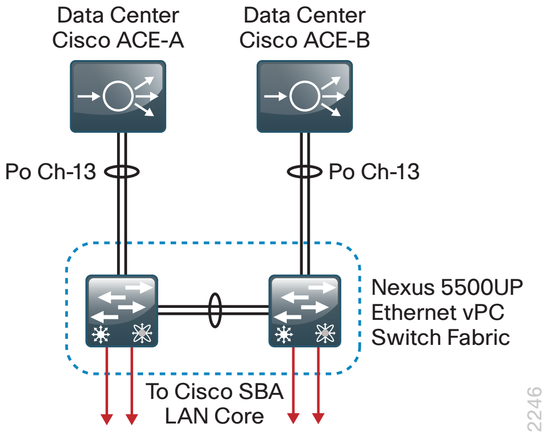

The use of EtherChannel links for connectivity to the core provides a resilient connection, load balances traffic over the links, and makes it easier to add bandwidth in the future.

The data center core Cisco Nexus 5500UP switches use Virtual Port Channel (vPC) for many dual-homed EtherChannel devices. If the vPC peer link between the data center core switches fails, one of the switches will go into error recovery and shut down interfaces associated with VLANs that are part of vPC connections to prevent any loops in the infrastructure. Because the Cisco ACE s are single-homed to each data center core switch and do not use a vPC for connectivity—but instead are using a VLAN that is part of other vPC connections—they are non-vPC ports, also called *vPC orphan ports*. Use the **vpc orphan-port suspend** command to shut down the EtherChannel interfaces to the attached Cisco ACE on each switch in the event that the vPC peer link is broken between the data center core switches and a switch goes into error recovery mode. The active Cisco ACE on the switch that remains in service will continue operating and provides the resiliency in the design.

Cisco ACE does support EtherChannel but does not support Link Aggregation Control Protocol (LACP). Therefore, the **channel-group mode** will be forced on.

Step 1: Configure physical interfaces to the port channels on the first Cisco Nexus 5500UP data center core switch. Use the **speed 1000** command to set the ports connected to Cisco ACE from the default of 10-Gigabit Ethernet to 1-Gigabit Ethernet.

Tech Tip

When configuring the interfaces, you must enter the **vpc orphan-port suspend** command before the **channel-group** command. If you enter the **channel-group** command on the interface first, the switch will not let you enter the **vpc orphan-port suspend** command on the interface.

You must enter the **vpc orphan-port suspend** on all physical interface members of this port-channel to ensure consistent and proper operation.

```
interface Ethernet1/3
  description ACE 1 Gig 1/1
  speed 1000
  vpc orphan-port suspend
  channel-group 13 mode on
!
interface Ethernet1/4
  description ACE 1 Gig 1/2
  speed 1000
  vpc orphan-port suspend
  channel-group 13 mode on
```

When you assign the channel group to a physical interface, it creates the logical EtherChannel (port-channel) interface. In the next step, configure the logical port-channel interfaces on both data center core switches. The physical interfaces tied to the port channel will inherit the settings.

Step 2: Configure the logical port-channel interface. Assign the QoS policy created in Procedure 3, "Configure QoS policies," to the port channel interface.

```
interface port-channel13
  switchport mode trunk
  switchport trunk allowed vlan 149,912
  spanning-tree port type edge trunk
  service-policy type qos input DC-FCOE+1P4Q_INTERFACE-DSCP-
QOS
```

Step 3: Configure an unused VLAN for the Cisco ACE fault-tolerant heartbeat VLAN.

```
vlan 912
name ACE-Heartbeat
```

Step 4: Apply the following configuration to the second Cisco Nexus 5500UP data center core switch.

```
interface Ethernet1/3
  description ACE 2 Gig 1/1
  speed 1000
  vpc orphan-port suspend
  channel-group 13 mode on
!
interface Ethernet1/4
  description ACE 2 Gig 1/2
  speed 1000
  vpc orphan-port suspend
  channel-group 13 mode on
!
interface port-channel13
  switchport mode trunk
  switchport trunk allowed vlan 149,912
  spanning-tree port type edge trunk  service-policy type qos
input DC-FCOE+1P4Q_INTERFACE-DSCP-QOS

!
vlan 912
name ACE-Heartbeat
```

Process

Configuring the Cisco ACE Network

1. Perform initial Cisco ACE setup
2. Configure high availability

Procedure 1 Perform initial Cisco ACE setup

Step 1: Connect to Cisco ACE via the console, perform the initial configuration, and then exit from the initial configuration dialog box at the prompt.

```
switch login: admin
Password: admin
Admin user is allowed to log in only from console until the default password is changed.
www user is allowed to log in only after the default password is changed.
Enter the new password for user "admin": password
Confirm the new password for user "admin": password
admin user password successfully changed.
Enter the new password for user "www": password
Confirm the new password for user "www": password
www user password successfully changed.
<text wall removed>
ACE>Would you like to enter the basic configuration dialog (yes/no) [y]: n
switch/Admin#
```

Step 2: In configuration mode, set the system host name.

```
hostname ACE4710-A
```

Step 3: Set up the basic network security policies. This allows for management access into Cisco ACE.

```
access-list ALL line 8 extended permit ip any any
class-map type management match-any remote_access
  2 match protocol xml-https any
  3 match protocol icmp any
  4 match protocol telnet any
  5 match protocol ssh any
  6 match protocol http any
  7 match protocol https any
  8 match protocol snmp any
policy-map type management first-match remote_mgmt_allow_policy
  class remote_access
    permit
```

Step 4: Configure port channel and trunking on the Gigabit Ethernet interfaces.

```
interface gigabitEthernet 1/1
  channel-group 1
  no shutdown
interface gigabitEthernet 1/2
  channel-group 1
  no shutdown
interface port-channel 1
  switchport trunk native vlan 1
  switchport trunk allowed vlan 149
  no shutdown
```

This configuration provisions a 2-Gbps port channel and is sufficient for Cisco ACE 4710 with up to a 2-Gbps license. If a 4-Gbps license is being used, include Gigabit Ethernet ports 1/3 and 1/4 for a total of 4 Gbps of throughput.

Step 5: Configure the VLAN 149 interface on the Cisco ACE for management access and general network connectivity.

```
interface vlan 149
  ip address 10.4.49.119 255.255.255.0
  access-group input ALL
  service-policy input remote_mgmt_allow_policy
  no shutdown
```

Step 6: Configure the default route.

```
ip route 0.0.0.0 0.0.0.0 10.4.49.1
```

Step 7: Configure NTP.

```
ntp server 10.4.48.17
```

Step 8: Configure SNMP.

```
snmp-server community cisco ro
```

The Cisco ACE appliance should now be reachable via the network. Repeat Step 1 to Step 8 on the second Cisco ACE, replacing the IP address in Step 5 with 10.4.49.120.

Procedure 2 Configure high availability

Next, you configure the Cisco ACE appliances as an active/standby failover pair. After you configure high availability, the devices will be synchronized and further configuration is only necessary on the primary Cisco ACE appliance. Start with the Cisco ACE appliance that you want to be primary. In this example, the primary is 10.4.49.119.

Step 1: Open a browser window and enter https://10.4.49.119 into the address field. The Cisco ACE GUI opens.

Step 2: In the **Username** box, type admin, in the **Password** box, type the password that you configured in Procedure 1, Step 1, and then click **Log In**.

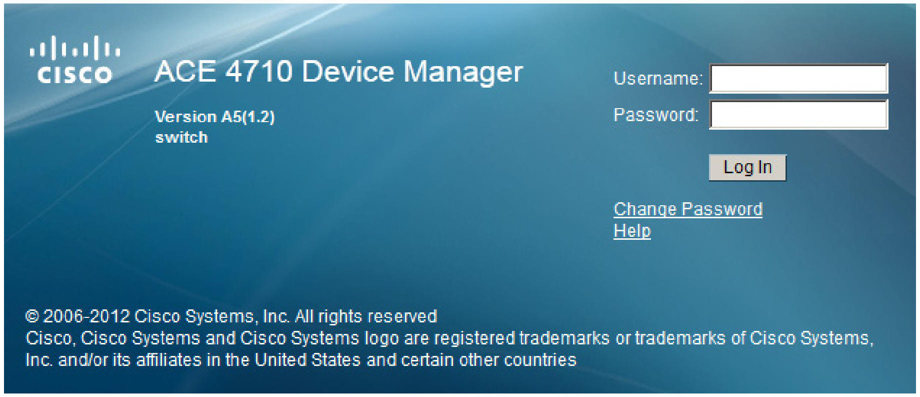

© 2006-2012 Cisco Systems, Inc. All rights reserved
Cisco, Cisco Systems and Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries

Step 3: Navigate to **Config > Virtual Contexts > High Availability (HA) > Setup,** and then click **Edit**.

