

DCUCD

Data Center Unified Computing Design

Volume 2

Version 4.0

Student Guide

Text Part Number: 97-3025-01




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

 Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS" AND AS SUCH MAY INCLUDE TYPOGRAPHICAL, GRAPHICS, OR FORMATTING ERRORS. CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 2

Design Cisco Unified Computing System Solution **3-1**

Overview	3-1
Objectives	3-1

Evaluating Cisco UCS C-Series Architecture **3-3**

Overview	3-3
Objectives	3-3
Evaluating the C-Series Server	3-4
Evaluating C-Series Memory	3-17
Evaluating C-Series LAN and SAN Connectivity	3-23
Evaluating C-Series Local Storage	3-32
Evaluating C-Series Management	3-36
Summary	3-41

Sizing the Cisco UCS C-Series Solution **3-43**

Overview	3-43
Objectives	3-43
Analyzing Requirements	3-44
Sizing the C-Series Solution	3-47
Creating the C-Series BOM	3-55
Summary	3-58

Evaluating Cisco UCS B-Series Architecture **3-59**

Overview	3-59
Objectives	3-59
Evaluating Cisco UCS B-Series	3-60
Cisco UCS 6100 Fabric Interconnects	3-65
Cisco UCS 6120XP Fabric Interconnect	3-65
Cisco UCS 6140XP Fabric Interconnect	3-67
Cisco UCS 5108 Chassis	3-72
Chassis Power	3-72
Chassis Cooling	3-73
Chassis Power Supply Modes	3-74
Chassis Power Connectivity	3-74
Chassis Management Controller	3-76
Cisco UCS B-Series Server Blades	3-80
Cisco IMC	3-80
Advanced Functionality	3-81
B200-M1 Blade	3-82
B200-M2 Blade	3-82
B230-M1 Blade	3-83
B250-M1 Blade	3-84
B250-M2 Blade	3-84
B440 M1 Blade	3-85
Cisco UCS B-Series Memory	3-87
Cisco UCS B-Series Adapters	3-89
Cisco UCS NIC 82598KR-CI	3-89
Cisco UCS NIC M51KR-B Broadcom BCM57711	3-89
Cisco UCS CNA M61KR-I Intel	3-90
Cisco UCS CNA M71KR	3-90
Cisco UCS CNA M72KR	3-90
Cisco UCS VIC M81KR	3-90
Cisco VN-Link Technology	3-92
VMware Integration	3-92
Hypervisor Bypass	3-92
Use Cases	3-93
Cisco UCS B-Series Management	3-96

Data Center Operating System Network—VLAN 4042	3-97
Adapter Management Network—VLAN 4043	3-97
CMC	3-98
BMC	3-98
Cisco UCS Manager Controller	3-110
Cisco NX-OS System Manager	3-110
Cisco UCS B-Series LAN Connectivity	3-112
EHV Mode	3-112
Switching Mode	3-113
VLANs	3-117
Uplink Ports	3-118
Port Channels	3-118
Server Blade to IOM Server Ports—vNIC	3-119
Cisco UCS CNA M71KR/M72KR, UCS 82598KR-CI Adapter	3-120
Cisco UCS VIC M81KR Adapter	3-120
Cisco UCS B-Series SAN Connectivity	3-128
VSANs	3-132
Summary	3-141
Sizing the Cisco UCS B-Series Solution	3-143
Overview	3-143
Objectives	3-143
Analyzing Requirements	3-144
Sizing the B-Series Solution	3-149
Creating the B-Series BOM	3-164
Summary	3-170
Planning Physical Deployment	3-171
Overview	3-171
Objectives	3-171
Cisco UCS Power Calculator Tool	3-172
Creating the Physical Deployment Plan	3-174
Summary	3-179
Examining the Cisco UCS Network and Storage	3-181
Overview	3-181
Objectives	3-181
Cisco UCS LAN	3-182
Cisco Nexus 1000V	3-192
Cisco Nexus 1000V VMware Integration	3-205
Cisco UCS SAN	3-225
Cisco UCS Storage	3-230
Summary	3-234
Designing the Cisco UCS Network and Storage	3-235
Overview	3-235
Objectives	3-235
Analyzing Requirements	3-236
Designing the Cisco UCS LAN	3-241
Designing the Cisco UCS SAN	3-254
Summary	3-261
Module Summary	3-263
Module Self-Check	3-265
Module Self-Check Answer Key	3-269

<i>Design Server Deployment</i>	<i>4-1</i>
Overview	4-1
Objectives	4-1
<i>Designing the Cisco UCS Server Deployment Model</i>	<i>4-3</i>
Overview	4-3
Objectives	4-3
Designing Basic B-Series Server Deployment	4-4
Advanced Server Deployment Model	4-11
Designing Identity Pools	4-14
Designing Resource Pools	4-25
Design Policies	4-29
Cisco UCS 1.4 Enhancements	4-38
Summary	4-40
Module Summary	4-41
Module Self-Check	4-43
Module Self-Check Answer Key	4-44
<i>Design Cisco UCS Solution Management</i>	<i>5-1</i>
Overview	5-1
Objectives	5-1
<i>Examining Cisco UCS Solution Management</i>	<i>5-3</i>
Overview	5-3
Objectives	5-3
Cisco UCS Management Aspects	5-4
Managing Multiple Cisco UCS Pods	5-21
Cisco UCS 1.4 Enhancements	5-26
Summary	5-29
<i>Designing Cisco UCS Solution Management</i>	<i>5-31</i>
Overview	5-31
Objectives	5-31
Analyzing Requirements	5-32
Designing Cisco UCS Management	5-36
Summary	5-42
Module Summary	5-43
Module Self-Check	5-45
Module Self-Check Answer Key	5-46
<i>Design Advanced Server Deployment</i>	<i>6-1</i>
Overview	6-1
Objectives	6-1
<i>Evaluating Cisco UCS Deployment with Microsoft Hyper-V</i>	<i>6-3</i>
Overview	6-3
Objectives	6-3
Assessing Microsoft Hyper-V R2 Requirements	6-4
Designing Microsoft Hyper-V R2 Deployment	6-8
Summary	6-15
<i>Evaluating Cisco UCS Integration with VMware vSphere</i>	<i>6-17</i>
Overview	6-17
Objectives	6-17
Assessing VMware vSphere Requirements	6-18
Designing VMware vSphere Deployment	6-22
Summary	6-48

Evaluating Cisco UCS and Cisco Nexus 1000V Integration with VMware vSphere 6-49

Overview	6-49
Objectives	6-49
Assessing VMware vSphere with Cisco Nexus 1000V Requirements	6-50
Designing VMware vSphere with Cisco Nexus 1000V Deployment	6-57
Summary	6-72
Module Summary	6-73
Module Self-Check	6-75
Module Self-Check Answer Key	6-77

Design Cisco Unified Computing System Solution

Overview

This module identifies and explains the Cisco Unified Computing System (UCS) C-Series and B-Series hardware components and employs design steps to create a computing solution for a given set of requirements.

Objectives

Upon completing this module, you will be able to understand the Cisco UCS hardware components and select the proper hardware for a given set of requirements. This includes the ability to meet these objectives:

- Assess the Cisco UCS C-Series hardware components, connectivity, high-availability options, and management
- Identify how to assemble a Cisco UCS C-Series solution for a given set of requirements
- Assess the Cisco UCS B-Series hardware components, connectivity, high-availability options, and management
- Discuss how to assemble a Cisco UCS B-Series solution for a given set of requirements
- Discuss a physical deployment plan for the Cisco UCS solution
- Discuss the LAN, SAN, application service, and storage products and technologies that are relevant to Cisco UCS for a given set of requirements
- Discuss design of the LAN and SAN components of the Cisco UCS solution

Evaluating Cisco UCS C-Series Architecture

Overview

This lesson assesses the Cisco UCS C-Series hardware components, connectivity, high-availability options, and management.

Objectives

Upon completing this lesson, you will be able to identify and describe the Cisco UCS C-Series. This includes the ability to meet these objectives:


- Evaluate C-Series rack server options
- Evaluate C-Series memory options
- Evaluate C-Series LAN and SAN connectivity options
- Evaluate C-Series RAID and disk options
- Evaluate the Cisco UCS IMC

Evaluating the C-Series Server

This topic identifies and describes C-Series rack server options.

UCS C-Series Rack Servers Overview

- CPU sockets for Intel Xeon architecture
 - Must be of the same type per server
- Internal SAS/SATA disk drive bays
- 2-, 4-, 8-, and 16-GB DDR3 memory modules
- In-chassis USB key slot for security key or hypervisor boot
- Support for statelessness (MAC, WWN, UUID, BIOS, boot order)
- Management via:
 - Cisco IMC
 - Cisco UCS Manager (requires v1.4)



Property	Value
Operating temperature	10 to 35 °C
Operating humidity	5 to 93% noncondensing
Operating altitude	0 to 3000 m

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-4

Cisco UCS C-Series rack-mount servers, which are based on Intel Xeon 5500, 5600 and 7500 series processors, extend unified computing innovations to an industry-standard form factor. Designed to operate both in standalone environments and as part of Cisco UCS, the series incorporates standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology. It supports an incremental deployment model with an easy migration path to unified computing.


Intel Xeon 5500, 5600 and 7500 series processors automatically and intelligently adjust server performance according to application needs, increasing performance when needed and achieving substantial energy savings when not. Performance and power settings can also be manually configured.

When deployed as standalone servers in a heterogeneous environment, Cisco UCS C-Series servers can be managed just like any other x86-architecture servers. Popular enterprise management tools using operating system-resident host agents work without modification. Cisco UCS Integrated Management Controller (IMC) gives administrators the tools they need to manually control server functions, including remote keyboard, video, and mouse (KVM); power on and off; and standard Simple Network Management Protocol (SNMP) traps for system monitoring.

C200 M1 and M2 Servers

CPU	2 Intel Xeon processor architecture 5500 or 5600 (4-core, 6-core)
Memory	12 DIMM slots—up to 96 GB RAM
Disk	4 3.5" hot-swappable disks (RAID 0, 1, 5, 6, 10)
On-board interfaces	2 1-Gigabit Ethernet, 1 mgmt Ethernet ports 2 USB, 1 serial
Options	CD/DVD drive 2 PCIe Gen 2 slots for additional: <ul style="list-style-type: none">• Gigabit Ethernet interfaces• Fibre Channel interfaces• 10 Gigabit Ethernet interfaces• RAID disk controller

1 RU



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-3-5

The Cisco UCS C200 M1 and M2 High-Density Rack-Mount Servers are high-density, two-socket, one-rack-unit (1 RU) rack-mount servers that are designed to balance simplicity, performance, and density. Powered with up to two quad-core Intel Xeon 5500 or 5600 series processors, the server is built to support production-level network infrastructure, web services, and mainstream data center, branch-office, and remote-office applications. Its compact size makes it useful for service providers offering dedicated or multitenant hosting.

Energy efficiency is even more important for high-density, 1-RU servers because of the number of processors that can be placed into a rack. Both the CPUs and power supplies contribute to the excellent energy efficiency of the Cisco UCS C200 M1 or M2 server. The Intel Xeon 5500 and 5600 series processors feature Intel Intelligent Power Technology that automatically adjusts processor energy consumption to match workload requirements. With a greater number of automated low-power states, processor cores can go to near-zero power consumption independently. Further, the Cisco UCS C200 M1 and M2 servers are equipped with a single or redundant pair of power supplies that meet Climate Smart specifications, increasing efficiency. Both these factors help enable increased server density in modern data centers.