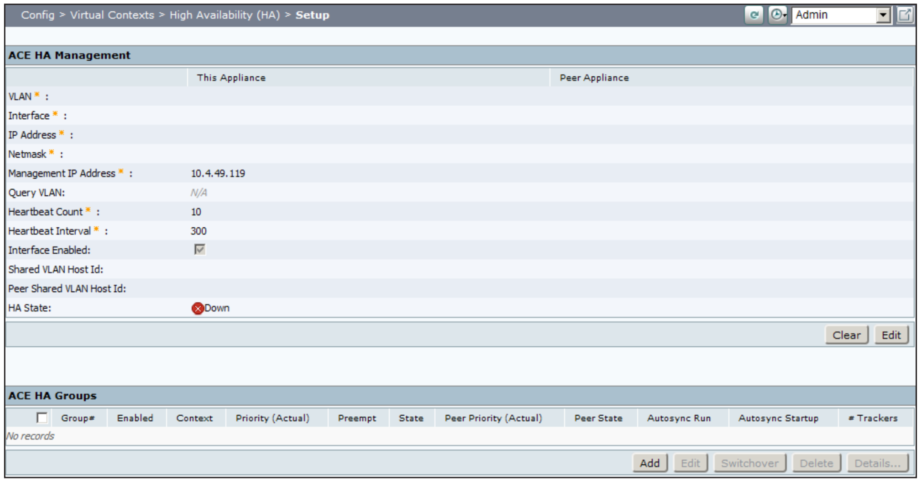

ACE HA Management

	This Appliance	Peer Appliance
VLAN * :		
Interface * :		
IP Address * :		
Netmask * :		
Management IP Address * :	10.4.49.119	
Query VLAN:	N/A	
Heartbeat Count * :	10	
Heartbeat Interval * :	300	
Interface Enabled:	☑	
Shared VLAN Host Id:		
Peer Shared VLAN Host Id:		
HA State:	Down	

ACE HA Groups

□	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
No records											

Step 4: In the ACE HA Management dialog box, enter the following values, and then click **Deploy Now**.

- VLAN—912
- Interface—Port Channel 1
- IP Address—10.255.255.1
- IP Address Peer Appliance—10.255.255.2
- Netmask—255.255.255.0
- Management IP Address—10.4.49.119
- Management IP Address Peer Appliance—10.4.49.120

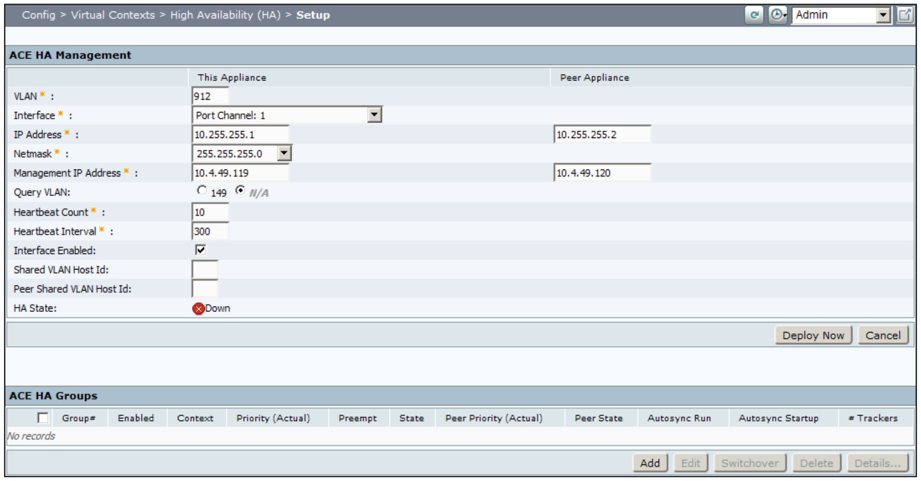
Config > Virtual Contexts > High Availability (HA) > **Setup**
Admin
Deploy Now
Cancel
Add
Edit
Switchover
Delete
Details...

ACE HA Management

	This Appliance	Peer Appliance
VLAN * :	912	
Interface * :	Port Channel: 1	
IP Address * :	10.255.255.1	10.255.255.2
Netmask * :	255.255.255.0	
Management IP Address * :	10.4.49.119	10.4.49.120
Query VLAN:	○ 149 ● N/A	
Heartbeat Count * :	10	
Heartbeat Interval * :	300	
Interface Enabled:	☑	
Shared VLAN Host Id:		
Peer Shared VLAN Host Id:		
HA State:	⊗Down	

ACE HA Groups

□	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
No records											

Step 5: In the ACE HA Groups dialog box, click **Add**.

Step 6: Leave all of the values at their defaults, and then click **Deploy Now**.

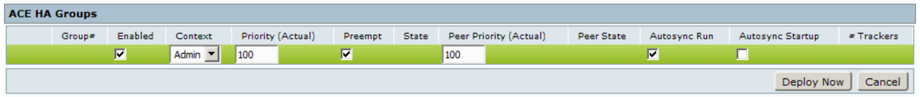

ACE HA Groups

	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
		☑	Admin	100	☑		100		☑	□	

High availability is now configured on the primary Cisco ACE appliance. To configure high availability on the secondary appliance, you must log in to the secondary Cisco ACE appliance.

Step 7: Open a browser window and enter https://10.4.49.120 into the address field. The Cisco ACE GUI opens.

Step 8: In the **Username** box, type admin, in the **Password** box, type the password that you configured in Procedure 1, Step 1, and then click **Log In**.

Step 9: Navigate to **Config > Virtual Contexts > High Availability (HA) > Setup**, and then click **Edit**.

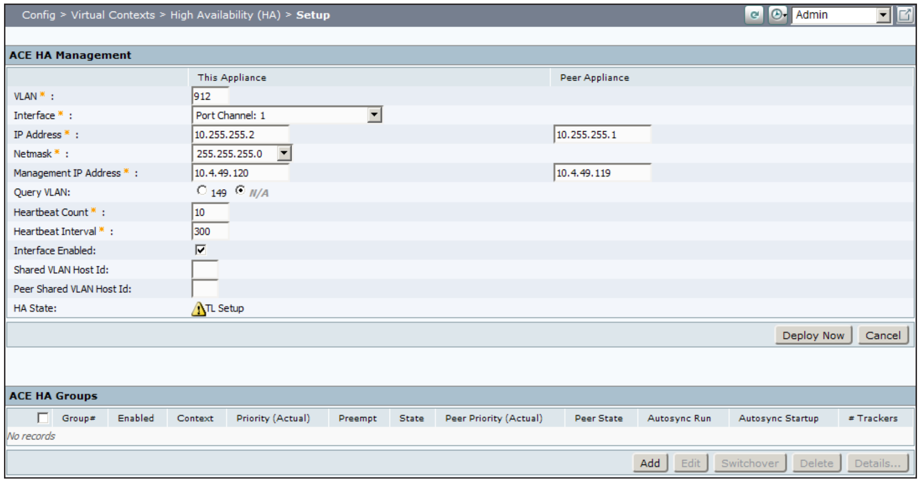
Config > Virtual Contexts > High Availability (HA) > **Setup**
Admin
Deploy Now
Cancel
Add
Edit
Switchover
Delete
Details...

ACE HA Management

	This Appliance	Peer Appliance
VLAN * :	912	
Interface * :	Port Channel: 1	
IP Address * :	10.255.255.2	10.255.255.1
Netmask * :	255.255.255.0	
Management IP Address * :	10.4.49.120	10.4.49.119
Query VLAN:	○ 149 ● N/A	
Heartbeat Count * :	10	
Heartbeat Interval * :	300	
Interface Enabled:	☑	
Shared VLAN Host Id:		
Peer Shared VLAN Host Id:		
HA State:	⚠TL Setup	

ACE HA Groups

□	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
No records											

Step 10: In the ACE HA Management dialog box, enter the following values, and then click **Deploy Now**.

- VLAN—912
- Interface—Port Channel 1
- IP Address—10.255.255.2
- IP Address Peer Appliance—10.255.255.1
- Netmask—255.255.255.0
- Management IP Address—10.4.49.120
- Management IP Address Peer Appliance—10.4.49.119

Step 11: In the ACE HA Groups dialog box, click **Add**.

Step 12: Leave all of the values at their defaults, and then click **Deploy Now**.

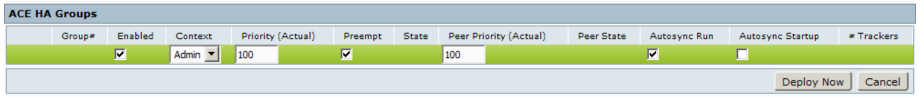

ACE HA Groups

	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
		☑	Admin	100	☑		100		☑	☐	

The two Cisco ACE appliances should be communicating and high availability should be up and active. The device you just finished configuring should show a state of "Standby Hot" and the peer should be "Active," as shown in the ACE HA Groups dialog box below.

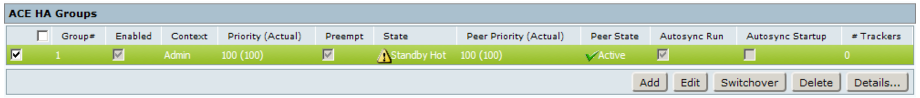

ACE HA Groups

☐	Group#	Enabled	Context	Priority (Actual)	Preempt	State	Peer Priority (Actual)	Peer State	Autosync Run	Autosync Startup	# Trackers
☑	1	☑	Admin	100 (100)	☑	⚠ Standby Hot	100 (100)	✓ Active	☑	☐	0

Make any additional configurations on the primary Cisco ACE appliance, 10.4.49.119. All changes are automatically replicated to the secondary Cisco ACE appliance, 10.4.49.120.