C200 M1 and M2 series servers have the following characteristics:

- **Flexible I/O and storage options:** With two PCI Express (PCIe) expansion slots, the server offers I/O flexibility and bandwidth, including the ability to integrate with traditional Gigabit Ethernet LANs and Fibre Channel SANs. The server hosts up to four internal SAS or Serial Advanced Technology Attachment (SATA) drives, providing internal storage capacity exceeding what is available in a corresponding blade form-factor server.
- **10 Gigabit unified network fabric:** When equipped with converged network adapters (CNAs) or the Cisco UCS P81E Virtual Interface Card, the server integrates with a low-latency, lossless 10-Gb Ethernet, and industry-standard Fibre Channel over Ethernet (FCoE) fabric. This technology enables a “wire-once” deployment model in which changing I/O configurations no longer means installing adapters and recabling racks and switches.

- **Virtualization optimization:** Cisco VN-Link technology, I/O virtualization, and Intel Xeon 5500 series processor features extend the network directly to virtual machines. This optimization enables a consistent and scalable operational model, helping increase security and efficiency while reducing complexity.
- **Unified management:** When integrated as a part of Cisco UCS, management is uniquely integrated into all components of the system. This enables the entire solution to be managed as a single entity through Cisco UCS Manager, improving operational efficiency and flexibility.
- **Service profiles:** When integrated as part of Cisco UCS, Cisco UCS Manager implements role- and policy-based management using service profiles and templates. Service profiles help automate provisioning and increase business agility, allowing data center managers to provision applications in minutes instead of days.
- **Up to two quad-core Intel Xeon 5500 or 5600 series processors:** These multicore processors automatically and intelligently adjust server performance according to application needs, increasing performance when needed and achieving substantial energy savings when not needed.
- **Up to 96 GB of industry-standard DDR3 main memory (using 12 8-GB DIMMs):**
 - 12 DIMM slots for up to 96 GB of memory using 8-GB DIMMs
 - Support for DDR3 registered DIMMs
 - ECC and ChipKill support
 - Mirroring option
- **Up to four internal SAS or SATA drives for up to a total of 4 TB with RAID 0, 1, 5, and 6 support.**
- **Support for up to one low-profile, half-length, and one full-height, half-length x8 PCIe card:** The full-height slot uses an x16 connector.
 - Two PCIe Gen 2.0 slots available
 - One x16 full-height and one x8 low-profile slots, both half-length
 - x16 connector on full-height slot and x8 connector on low-profile slot
- **Two integrated Gigabit Ethernet ports and one 10/100-Mb/s Ethernet management port for accessing the Cisco UCS Integrated Management Controller.**
- **Cisco UCS Integrated Management Controller:**
 - Integrated ServerEngines Pilot-2 baseboard management controller (BMC)
 - Intelligent Platform Management Interface (IPMI) 2.0 compliant for management and control
 - One 10/100BASE-T out-of-band management interface
 - CLI and WebGUI management tool for automated, lights-out management
 - KVM
- **Front-panel CD/DVD drive, locator LED, and interface with video, two USB, and serial port connections.**
- **Back-panel video, two USB, and serial port connectors.**
- **Increased reliability, availability, and serviceability through optional dual-redundant power supplies.**

C200 M1 Server Processor Support

Processor support is as follows:

- One or two Intel Xeon Series 5500 processors
- Choice of processors: Intel Xeon X5570, X5550, E5540, E5520, L5520, or E5504

C200 M2 Server Processor Support

Processor support is as follows:

- One or two Intel Xeon Series 5500 or 5600 processors
- Choice of processors: Intel Xeon X5570, X5550, E5540, E5520, L5520, E5504, or Intel Xeon X5670, X5650, L5640, E5640, E5620, or E5506

C210 M1 and M2 Servers

CPU	2 Intel Xeon processor architecture 5500 or 5600 (4-core, 6-core)
Memory	12 DIMM slots—up to 96 GB RAM
Disk	16 2.5" hot-swappable disks (RAID 0, 1, 5, 6, 10, 50, 60)
On-board interfaces	2 1-Gigabit Ethernet, 1 mgmt Ethernet port 2 USB, 1 serial
Options	CD/DVD drive 5 PCIe Gen 2 slots for additional <ul style="list-style-type: none">• Gigabit Ethernet interfaces• Fibre Channel interfaces• 10 Gigabit Ethernet interfaces• RAID disk controller

2 RU



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-6

The Cisco UCS C210 M1 and M2 General-Purpose Rack-Mount Servers are general-purpose, two-socket, 2-RU rack-mount servers that are designed to balance performance, density, and efficiency for workloads requiring economical, high-capacity, reliable internal storage.

Designed for applications such as virtualization, network file servers, application servers, appliances, storage servers, database servers, and content-delivery servers, the Cisco UCS C210 M1 server packs up to 16 small form-factor SAS or SATA disk drives into only 2 RU, for a total of up to 8 TB of storage.

Not all storage-intensive workloads are alike, and the server disk configuration of the Cisco UCS C210 M1 and M2 servers adapts to balance performance and economy to best meet individual workload requirements. Both 10,000- and 15,000-rpm SAS drives deliver a high number of I/O operations per second for transactional workloads such as database management systems. High-capacity SATA drives provide an economical solution for applications, including content-delivery servers. A choice of three RAID controller options helps increase performance and reliability and the flexibility of the servers.

C210 M1 and M2 series servers have the following characteristics:

- **Flexible I/O and storage options:** With five PCI Express (PCIe) expansion slots, the server offers I/O flexibility and bandwidth, including the ability to integrate with traditional Gigabit Ethernet LANs and Fibre Channel SANs and incorporate a range of RAID controller options to support up to 16 drives.
- **10 Gigabit unified network fabric:** When equipped with CNAs or the Cisco UCS P81E Virtual Interface Card, the server integrates with a low-latency, lossless 10 Gigabit Ethernet and industry-standard FCoE fabric. This technology enables a “wire-once” deployment model in which changing I/O configurations no longer means installing adapters and recabling racks and switches.
- **Virtualization optimization:** Cisco VN-Link technology, I/O virtualization, and Intel Xeon 5500 series processor features extend the network directly to virtual machines. This optimization enables a consistent and scalable operational model, helping increase security and efficiency while reducing complexity.

- **Unified management:** When integrated as a part of Cisco UCS, management is uniquely integrated into all components of the system. This enables the entire solution to be managed as a single entity through Cisco UCS Manager, improving operational efficiency and flexibility.
- **Service profiles:** When integrated as part of Cisco UCS, Cisco UCS Manager implements role- and policy-based management using service profiles and templates. Service profiles help automate provisioning and increase business agility, allowing data center managers to provision applications in minutes instead of days.
- Up to two quad-core Intel Xeon 5500 or 5600 series processors: These multicore processors automatically and intelligently adjust server performance according to application needs, increasing performance when needed and achieving substantial energy savings when not needed.
- Up to 96 GB of industry-standard DDR3 main memory (using 12 8-GB DIMMs)
- Up to 16 internal SFF SAS or SATA drives for a total of up to 8 TB with RAID 0, 1, 5, 6, 10, 50, and 60 support.
- Room for up to five full-height PCIe cards: two full-height, full-length x8 cards, and three full-height, half-length x8 cards; all slots use x16 connectors
- Two integrated Gigabit Ethernet ports and one 10/100-Mb/s Ethernet management port for accessing the Cisco UCS Integrated Management Controller
- Cisco UCS Integrated Management Controller
 - Integrated ServerEngines Pilot-2 BMC
 - IPMI 2.0 compliant for management and control
 - One 10/100BASE-T out-of-band management interface
 - CLI and WebGUI management tool for automated, lights-out management
 - KVM
- Optional front-panel CD/DVD drive, locator LED, and interface with video, two USB, and serial port connections
- Back-panel video, two USB, and serial port connectors
- Increased reliability, availability, and serviceability through optional dual-redundant power supplies that meet Climate Smart specifications

C250 M1 and M2 Servers

CPU	2 Intel Xeon processor architecture 5500/5600 (4-core, 6-core)
Memory	48 DIMM slots—up to 384GB RAM
Disk	8 2.5" hot-swappable disks (RAID 0, 1, 5, 6, 10, 50, 60)
On-board interfaces	4 1-Gigabit Ethernet, 2 mgmt Ethernet ports 2 USB, 1 serial
Options	CD/DVD drive 5 PCIe Gen 2 slots for additional <ul style="list-style-type: none">• Gigabit Ethernet interfaces• Fibre Channel interfaces• 10 Gigabit Ethernet interfaces• RAID disk controller

2 RU



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-7

The Cisco UCS C250 M1 and M2 Extended-Memory Rack-Mount Servers are high-performance, memory-intensive, two-socket, 2-RU rack-mount servers that are designed to increase performance and capacity for demanding virtualization and large-data-set workloads. They can also reduce the cost of smaller memory footprints by using lower-cost, lower-density memory. The system is built for virtualized workloads in enterprise data centers, service provider environments, and virtual desktop hosting. The system also helps increase performance for large-data-set workloads, including database management systems and modeling and simulation applications.

The Cisco UCS C250 M1 and M2 servers extend the Cisco product portfolio to meet the needs of customers that choose to deploy rack-mount servers. The server enables organizations to deploy systems incrementally—using as many or as few servers as needed—on a schedule that best meets the timing and budget of the organization.

Designed to operate both in standalone environments and as part of Cisco UCS, the servers combine flexible disk storage and I/O configurations with Cisco innovations including patented Cisco Extended Memory technology, a unified network fabric, and network-aware Cisco VN-Link technology. The server brings differentiation and value to what has been a commodity market with products not optimized to meet the needs of virtualized data centers. Available from Cisco and its data center network infrastructure (DCNI) partners, the server advances the rack-mount server market.

C250 M1 and M2 series servers have the following characteristics:


- **Cisco Extended Memory Technology:** This technology offers twice as much memory (384 GB) as traditional two-socket servers, or a more economical (192 GB) memory footprint using 4-GB rather than 8-GB DIMMs. Designed to strike a balance between processing power and memory capacity appropriate for virtualized and large-data-set workloads, the technology eliminates the need to upgrade to more expensive four-socket servers just to establish a large memory footprint. This technology helps lower both capital and operating costs.

- **Flexible I/O and storage options:** With five PCIe expansion slots, the server offers I/O flexibility and bandwidth, including the ability to integrate with traditional Gigabit Ethernet LANs and Fibre Channel SANs. The server hosts up to eight internal small form-factor SAS or SATA drives, providing internal storage capacity that exceeds what is available in a corresponding blade form-factor server.
- **10 Gigabit unified network fabric:** When equipped with CNAs or the Cisco UCS P81E Virtual Interface Card, the server integrates with a low-latency, lossless 10-Gigabit Ethernet, and industry-standard FCoE fabric. This technology enables a “wire-once” deployment model in which changing I/O configurations no longer means installing adapters and recabling racks and switches.
- **Virtualization optimization:** Cisco VN-Link technology, I/O virtualization, and Intel Xeon 5500 series processor features extend the network directly to virtual machines. This optimization enables a consistent and scalable operational model, helping increase security and efficiency while reducing complexity.
- **Unified management:** When integrated as a part of Cisco UCS, management is uniquely integrated into all components of the system, enabling the entire solution to be managed as a single entity through Cisco UCS Manager, improving operational efficiency and flexibility.
- **Service profiles:** When integrated as part of Cisco UCS, Cisco UCS Manager implements role- and policy-based management using service profiles and templates. Service profiles help automate provisioning and increase business agility, allowing data center managers to provision applications in minutes instead of days.
- Up to two quad-core Intel Xeon 5500 or 5600 series processors; these multicore processors automatically and intelligently adjust server performance according to application needs, increasing performance when needed and achieving substantial energy savings when not needed.
- Up to 384 GB of industry-standard DDR3 main memory (using 48 8-GB DIMMs)
- Up to eight internal SFF SAS or SATA drives for a total of up to 4 TB with RAID 0, 1, 5, 6, 10, 50, and 60 support.
- Support for up to five PCIe cards in three low-profile, half-length x8 and two full-height, half-length x16 slots; all slots use x16 connectors
- Four integrated Gigabit Ethernet ports and two 10/100-Mb/s Ethernet management ports for accessing the Cisco UCS Integrated Management Controller
- Cisco UCS Integrated Management Controller
 - Integrated ServerEngines Pilot-2 BMC
 - IPMI 2.0 compliant for management and control
 - One 10/100BASE-T out-of-band management interface
 - CLI and WebGUI management tool for automated, lights-out management
 - KVM
- Optional front-panel CD/DVD drive, locator LED, and interface with video, two USB, and serial port connections
- Back-panel video, two USB, and serial port connectors
- Increased reliability, availability, and serviceability through optional dual-redundant power supplies that meet Climate Smart specifications