Process

Setting Up Load Balancing for HTTP Servers

1. Configure health probes
2. Configure real servers
3. Configure a server farm
4. Configure Inband-Health checking
5. Configure a NAT pool
6. Configure a virtual server

Procedure 1 Configure health probes

Health probes poll the servers or applications to make sure that the server or service is available and to allow the system to remove failed devices. For this configuration, you will build an Internet Control Message Protocol (ICMP) and an HTTP probe.

Step 1: Open a browser window and enter https://10.4.49.119 into the address field. The Cisco ACE GUI opens.

Step 2: In the **Username** box, type admin, in the **Password** box, type the password you configured in Procedure 1,Step 1, and then click **Log In**.

Step 3: Navigate to **Config > Virtual Contexts > Load Balancing > Health Monitoring**, and then click **Add**.

Step 4: In the New Health Monitoring dialog box, in the Name box, enter icmp-probe, and then, in the Type list, choose **ICMP**.

Step 5: Click **Deploy Now**.

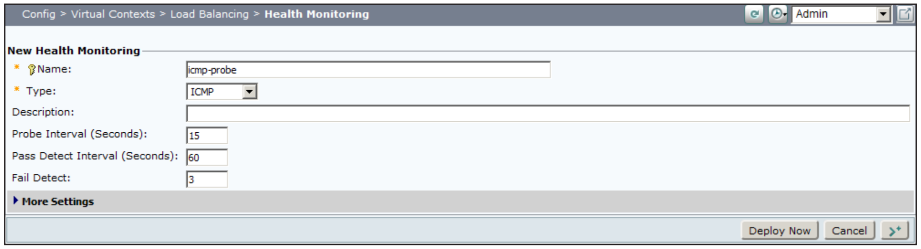

Config > Virtual Contexts > Load Balancing > Health Monitoring
Admin

New Health Monitoring

- * Name: icmp-probe
- * Type: ICMP
- Description:
- Probe Interval (Seconds): 15
- Pass Detect Interval (Seconds): 60
- Fail Detect: 3

More Settings
Deploy Now
Cancel

Step 6: Navigate to **Config > Virtual Contexts > Load Balancing > Health Monitoring**, and then click **Add**.

Step 7: In the New Health Monitoring dialog box, in the **Name** box, enter http-probe, and then, in the Type list, choose **HTTP**.

Step 8: Click **Deploy Now**.

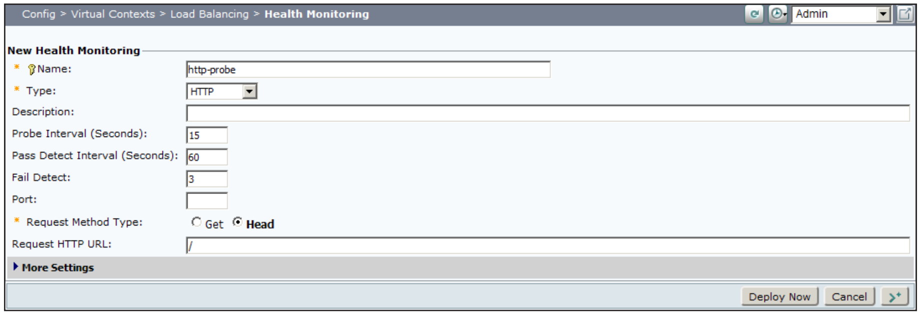

Config > Virtual Contexts > Load Balancing > Health Monitoring
Admin

New Health Monitoring

- * Name: http-probe
- * Type: HTTP
- Description:
- Probe Interval (Seconds): 15
- Pass Detect Interval (Seconds): 60
- Fail Detect: 3
- Port:
- * Request Method Type: ○ Get ● Head
- Request HTTP URL: /

More Settings
Deploy Now
Cancel

Step 9: Click the **Expect Status** tab, and then click **Add**.

Step 10: For both the maximum and minimum status codes, enter 200, and then click **Deploy Now**.

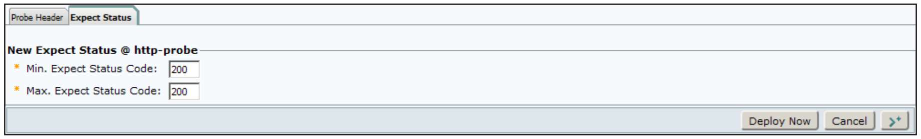

New Expect Status @ http-probe

* Min. Expect Status Code: 200

* Max. Expect Status Code: 200

You have now created the ICMP and HTTP probes, which will be used to monitor the real and virtual servers in the load balancing server farm.

Procedure 2 Configure real servers

In this procedure, you add the real servers across which Cisco ACE load balances client connections.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Real Servers**, and then click **Add**.

Step 2: In the New Real Server dialog box, enter the values below, and then click **Deploy Now**.

- Name—webserver1
- IP Address—10.4.49.111
- Probes—icmp-probe

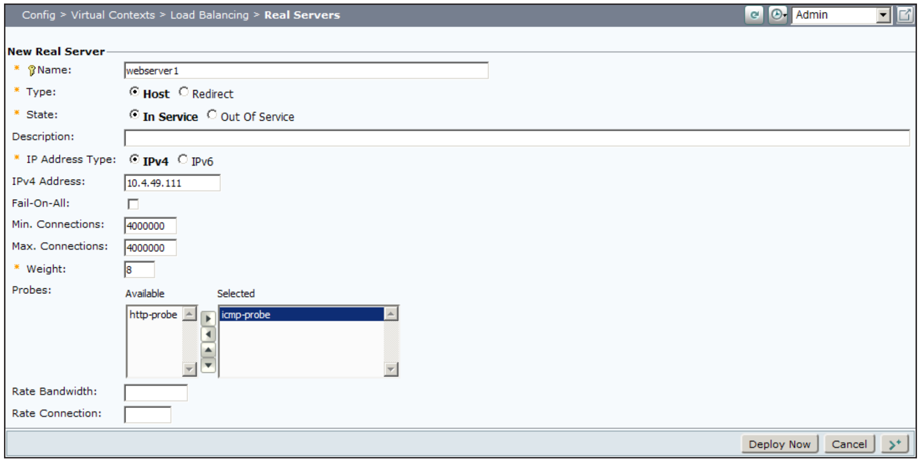

New Real Server

* Name: webserver1

* Type: ● Host ○ Redirect

* State: ● In Service ○ Out Of Service

Description:

* IP Address Type: ● IPv4 ○ IPv6

IPv4 Address: 10.4.49.111

Fail-On-All: □

Min. Connections: 4000000

Max. Connections: 4000000

* Weight: 8

Probes: Available: http-probe Selected: icmp-probe

Rate Bandwidth:

Rate Connection:

Step 3: Navigate to **Config > Virtual Contexts > Load Balancing > Real Servers**, and then click **Add**.

Step 4: In the New Real Server dialog box, enter the values below, and then click **Deploy Now**.

- Name—webserver2
- IP Address—10.4.49.112
- Probes—icmp-probe

This example uses the ICMP probe to monitor the real servers configured in this example, thereby ensuring the server is monitored rather than a specific service. This is the most flexible configuration and allows load balancing for multiple services on a single physical or virtual server.

You have just configured the two web servers. If you have additional servers that you plan on using, you can configure them now by repeating Procedure 2.

Procedure 3 Configure a server farm

A server farm on Cisco ACE is a pool of real servers that you can use to connect to the virtual IP address that the clients will use to connect to the HTTP service.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Server Farms**, and then click **Add**.

Step 2: In the New Server Farm dialog box, enter the values below, and then click **Deploy Now**.

- Name—**webfarm**
- Probes—**http-probe**

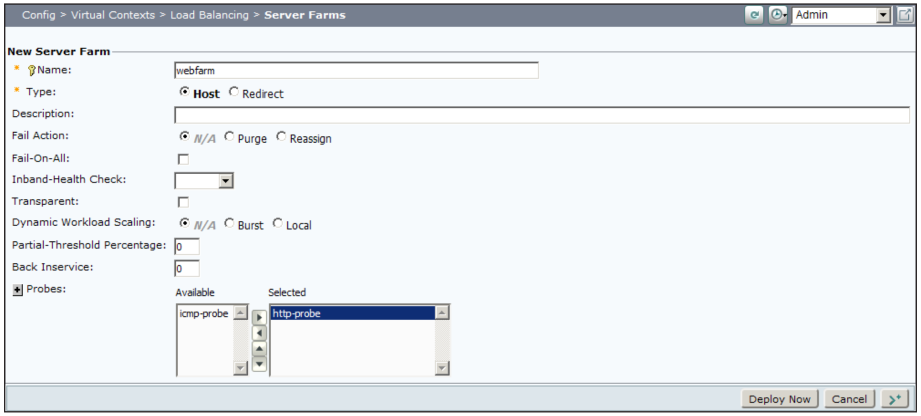
Config > Virtual Contexts > Load Balancing > **Server Farms**
Admin
Deploy Now
Cancel

New Server Farm

* Name: webfarm
* Type: ● **Host** ○ Redirect
Description:
Fail Action: ● *N/A* ○ Purge ○ Reassign
Fail-On-All: □
Inband-Health Check:
Transparent: □
Dynamic Workload Scaling: ● *N/A* ○ Burst ○ Local
Partial-Threshold Percentage: 0
Back Inservice: 0
Probes: Available: icmp-probe — Selected: http-probe

Step 3: Click the **Real Server** tab, and then click **Add**.