C460 M1 Server

CPU	4 Intel Xeon processor architecture 7500 (4-, 6-, 8- core)
Memory	64 DIMM slots—up to 512GB RAM
Disk	12 2.5" hot-swappable disks (RAID 0,1,5,6,10,50,60)
On-board interfaces	2 10-Gigabit Ethernet, 2 OOB 1 Gigabit Ethernet interfaces 2 USB, 1 serial
Options	CD/DVD drive 10 PCIe Gen 2 slots for additional (4 are hot-pluggable) <ul style="list-style-type: none">• Gigabit Ethernet interfaces• Fibre Channel interfaces• 10 Gigabit Ethernet interfaces• RAID disk controller

4 RU



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-8

The Cisco UCS C460 M1 High-Performance Rack-Mount Server is a high-performance, high-memory-capacity server that is designed with the performance and reliability to power compute-intensive, enterprise-critical standalone applications and virtualized workloads. The system is a 4-RU rack-mount server supporting up to four Intel Xeon 7500 series processors, up to 512 GB of Samsung 40-nm class 1.35V double-data rate DDR3 memory in 64 slots, and 12 small form-factor hot-pluggable SAS and SATA disk drives. Abundant I/O capability is provided by 10 PCIe slots supporting the Cisco UCS C-Series network adapters, with an 11th PCIe slot reserved for SAS drive controller cards. Additional I/O is provided by four Ethernet LAN-on-motherboard (LOM) ports: two 10 Gigabit Ethernet and two 1 Gigabit Ethernet.

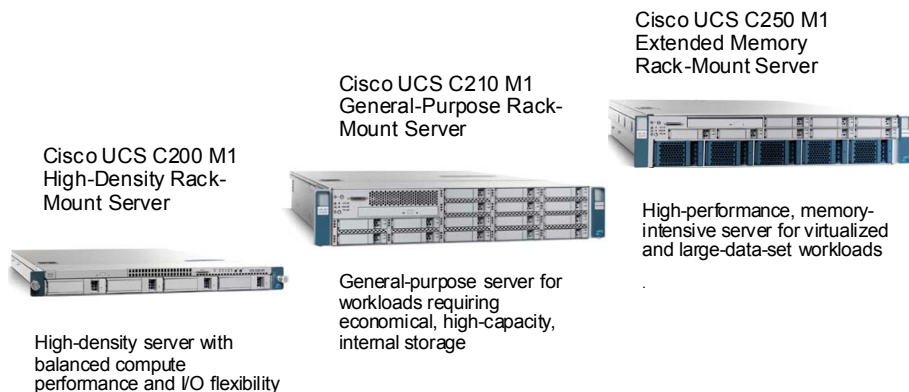
The Cisco UCS C460 M1 server extends the Cisco product portfolio to meet the needs of customers who choose to deploy rack-mount servers. The server enables organizations to deploy systems incrementally—using as many or as few servers as needed—on a schedule that best meets the timing and budget of the organization.

C460 M1 series servers have the following characteristics:

- **High performance:** With up to four Intel Xeon 7500 series processors with intelligent performance that automatically adapts to the diverse needs of a virtualized environment, and advanced reliability and exceptional scalability for the most data-demanding applications, the Cisco UCS C460 M1 server unifies the most performance-intensive and enterprise critical applications in the Cisco UCS architecture.
- **Exceptional memory capacity:** With up to 512 GB of Samsung high-efficiency memory, the Cisco UCS C460 M1 is designed to balance processing power and memory capacity for large-data-set, transaction-intensive, mission-critical workloads and provide headroom for jumbo virtual machines and greater levels of server consolidation.
- **Flexible I/O and storage options:** With 10 PCIe expansion slots, the server offers I/O flexibility and bandwidth, including the capability to integrate with both traditional Gigabit and 10 Gigabit Ethernet LANs and Fibre Channel SANs. The server hosts up to 12 internal small form-factor SAS or SATA drives, providing internal storage capacity exceeding what is available in a corresponding blade form-factor server.

- **10 Gigabit unified network fabric:** When equipped with CNAs, the server integrates with a low-latency, lossless 10 Gigabit Ethernet, and industry-standard FCoE fabric. This technology enables a “wire-once” deployment model in which changing I/O configurations no longer means installing adapters and recabling racks and switches.
- **Virtualization optimization:** Cisco VN-Link technology, I/O virtualization, and Intel Xeon 7500 series processor features extend the network directly to virtual machines. This optimization enables a consistent and scalable operational model, helping increase security and efficiency while reducing complexity.
- **Unified management:** When the server is integrated into Cisco UCS, Cisco UCS Manager provides management. Management is uniquely integrated into all components of the system, enabling the entire solution to be managed as a single entity through Cisco UCS Manager, improving operational efficiency and flexibility.
- **Service profiles:** When the server is integrated into Cisco UCS, Cisco UCS Manager implements role- and policy-based management using service profiles and templates. Service profiles help automate provisioning and increase business agility, allowing data center managers to provision applications in minutes instead of days.
- Support for up to 10 PCIe cards in four half-length and six three-quarter-length slots; an 11th slot is reserved for a SAS drive controller card
- Dual-port 1 Gigabit Ethernet LOM, dual-port 10 Gigabit Ethernet LOM, and two dedicated out-of-band management ports
- Cisco UCS Integrated Management Controller
 - Integrated ServerEngines Pilot-2 BMC
 - IPMI 2.0 compliant for management and control
 - One 10/100BASE-T out-of-band management interface
 - CLI and WebGUI management tool for automated, lights-out management
 - KVM
- Front-panel CD/DVD drive, locator LED, and interface with video, two USB, and serial port connections
- Back-panel video, two USB, and serial port connectors
- Increased reliability, availability, and serviceability through optional dual-redundant power supplies meeting Climate Saver specifications and front panel-accessible hot-swap cooling fans

C-Series Product Positioning



C200 M1 and M2 Positioning

The Cisco UCS C200 M1 and M2 servers are two-socket, 1-RU rack-mount servers that are designed to balance simplicity, performance, and density for web infrastructure and other mainstream data center workloads. These servers increase speed in standard and virtualized environments while helping optimize performance of storage-intensive applications

- Virtualization workloads using a single or a large pool of servers: the optional Cisco UCS P81E Virtual Interface Card brings the full power of Cisco UCS to the platform. This includes the capability to support up to 128 Ethernet or Fibre Channel virtual interfaces that are programmed on demand to meet the needs of both virtualized and nonvirtualized environments, and Intel Virtualization Technology for Direct I/O further speeds virtual machine I/O operations by facilitating direct control over physical interfaces from virtual machines.
- Database management systems can thrive on the abundant internal storage of the servers.
- Ideal as application servers, where multiple processor cores contribute directly to performance.
- I/O-intensive applications, including data warehousing, medical imaging, video surveillance, document imaging, and content distribution, have up to 8 TB of disk storage for application data.

C210 M1 and M2 Positioning

The Cisco UCS C210 M1 and M2 servers are general-purpose, two-socket, 2-RU rack-mount servers housing up to 16 internal small form-factor SAS or SATA disk drives for a total of up to 8 TB of storage. These servers are designed to balance performance, density, and efficiency for workloads requiring economical, high-capacity, reliable, internal storage. These servers increase speed in standard and virtualized environments while helping optimize performance of storage-intensive applications:

- Virtualization workloads using a single or a large pool of servers: the optional Cisco UCS P81E Virtual Interface Card brings the full power of Cisco UCS to the platform. This includes the capability to support up to 128 Ethernet or Fibre Channel virtual interfaces that are programmed on demand to meet the needs of both virtualized and nonvirtualized environments, and Intel Virtualization Technology for Direct I/O further speeds virtual machine I/O operations by facilitating direct control over physical interfaces from virtual machines.
- Database management systems can thrive on the abundant internal storage of the servers.
- The Cisco UCS C210 M2 is an ideal application server, where multiple processor cores contribute directly to performance.
- I/O-intensive applications, including data warehousing, medical imaging, video surveillance, document imaging, and content distribution, have up to 8 TB of disk storage for application data.

C250 M1 and M2 Positioning

The Cisco UCS C250 M1 and M2 servers are two-socket 2-RU rack-mount server featuring patented Cisco Extended Memory Technology. The servers are designed to increase performance and capacity for demanding virtualization and large-data-set workloads and to deliver a more cost-effective memory footprint for less-demanding workloads.

From a memory capacity perspective, this platform can alleviate memory bottlenecks in situations in which costly four-socket servers might otherwise be necessary, helping improve the price-to-performance ratio for running large-memory-footprint applications. From a memory-cost perspective, the servers can be populated with low-cost 4-GB DIMMs for a total of up to 192 GB of main memory. This memory configuration delivers a memory footprint that other two-socket, Intel Xeon 550/56000 series processor-based systems require 16-GB DIMMs to achieve. From a memory capacity perspective, the server can be populated with 8-GB DIMMs for a total of up to 384 GB of memory.

These benefits of Cisco Extended Memory Technology can be harnessed by customers when very large memory footprints are required, or when large, low-cost memory footprints are desirable, as in the following examples:

- Large virtualized environments can host more or larger virtual machines with the larger memory footprint, and with higher performance in cases in which existing implementations are memory-bound.
- Database applications will thrive in virtualized and nonvirtualized environments, as the server uses the combination of a large memory footprint and the fastest Intel processors.
- Traditional high-performance computing applications can benefit from the performance and memory footprint, including memory-intensive engineering design automation and geophysical modeling applications. Other memory-bound high-performance computing (HPC) applications are likely to see performance accelerations on a Cisco UCS C250 M1 server.
- Enterprise resource planning applications can run with improved performance with large data sets in main memory when hosted on the Cisco UCS C250 M1 server.

C460 M1 Positioning

The Cisco UCS C460 M1 server is designed with the performance and reliability to power compute-intensive, enterprise-critical standalone applications and virtualized workloads.

Supporting up to 512 GB of DDR3 memory in 64 DIMM slots, the Cisco UCS C460 M1 server, when combined with the Intel Xeon 7500 series processors, enables customers to manage the most demanding applications. Applications that are memory-bound today will benefit from the increased performance and memory, allowing a wider range of performance-intensive and enterprise-critical applications as well as increased virtual machine deployments and greater server consolidation.

From a memory-cost perspective, the server can be populated with low-cost 4-GB DIMMs, for a total of up to 256 GB of main memory, delivering exceptional value to Cisco customers.

Customers gain the benefits of the high-capacity memory of the Cisco UCS C460 M1 server when very large memory footprints are required, or when large, low-cost memory footprints are desirable, as in the following examples:

- Large virtualized environments
- Database applications
- Traditional HPC applications
- Enterprise resource planning (ERP) applications

Evaluating C-Series Memory

This topic identifies and describes C-Series memory options.

C-200 and C-210 Memory Installation

2-GB, 4-GB, 8-GB DIMM modules (1 Gb, 2 Gb DRAM modules)

800 MHz, 1066 MHz, 1333 MHz

ECC—uncorrectable or correctable

Reliability, availability and serviceability (RAS)

- Independent Channel Mode
- Mirroring (half of installed memory is used)

Nonuniform memory architecture (NUMA)

Mixed memory supported

- Recommend matching on channel

Speed

- Single rank, dual rank (one DIMM per channel—maximum 1333 MHz, two DIMMs per channel—maximum 1066 MHz)
- Quad rank DIMMs (one DIMM per channel—maximum 1066 MHz, two DIMMs per channel—maximum 800 MHz)
- Different maximum memory controller frequency, depends on processor SKU

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-11

Modern computer systems use DRAM technology to implement the main memory. On a computer system motherboard, there are main memory slots that allow system engineers or administrators to insert DIMM sticks. Each DIMM stick consists of a number of DRAM integrated circuit chips, called DRAM devices.