Step 4: In the New Real Server dialog box, next to Name, select **webserver1**, and in the Port box, enter **80** for HTTP.

Step 5: Click **Deploy Now**.

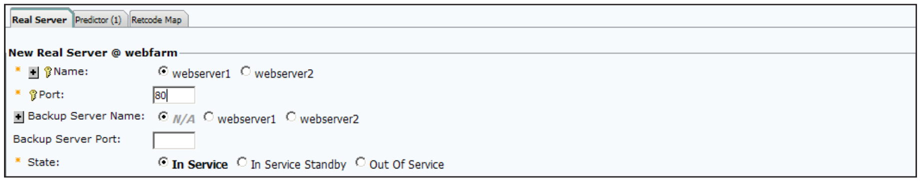

New Real Server @ webfarm

* Name: ● webserver1 ○ webserver2
* Port: 80
Backup Server Name: ● *N/A* ○ webserver1 ○ webserver2
Backup Server Port:
* State: ● **In Service** ○ In Service Standby ○ Out Of Service

Step 6: Click the **Real Server** tab, and then click **Add**.

Step 7: In the New Real Server dialog box, next to Name, select **webserver2**, and in the Port box, enter **80.**

Step 8: Click **Deploy Now**.

Step 9: On the Edit Server Farm dialog box, click **Deploy Now**.

You have just created the server-farm, webfarm, with the real-server members, webserver1 and webserver2, for HTTP on port 80. The http-probe will monitor all of the servers in the server farm to ensure that the HTTP service is available.

Procedure 4 Configure Inband-Health checking

Inband-health checking on Cisco ACE monitors return traffic and looks for failures from the real servers to the clients. It can identify, faster than active probes, when a server is having issues. When a failure is detected, the following modes are available:

- **Count**—Logs the failures locally on Cisco ACE, allowing you to view server issues from the CLI.
- **Log**—Triggers a syslog message to be sent to a Network Management System (NMS), as well as keeping the log locally on Cisco ACE.
- **Remove**—Triggers a log and takes the server out of service.

In this procedure, Log mode is used. This is because a small amount of errors of this type are normal on servers. Without more information about the server farm, using Remove mode could mean that the threshold would be too low and would take a system out of service unnecessarily, or too high and not take a failing server out of service. Log mode allows you to see errors and identify which real sever is having problems.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Server Farms,** select **webfarm**, and then click **View/Edit**.

Step 2: In the Edit Server Farm dialog box, enter the values below, and then click **Deploy Now**.

- Inband-Health Check—Log
- Connection Failure Threshold Count—5
- Reset Timeout (Milliseconds)—500

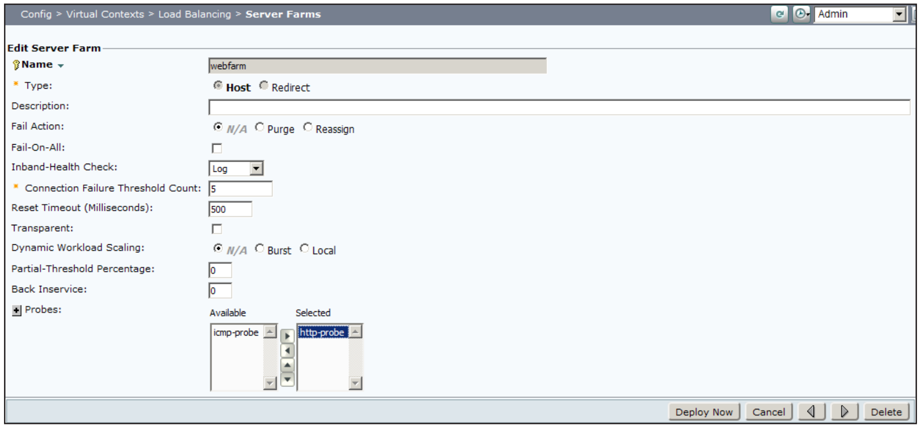

Edit Server Farm

Name: webfarm
* Type: ● **Host** ○ Redirect
Description:
Fail Action: ● N/A ○ Purge ○ Reassign
Fail-On-All: □
Inband-Health Check: Log
* Connection Failure Threshold Count: 5
Reset Timeout (Milliseconds): 500
Transparent: □
Dynamic Workload Scaling: ● N/A ○ Burst ○ Local
Partial-Threshold Percentage: 0
Back Inservice: 0
Probes: Available: icmp-probe Selected: http-probe

Servers in the webfarm are now being monitored for TCP errors. If five errors occur within a 500-ms period, a syslog message will be sent to the NMS. If there is not a syslog server available on the network, the inband-health check can be set to use Count mode, and local statistics will be maintained on Cisco ACE and can be checked from the CLI.

Step 3: At the bottom of the Server Farm dialog box, click the **Retcode Map** tab, and then click **Add**.

Step 4: In the New Retcode Map dialog box, enter the values below, and then click **Deploy Now**.

- Lowest Retcode—404
- Highest Retcode—404
- Type—Log
- Threshold—5
- Reset—10

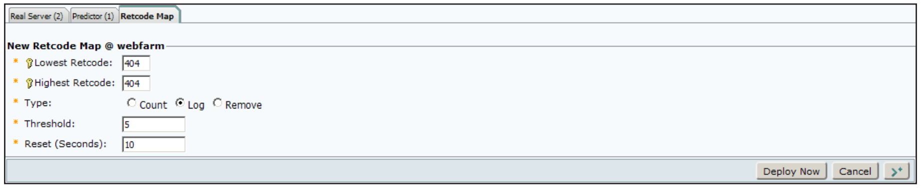

New Retcode Map @ webfarm

* Lowest Retcode: 404
* Highest Retcode: 404
* Type: ○ Count ● Log ○ Remove
* Threshold: 5
* Reset (Seconds): 10

If a server in the webfarm responds to a client with the HTTP return code 404 five times in 10 seconds, a syslog message will be sent to the NMS.

Step 5: At the bottom of the Server Farm dialog box, click the **Retcode Map** tab, and then click **Add**.

Step 6: In the New Retcode Map dialog box, enter the values below, and then click **Deploy Now**.

- Lowest Retcode—500
- Highest Retcode—505
- Type—Log
- Threshold—5
- Reset—10

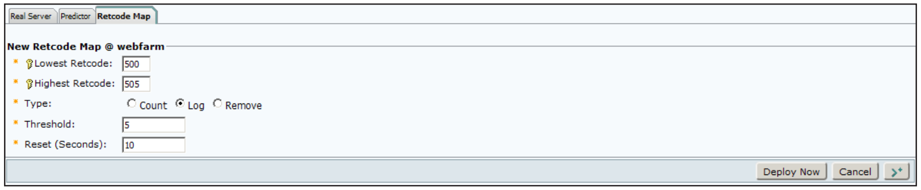

New Retcode Map @ webfarm

* Lowest Retcode: 500
* Highest Retcode: 505
* Type: ○ Count ● Log ○ Remove
* Threshold: 5
* Reset (Seconds): 10

If, within a 10-second period, a server in the webfarm responds to a client five times with the HTTP return code in the range of 500 to 505, a syslog message will be sent to the NMS.

Step 7: Navigate to **Config > Virtual Contexts > System > Syslog**, and then select **Enable Syslog**.

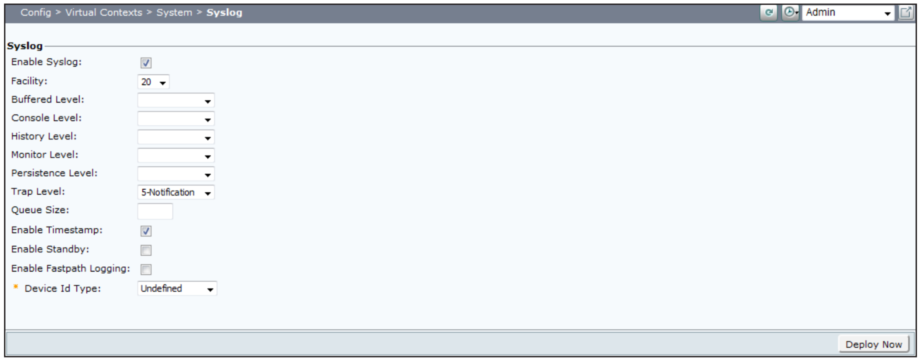

Syslog

Enable Syslog: ☑
Facility: 20
Buffered Level:
Console Level:
History Level:
Monitor Level:
Persistence Level:
Trap Level: 5-Notification
Queue Size:
Enable Timestamp: ☑
Enable Standby: □
Enable Fastpath Logging: □
* Device Id Type: Undefined

Step 8: On the Log Host tab, click **Add**, enter 10.4.48.35, and then click **Deploy Now**.

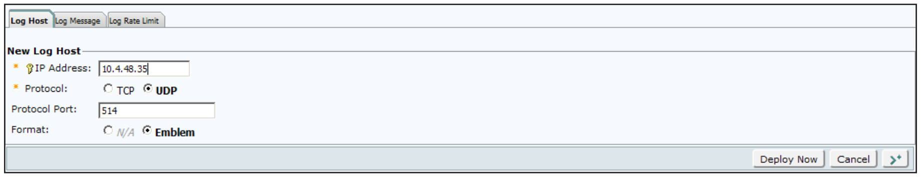

New Log Host

* IP Address: 10.4.48.35
* Protocol: ○ TCP ● **UDP**
Protocol Port: 514
Format: ○ N/A ● **Emblem**

Step 9: In the Syslog dialog box, click **Deploy Now**.

Now the syslog messages that are triggered by the inband-health checks are sent to the syslog server at 10.4.48.35.

Procedure 5 Configure a NAT pool

Step 1: Navigate to **Config > Virtual Contexts > Network > NAT Pools**, and then click **Add**.

Step 2: In the New NAT Pool dialog box, enter the following values, and then click **Deploy Now**.

- Start IP Address—10.4.49.99
- End IP Address—10.4.49.99
- Netmask—255.255.255.0

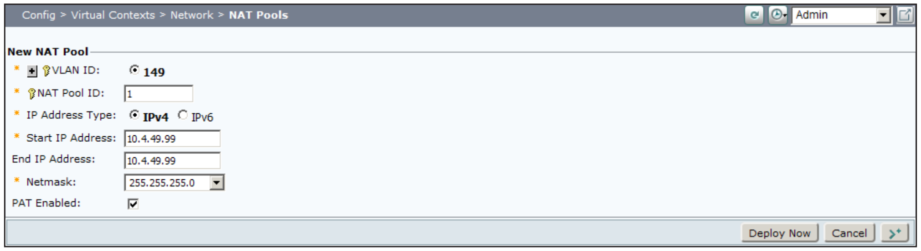

New NAT Pool

* VLAN ID: ● **149**
* NAT Pool ID: 1
* IP Address Type: ● **IPv4** ○ IPv6
* Start IP Address: 10.4.49.99
End IP Address: 10.4.49.99
* Netmask: 255.255.255.0
PAT Enabled: ☑

Procedure 6 Configure a virtual server

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Virtual Servers**, and then click **Add**.

Step 2: In the Properties dialog box, enter the following values:

- Virtual Server Name—http-vip
- Virtual IP Address—10.4.49.100
- VLAN—149

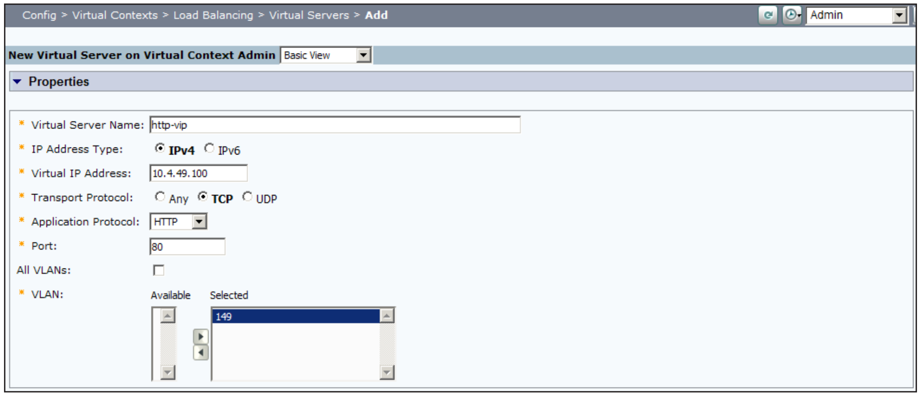
Config > Virtual Contexts > Load Balancing > Virtual Servers > **Add**
Admin

New Virtual Server on Virtual Context Admin Basic View

Properties

* Virtual Server Name: http-vip
* IP Address Type: ● **IPv4** ○ IPv6
* Virtual IP Address: 10.4.49.100
* Transport Protocol: ○ Any ● **TCP** ○ UDP
* Application Protocol: HTTP
* Port: 80
All VLANs: □
* VLAN: Available Selected 149

Step 3: In the Default L7 Load-Balancing Action dialog box, in the **Server Farm** list, choose **webfarm**, and then select **Deflate**.

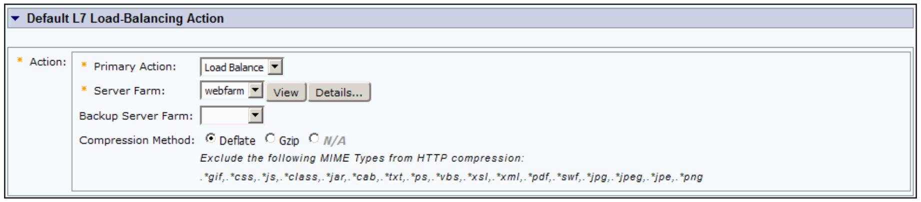

Default L7 Load-Balancing Action

* Action:

* Primary Action: Load Balance

* Server Farm: webfarm

Backup Server Farm:

Compression Method: ● Deflate ○ Gzip ○ N/A

Exclude the following MIME Types from HTTP compression:

*.*gif,.*css,.*js,.*class,.*jar,.*cab,.*txt,.*ps,.*vbs,.*xsl,.*xml,.*pdf,.*swf,.*jpg,.*jpeg,.*jpe,.*png*

Step 4: In the NAT dialog box, click **Add**, click **OK**, and then click **Deploy Now**.

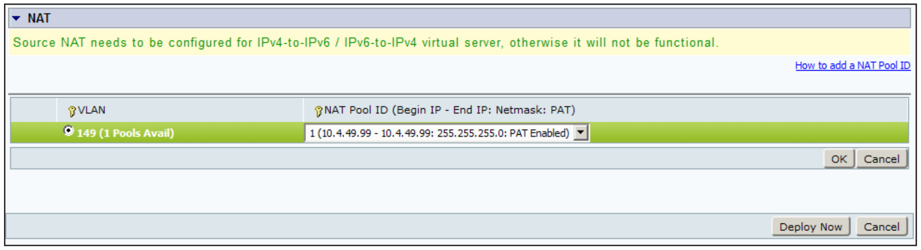

NAT

Source NAT needs to be configured for IPv4-to-IPv6 / IPv6-to-IPv4 virtual server, otherwise it will not be functional.

	VLAN	NAT Pool ID (Begin IP - End IP: Netmask: PAT)
●	149 (1 Pools Avail)	1 (10.4.49.99 - 10.4.49.99: 255.255.255.0: PAT Enabled)

Clients going to the virtual IP 10.4.49.100 on port 80 will be load balanced across the real servers webserver1 and webserver2 in the server farm webfarm.

Process

Load Balancing and SSL Offloading for HTTPS Servers

1. Configure real servers
2. Configure a server farm
3. Configure SSL proxy service
4. Configure HTTP-cookie sticky service
5. Configure a virtual server
6. Configure an HTTP-to-HTTPS Redirect

You can configure a group of servers for load balancing, in which the Cisco ACE appliance performs all of the SSL processing, thereby offloading it from the servers.

Procedure 1 Configure real servers

In this procedure, you add the real servers across which Cisco ACE load balances client SSL connections.

Step 1: Open a browser window and enter https://10.4.49.119 into the address field. The Cisco ACE GUI opens.

Step 2: In the **Username** box, type admin, in the **Password** box, type the password you configured in Procedure 1,,Step 1, and then click **Log In**.

Step 3: Navigate to **Config > Virtual Contexts > Load Balancing > Real Servers**, and then click **Add**.

Step 4: In the New Real Server dialog box, enter the values below, and then click **Deploy Now**.

- Name—**webserver3**
- IP Address—**10.4.49.113**
- Probes—**icmp-probe**

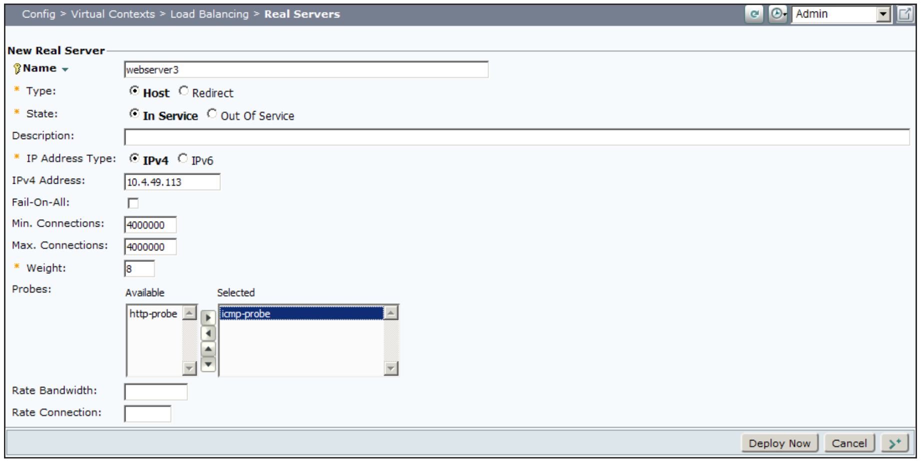

Config > Virtual Contexts > Load Balancing > **Real Servers**
Admin
Deploy Now
Cancel

New Real Server

Name	webserver3
* Type:	● **Host** ○ Redirect
* State:	● **In Service** ○ Out Of Service
Description:	
* IP Address Type:	● **IPv4** ○ IPv6
IPv4 Address:	10.4.49.113
Fail-On-All:	□
Min. Connections:	4000000
Max. Connections:	4000000
* Weight:	8
Probes:	Available: http-probe Selected: icmp-probe
Rate Bandwidth:	
Rate Connection:	

Step 5: Navigate to **Config > Virtual Contexts > Load Balancing > Real Servers**, and then click **Add**.