Modern microprocessors operate at a much higher frequency than the DRAM modules. Hence, processors communicate with memory subsystems via multiple channels to increase memory access bandwidth.

In the Cisco UCS C200 M1 and M2 and C210 M1 and M2 servers, each processor can be directly connected to three memory channels, but each channel will support up to two DIMMs.

C-200 and C-210 Memory Population Rules

Socket 1						Socket 2					
Channel A		Channel B		Channel C		Channel D		Channel E		Channel F	
A1	A2	B1	B2	C1	C2	D1	D2	E1	E2	F1	F2

- **Regular**
 - Processor socket 1—populate first
 - A1—populate first, color coded blue, “farthest first”
 - If socket 1 and socket 2 are installed, DIMMs only on socket 1 are OK
- **Mirroring**
 - DIMMs must be identical across CPUs
 - Minimum population in mirrored mode: socket 1: A1, B1
 - Upgrade—socket 1: A1, B1; socket 2: D1, E1
 - Upgrade—socket 1: A1, A2, B1, B2
 - Maximum—socket 1: A1, A2, B1, B2; socket 2: D1, D2, E1, E2

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-12

When a blade server has just one processor that is populated in the CPU0 socket, only the DIMM slots (A1, A2, B1, B2, C1, and C2) associated with CPU0 will be visible to the operating system (because the QPI link between CPU0 and CPU1 is broken). Cisco UCS C200 M1 and M2 and C210 M1 and M2 servers support up to two DIMMs per channel (DPC).

The actual running speed of DDR3 DIMMs can be different from the labeled speed. It depends on the CPU types, DIMM population patterns, and DIMM types.

In a Cisco UCS C200 M2 server with X5570 processors and identical DIMMs, all DIMMs will run at their labeled speed if there is one DIMM per channel, and will run at 1066 million transfers per second (MT/s) if there are two DIMMS per channel (except if using quad-rank DIMMs). In a UCS C200 M1 server with X5540, X5520, or L5520 processors with identical DIMMs, all DIMMs will run 1066 MT/s speed except when quad-rank DIMMs are used in a setup with two DIMMS per channel.

The maximum performance configuration is when six DDR3-1333 DIMMs are populated as 1 DPC (1-1-1/1-1-1 pattern) along with an Intel Xeon processor.

Performance Configuration Guidelines

There are many options to populate DIMM slots in UCS C200 M1 and M2 and C210 M1 and M2 servers. One should always assess the customer requirements, especially on the bandwidth and capacity, before purchasing DIMM and making configuration decisions.

Because of the three channel per direct-linked processor socket memory architecture, one should always try to perform “multiples of 6” balanced configuration, that is, populate all six channels with identical DIMMs.

The following guidelines should be observed:

- Use identical DIMM types throughout the system: same size, speed, and number of ranks.
- Populate the same DIMMs for each channel and each processor socket.
- Populate DIMMs to maximize the number of channels to achieve the highest degree of memory interleaving for the highest memory bandwidth.

If two DIMMs in one channel have different speeds, such as DDR3-1333 and DDR3-1066, that channel will operate at the lowest speed (1066 MT/s).

If the two CPUs are not X5570, there is no need to populate DDR3-1333 DIMMs.

Performance degradation will occur if any of the following occur:

- A processor socket does not have three channels fully populated (such as a 2-2-0/2-2-0 pattern where the third channel DIMM slots are empty).
- DIMMs with different sizes and densities are mixed in one memory bank.
- DIMMs are unevenly populated between two processor sockets.

Memory Mirroring

DIMMs in Cisco UCS C200 M1 and M2 and C210 M1 and M2 servers can be configured in mirroring mode for reliability reasons. All DIMMs must be identical in this configuration. Only two channels on each processor side can be populated for mirroring.

In the BIOS settings, the mirroring option will not be available when the third channel is populated (that is., in Independent Channel Mode). When identical DIMMs are populated in two channels for each CPU, for example, channels A1 and B1 for CPU0 and channels D1 and E1 for CPU1, the BIOS will have the memory mirroring option available.

If the BIOS finds nonidentical DIMMs populated in the intended mirroring channels, the memory mirroring option will not be available. The memory system will fall into the Independent Channel Mode.

In the mirroring mode, two identical images of memory channel data are maintained. Write transactions are issued to both memory channels. Read transactions are alternated between both channels. This data redundancy comes with a price: the total effective memory size will be reduced to half.

C-250 Memory Installation

48 DIMMs

- 4-GB, 8-GB DIMM modules
- ECC—uncorrectable or correctable

RAS

- Independent Channel Mode
- Mirroring (half of installed memory is used)

NUMA

Mixed memory supported

- Recommend matching on channel

Speed

- BIOS will limit the bus speed to 1066 MHz
- Dual rank DIMMs (quad rank not supported)
- Maximum bus frequency depends on processor SKU and DIMM capability

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-13

The Cisco Extended Memory Architecture employs a custom memory controller chipset that virtualizes DIMMs as seen by the CPU memory controller. With the technology, up to eight DIMMs can be installed per CPU memory channel without reducing the bus speed.

The Cisco UCS C250 M1 and M2 servers that are equipped with two sockets use a Cisco Extended Memory Technology that allows up to eight DIMMs per channel, resulting in 24 DIMMs per socket. In a two-socket server, this counts as 48 DIMMs, which results in 384 GB of memory per server if 8-GB DIMMS are used.

The benefit of large number of DIMMs is also per-server memory cost reduction. Note that with 48 DIMMs, 2-GB DIMMs can be used to deploy 96 GB of memory per server. The cost of a 2-GB DIMM compared to the 8-GB DIMM is much smaller.