Step 6: In the New Real Server dialog box, enter the values below, and then click **Deploy Now**.

- Name—**webserver4**
- IP Address—**10.4.49.114**
- Probes—**icmp-probe**

In this example, the ICMP-probe monitors the real servers, thereby ensuring that the server is monitored, rather than a specific service. This is the most flexible configuration and allows load-balancing for multiple services on a single physical or virtual server.

You have just configured the two web servers. If you have additional servers that you plan on using, you can configure them now by repeating Procedure 1.

Procedure 2 Configure a server farm

A server farm is a pool of real servers that you can use to connect to the VIP-address that the clients will use to connect to the HTTP service.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Server Farms**, and then click **Add**.

Step 2: In the New Server Farm dialog box, enter the values below, and then click **Deploy Now**.

- Name—**appfarm**
- Probes—**http-probe**

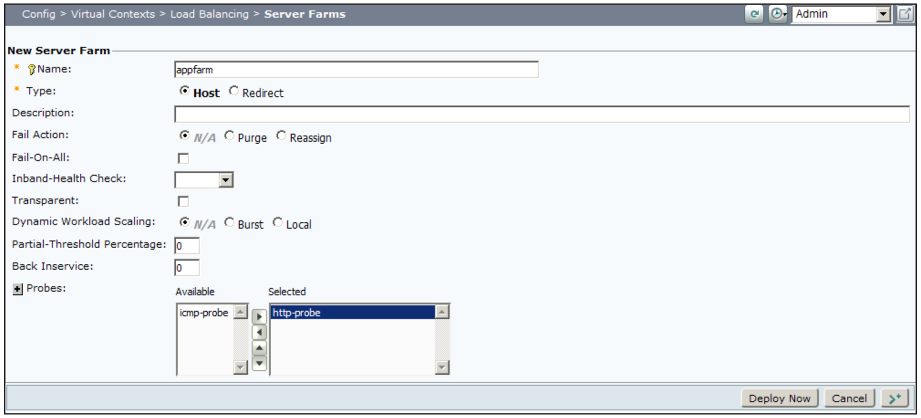

Config > Virtual Contexts > Load Balancing > **Server Farms**
Admin
Deploy Now
Cancel

New Server Farm

* Name:	appfarm
* Type:	● **Host** ○ Redirect
Description:	
Fail Action:	● *N/A* ○ Purge ○ Reassign
Fail-On-All:	□
Inband-Health Check:	
Transparent:	□
Dynamic Workload Scaling:	● *N/A* ○ Burst ○ Local
Partial-Threshold Percentage:	0
Back Inservice:	0
Probes:	Available: icmp-probe Selected: http-probe

Step 3: On the Real Server tab, click **Add**.

Step 4: In the New Real Server dialog box, in the **Name** list, choose **webserver3**, and then in the **Port** box, enter **80** for HTTP.

Step 5: Click **Deploy Now**.

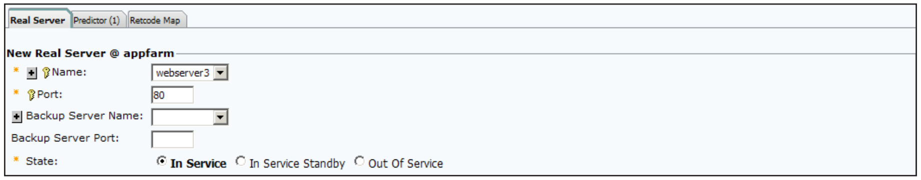

New Real Server @ appfarm

* Name:	webserver3
* Port:	80
Backup Server Name:	
Backup Server Port:	
* State:	● **In Service** ○ In Service Standby ○ Out Of Service

Step 6: Click **Deploy Now** for the newly created server farm. This saves your changes.

Step 7: On the **Real Server** tab, and then click **Add**.

Step 8: In the New Real Server dialog box, in the Name list, choose **webserver4**, and then in the **Port** box, enter **80**.

Step 9: Click **Deploy Now**.

Step 10: On the Edit Server Farm dialog box, click **Deploy Now**.

You have just created the server farm, appfarm, with the real-server members, webserver3 and webserver4, for HTTP on port 80. The Cisco ACE appliance will perform all of the SSL-processing so, even though clients will access the application on these servers via HTTPS, the traffic from Cisco ACE to the servers will happen over port 80. The http-probe will monitor all of the servers in the server farm to ensure that the HTTP service is available.

Procedure 3 Configure SSL proxy service

In order for Cisco ACE to offload the SSL processing, you need to configure an SSL proxy service. In this guide, the Cisco sample key and certificate is used. However, in a production deployment, you would most likely purchase a certificate from a trusted certificate authority (CA).

Step 1: Navigate to **Config > Virtual Contexts > SSL > Proxy Service**, and then click **Add**.

Step 2: In the New Proxy Service dialog box, in the **Name** box, enter **app-ssl-proxy**.

Step 3: Select both **cisco-sample-key** and **cisco-sample-cert**, and then click **Deploy Now**.

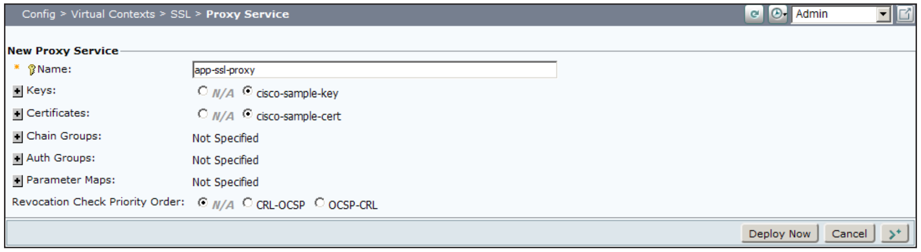

New Proxy Service

- * Name: app-ssl-proxy
- Keys: ○ N/A ● cisco-sample-key
- Certificates: ○ N/A ● cisco-sample-cert
- Chain Groups: Not Specified
- Auth Groups: Not Specified
- Parameter Maps: Not Specified
- Revocation Check Priority Order: ● N/A ○ CRL-OCSP ○ OCSP-CRL

Procedure 4 Configure HTTP-cookie sticky service

The HTTP cookie sticky service keeps traffic from a client "stuck" to a single real server. This is useful for applications where state could be lost if the client connection was balanced across several servers.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Stickiness**, and then click **Add**.

Step 2: In the New Sticky Group dialog box, in the **Group Name** box, enter **app-sticky**.

Step 3: In the Type list, choose **HTTP Cookie**, and in the **Cookie Name** box, enter **APPSESSIONID.**

Step 4: Select both **Enable Insert** and **Browser Expire.**

Step 5: Next to Sticky Server Farm, select **appfarm**, and then click **Deploy Now.**

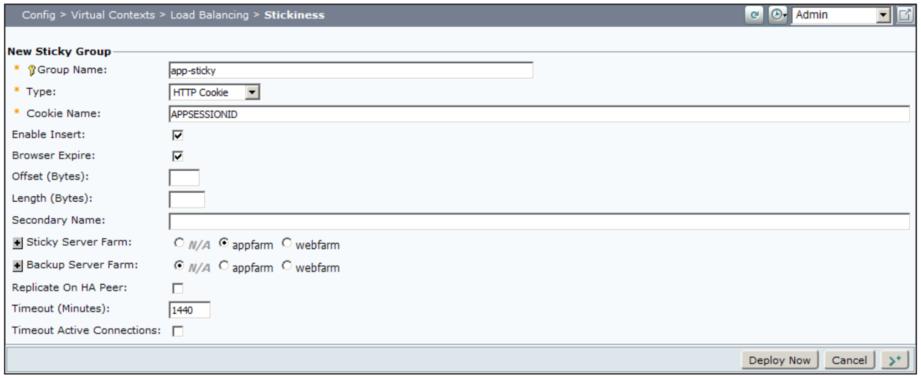

New Sticky Group

- * Group Name: app-sticky
- * Type: HTTP Cookie
- * Cookie Name: APPSESSIONID
- Enable Insert: ☑
- Browser Expire: ☑
- Offset (Bytes):
- Length (Bytes):
- Secondary Name:
- Sticky Server Farm: ○ N/A ● appfarm ○ webfarm
- Backup Server Farm: ● N/A ○ appfarm ○ webfarm
- Replicate On HA Peer: □
- Timeout (Minutes): 1440
- Timeout Active Connections: □

Procedure 5 Configure a virtual server

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Virtual Servers**, and then click **Add**.

Step 2: In the Properties dialog box, enter the following values:

- Virtual Server Name—https-vip
- Virtual IP Address—10.4.49.101
- Application Protocol—HTTPS
- VLAN—149

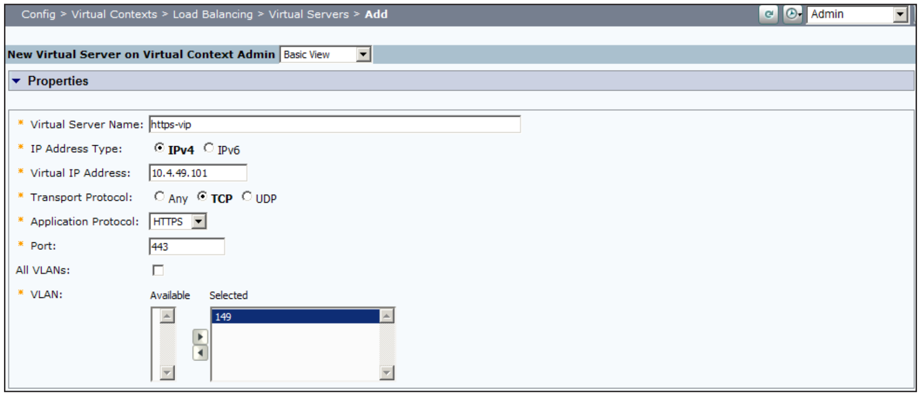
Config > Virtual Contexts > Load Balancing > Virtual Servers > **Add**

New Virtual Server on Virtual Context Admin Basic View

Properties

* Virtual Server Name: https-vip

* IP Address Type: ● IPv4 ○ IPv6

* Virtual IP Address: 10.4.49.101

* Transport Protocol: ○ Any ● TCP ○ UDP

* Application Protocol: HTTPS

* Port: 443

All VLANs: □

* VLAN: Available Selected 149

Step 3: In the SSL Termination dialog box, in the **Proxy Service Name** list, choose **app-ssl-proxy**.

Step 4: In the Default L7 Load-Balancing Action dialog box, in the **Primary Action** list, choose **Sticky**.

Step 5: In the **Sticky Group** list, choose **app-sticky (HTTP Cookie)**, and then select **Deflate**.

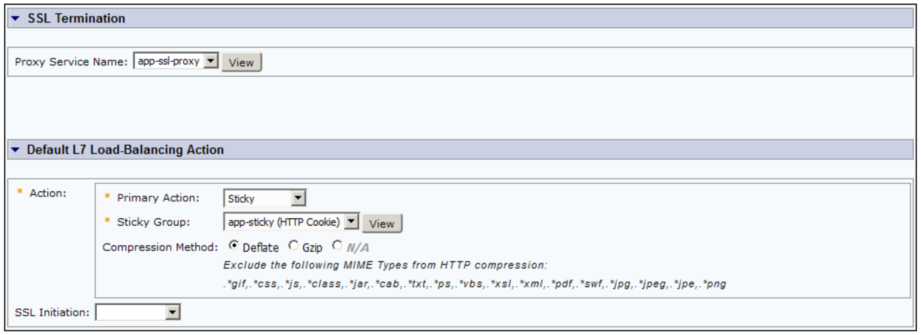
SSL Termination

Proxy Service Name: app-ssl-proxy

Default L7 Load-Balancing Action

* Action:

* Primary Action: Sticky

* Sticky Group: app-sticky (HTTP Cookie)

Compression Method: ● Deflate ○ Gzip ○ N/A

Exclude the following MIME Types from HTTP compression:

*.*gif,.*css,.*js,.*class,.*jar,.*cab,.*txt,.*ps,.*vbs,.*xsl,.*xml,.*pdf,.*swf,.*jpg,.*jpeg,.*jpe,.*png*

SSL Initiation:

Step 6: In the NAT dialog box, click **Add**, click **OK**, and then click **Deploy Now**.

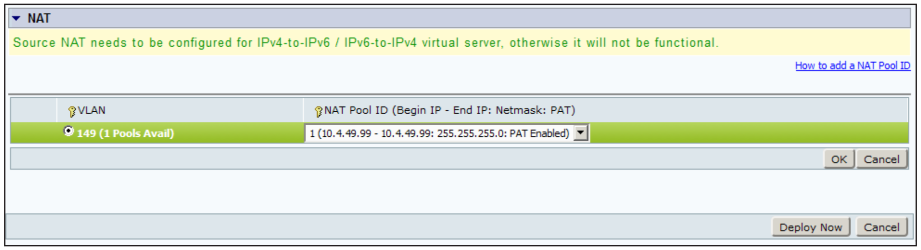
NAT

Source NAT needs to be configured for IPv4-to-IPv6 / IPv6-to-IPv4 virtual server, otherwise it will not be functional.

	VLAN	NAT Pool ID (Begin IP - End IP: Netmask: PAT)
	● 149 (1 Pools Avail)	1 (10.4.49.99 - 10.4.49.99: 255.255.255.0: PAT Enabled)

Clients going to the virtual IP, 10.4.49.101 on port 443, will be load-balanced across the real-servers, webserver3 and webserver4, in the server farm, appfarm. Cisco ACE will terminate the SSL session and load-balance the connections to the real-servers over standard HTTP on TCP port 80.

Procedure 6 Configure an HTTP-to-HTTPS Redirect

(Optional)

It is often preferable to have HTTP traffic redirected to HTTPS to ensure that connections to that service are encrypted. By following this procedure, you can create a service that redirects any HTTP traffic directed to 10.4.49.101 to the HTTPS service configured above.

Step 1: Navigate to **Config > Virtual Contexts > Load Balancing > Real Servers**, and then click **Add**.

Step 2: In the New Real Server dialog box, enter the values below, and then click **Deploy Now**.

- Name—**redirect1**
- Type—**Redirect**
- Web Host Redirection—**https://%h%p**
- Redirection Code—**302**

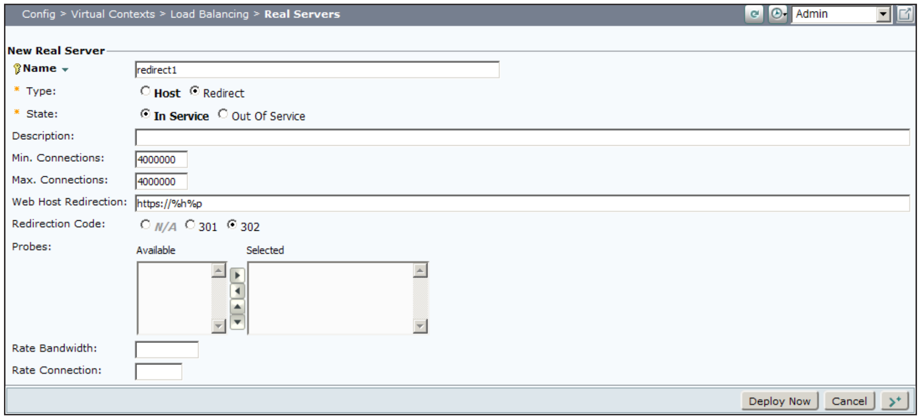

Config > Virtual Contexts > Load Balancing > **Real Servers**

Admin

New Real Server

Name: redirect1

* Type: ○ Host ● Redirect

* State: ● In Service ○ Out Of Service

Description:

Min. Connections: 4000000

Max. Connections: 4000000

Web Host Redirection: https://%h%p

Redirection Code: ○ N/A ○ 301 ● 302

Probes: Available Selected

Rate Bandwidth:

Rate Connection:

Step 3: Navigate to **Config > Virtual Contexts > Load Balancing > Server Farms**, and then click **Add**.

Step 4: In the New Server Farm dialog box, enter the values below, and then click **Deploy Now**.

- Name—**http-redirect**
- Type—**Redirect**

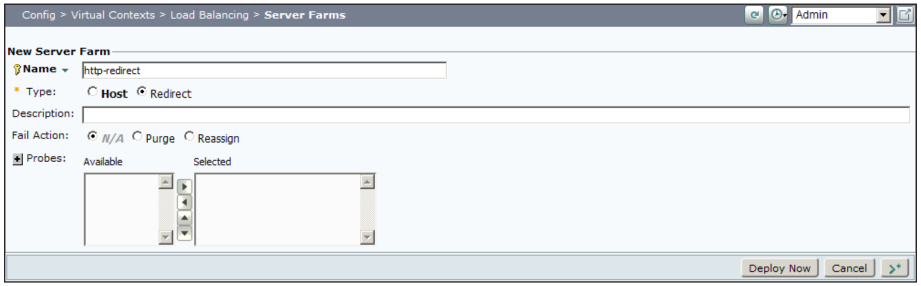

Config > Virtual Contexts > Load Balancing > **Server Farms**

Admin

New Server Farm

Name: http-redirect

* Type: ○ Host ● Redirect

Description:

Fail Action: ● N/A ○ Purge ○ Reassign

Probes: Available Selected

Step 5: Click the **Real Server** tab, and then click **Add**.

Step 6: In the New Real Server dialog box, select **redirect1**, and then click **Deploy Now**.

Step 7: On the Edit Server Farm dialog box, click **Deploy Now**.

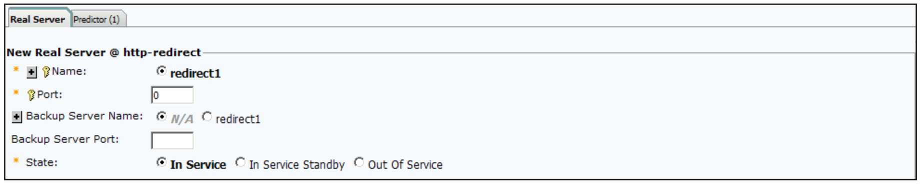

New Real Server @ http-redirect

* Name: ● redirect1

* Port: 0

Backup Server Name: ● N/A ○ redirect1

Backup Server Port:

* State: ● In Service ○ In Service Standby ○ Out Of Service

Step 8: Navigate to **Config > Virtual Contexts > Load Balancing > Virtual Servers**, and then click **Add**.

Step 9: In the Properties dialog box, enter the following values:

- Virtual Server Name—**http-vip-redirect**
- Virtual IP Address—**10.4.49.101**
- VLAN—**149**

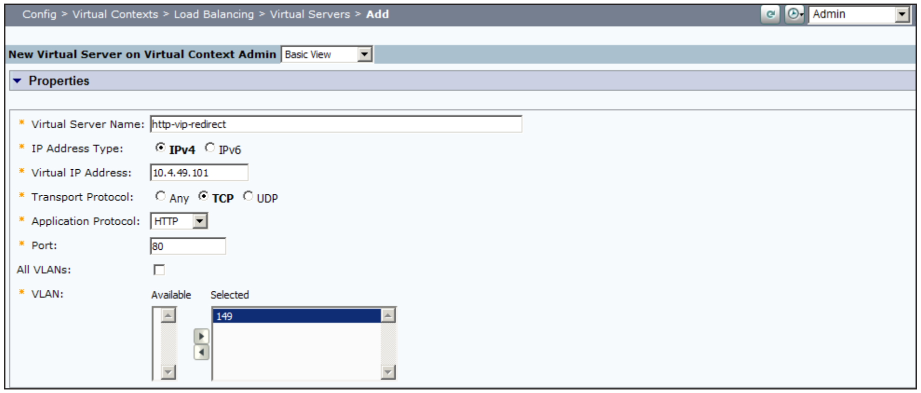

Config > Virtual Contexts > Load Balancing > Virtual Servers > **Add**

Admin

New Virtual Server on Virtual Context Admin Basic View

Properties

* Virtual Server Name: http-vip-redirect

* IP Address Type: ● IPv4 ○ IPv6

* Virtual IP Address: 10.4.49.101

* Transport Protocol: ○ Any ● TCP ○ UDP

* Application Protocol: HTTP

* Port: 80

All VLANs: □

* VLAN: Available Selected 149

Step 10: In the Default L7 Load-Balancing Action dialog box, in the **Server Farm** list, choose **http-redirect**, and then click **Deploy Now**.

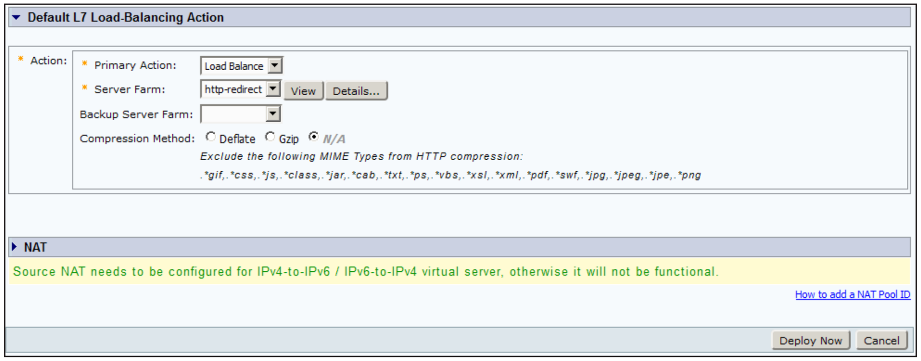

Default L7 Load-Balancing Action

* Action:

* Primary Action: Load Balance

* Server Farm: http-redirect

Backup Server Farm:

Compression Method: ○ Deflate ○ Gzip ● N/A

Exclude the following MIME Types from HTTP compression:

*.*gif,.*css,.*js,.*class,.*jar,.*cab,.*txt,.*ps,.*vbs,.*xsl,.*xml,.*pdf,.*swf,.*jpg,.*jpeg,.*jpe,.*png*

NAT

Source NAT needs to be configured for IPv4-to-IPv6 / IPv6-to-IPv4 virtual server, otherwise it will not be functional.

Notes

Appendix A: Product List

Data Center Core

Functional Area	Product Description	Part Numbers	Software
Core Switch	Cisco Nexus 5596 up to 96-port 10GbE, FCoE, and Fibre Channel SFP+	N5K-C5596UP-FA	NX-OS 5.1(3)N1(1a) Layer 3 License
	Cisco Nexus 5596 Layer 3 Switching Module	N55-M160L30V2	
	Cisco Nexus 5548 up to 48-port 10GbE, FCoE, and Fibre Channel SFP+	N5K-C5548UP-FA	
	Cisco Nexus 5548 Layer 3 Switching Module	N55-D160L3	
Ethernet Extension	Cisco Nexus 2000 Series 48 Ethernet 100/1000BASE-T Fabric Extender	N2K-C2248TP-1GE	—
	Cisco Nexus 2000 Series 48 Ethernet 100/1000BASE-T (enhanced) Fabric Extender	N2K-C2248TP-E	
	Cisco Nexus 2000 Series 32 1/10 GbE SFP+, FCoE capable Fabric Extender	N2K-C2232PP-10GE	

Data Center Services

Functional Area	Product Description	Part Numbers	Software
Application Resiliency	Cisco ACE 4710 Application Control Engine 2Gbps	ACE-4710-02-K9	A5(1.2)
	Cisco ACE 4710 Application Control Engine 1Gbps	ACE-4710-01-K9	
	Cisco ACE 4710 Application Control Engine 1Gbps 2-Pack	ACE-4710-2PAK	
	Cisco ACE 4710 Application Control Engine 500 Mbps	ACE-4710-0.5-K9	
Firewall	Cisco ASA 5585-X Security Plus IPS Edition SSP-40 and IPS SSP-40 bundle	ASA5585-S40P40-K9	ASA 8.4.3, IPS 7.1(4) E4
	Cisco ASA 5585-X Security Plus IPS Edition SSP-20 and IPS SSP-20 bundle	ASA5585-S20P20X-K9	
	Cisco ASA 5585-X Security Plus IPS Edition SSP-10 and IPS SSP-10 bundle	ASA5585-S10P10XK9	

Storage Network Extension

Functional Area	Product Description	Part Numbers	Software
Fibre-channel Switch	Cisco MDS 9148 Multilayer Fibre Channel Switch	DS-C9148D-8G16P-K9	NX-OS 5.0(7)
	Cisco MDS 9124 Multilayer Fibre Channel Switch	DS-C9124-K9	

Computing Resources

Functional Area	Product Description	Part Numbers	Software
UCS Fabric Interconnect	Cisco UCS up to 48-port Fabric Interconnect	UCS-FI-6248UP	2.0(2q) Cisco UCS Release
	Cisco UCS 20-port Fabric Interconnect	N10-S6100	
	Cisco UCS 6100 6-port Fibre Channel Expansion Module	N10-E0060	
UCS B-Series Blade Servers	Cisco UCS Blade Server Chassis	N20-C6508	2.0(2q) Cisco UCS Release
	Cisco UCS 8-port 10GbE Fabric Extender	UCS-IOM2208XP	
	Cisco UCS 4-port 10GbE Fabric Extender	UCS-IOM2204XP	
	Cisco UCS 4-port 10GbE First Generation Fabric Extender	N20-I6584	
	Cisco UCS B200 M2 Blade Server	N20-B6625-1	
	Cisco UCS B250 M2 Blade Server	N20-B6625-2	
	Cisco UCS M81KR Virtual Interface Card	N20-AC0002	
UCS C-Series Rack-mount Servers	Cisco UCS C200 M2 Rack Mount Server	R200-1120402W	1.4.1e Cisco UCS CIMC Release
	Cisco UCS C210 M2 Rack Mount Server	R210-2121605W	
	Cisco UCS C250 M2 Rack Mount Server	R250-2480805W	

Appendix B: Changes

This appendix summarizes the changes to this guide since the previous Cisco SBA series.

- We updated the “Ethernet Infrastructure” chapter with a QoS procedure to protect multimedia, control, and storage transport traffic. It also now includes configuration guidance on using enhanced VPC for Cisco Nexus FEX connectivity.
- The Computing Resources module is now named “Compute Connectivity” and we updated it to explain various methods for connecting servers to the data center network.
- The “Application Resiliency” chapter now includes an Inband Health check procedure to create a more robust server status check to detect failures faster.

Notes

Feedback

Click here to provide feedback to Cisco SBA.

SMART BUSINESS ARCHITECTURE

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLiERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

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.