C-250 Memory Population Rules

- CPU 1

Channel A								Channel B								Channel C							
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
- CPU 2

Channel D								Channel E								Channel F							
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
- Within one channel

DIMMS	0	1	6	7	3	2	5	4
2	x							x
4	x	x					x	x
8	x	x	x	x	x	x	x	x
- CPU 1—populate first
- A0, A4—populate first, “from outside in”
- If CPU 1 and CPU 2 are installed, DIMMs only on CPU 1 are OK

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-14

Like with the C200 M1 and M2 and C210 M1 and M2 servers, you must follow the recommended DIMM slot population scheme with C250 M1 and M2 servers in order to achieve the best performance.

C-250 Mirroring Population Rules

- Minimum population in mirrored mode:
 - CPU 1: A0, A4, B0, B4
- DIMMs must be identical across CPUs.
- Population must be symmetrical on channel A and channel B.
- If channel C or channel F are populated, BIOS will disable mirroring.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-15

DIMMs in Cisco UCS C250 M1 and M2 servers can also be configured in mirroring mode for reliability reasons. All DIMMs must be identical in this configuration.

At minimum, single CPU A0, A4, B0, and B4 DIMM banks have to be populated. When two CPUs are used, the same population scheme must be honored.

Like with C200 M1 and M2 and C210 M1 and M2 server, in the BIOS settings, the mirroring option will not be available with an improper population scheme. If the BIOS finds nonidentical DIMMs populated in the intended mirroring channels, the memory mirroring option will not be available. The memory system will fall into the Independent Channel Mode.


Again, in the mirroring mode, two identical images of memory channel data are maintained. Write transactions are issued to both memory channels. Read transactions are alternated between both channels. This data redundancy comes with a price: the total effective memory size will be reduced to half.

Evaluating C-Series LAN and SAN Connectivity

This topic identifies and describes C-Series LAN and SAN connectivity options.

Broadcom Ethernet Adapters

Adapter	Ports	Type	Features
Broadcom NetXtreme II 5709 PCIe (x4)	4	10/100/1000BASE-T Ethernet	<ul style="list-style-type: none">▪ TCP Offload▪ iSCSI boot
BCM95709A0907G PCIe (x4)	2	10/100/1000BASE-T Ethernet	<ul style="list-style-type: none">▪ TCP Offload▪ iSCSI boot
BCM957711A1113G PCIe (x8)	2	10GbE SFP+ (Optical modules vs. copper direct attach cable)	<ul style="list-style-type: none">▪ TCP Offload▪ iSCSI boot



Common features

- 802.1Q VLAN tagging (up to 64 VLANs)
- WOL

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-17

One of the benefits of rack-mount servers is the capability to configure a range of I/O options to meet specific workload requirements. The Cisco UCS C200 M1 server offers a range of flexible I/O options through its two PCIe expansion slots. Cisco supports adapters through arrangements with original equipment manufacturers (OEMs).

Cisco UCS C-Series servers are designed to operate both in standalone environments and as part of Cisco UCS. The series employs Cisco technology to help customers manage the most challenging workloads. The series incorporates a standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology. It supports an incremental deployment model and protects customer investments with an easy migration path to unified computing.

UCS C-Series Adapter Benefits

These are additional benefits:

- **Total cost of ownership:** Fewer network interface cards (NICs), HBAs, cables, and switches need to be powered, cooled, configured, and so on
- **Simplified operations:** Compatibility, flexibility, and unique virtualization support
- **Performance and availability:** 10 Gigabit Ethernet and up to 128 virtual interfaces on a single card

Cisco offers a choice of four types of PCIe adapters for use with Cisco UCS C-Series rack-mount servers, so that organizations can choose the technology most appropriate for their data centers and applications. The adapters can be mixed and matched in the same server.

Discrete I/O Adapters

Discrete I/O adapters further enhance customer flexibility and choice with Gigabit Ethernet, 10 Gigabit Ethernet, and 4 Gigabit Fibre Channel interfaces from industry-leading vendors including Broadcom, Emulex, and QLogic.

Industry-Standard Network Interface Cards

When using the Cisco UCS C-Series in a traditional computing environment, separate, fixed Ethernet and Fibre Channel adapters may be used, in keeping with existing architectural practices. There are several efficient, high-performance options, selected for compatibility with existing drivers of each type:

- Ethernet adapters:
 - Broadcom NetXtreme II 5709 Quad Port Ethernet PCIe Adapter Card with TOE and iSCSI boot
 - BCM95709A0907G Dual Port 1 Gigabit Ethernet PCIe Adapter Card with TOE and iSCSI boot
 - BCM957711A1113G Dual Port 10 Gigabit Ethernet PCIe Adapter Card with TOE and iSCSI boot

Intel Ethernet Adapters

Gigabit ET dual-port and quad-port PCIe (x4) adapter

- 10/100/1000BASE-T Ethernet

Gigabit EF dual-port PCIe (x4) adapter

- 1000BASE-SX Ethernet (62.5- or 50-micrometer MMF)

Common features

- LinkSec Layer 2 data protection
- PXE
- iSCSI boot
- 802.1Q VLAN tagging (up to 4096 VLANs)
- VMDq
- WOL
- IPv6 offload



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-18

The Intel Gigabit ET and EF Multi-Port Server Adapters are the third generation of PCIe Gigabit Ethernet network adapters. Built with the Intel 82576 Gigabit Ethernet Controller, these new adapters showcase the next evolution in Gigabit Ethernet networking features for the enterprise network and data center. These features include support for multicore processors and optimization for server virtualization.

The Intel Gigabit Ethernet adapters have the following common characteristics:

- Intel 82576 Gigabit Ethernet Controller
- Preboot Execution Environment (PXE) and iSCSI remote boot support
- Virtual Machine Device Queues (VMDq)
- Wake-on-LAN (WOL)
- IPv6 offload
- IEEE 802.1Q VLAN tagging with up to 4096 VLANs

Intel X520 Network Adapters

Common features

- 802.3ad link aggregation
- 802.1Q VLAN tagging (up to 4096 VLANs)
- VMDq
- IPv6 offload
- PXE and iSCSI boot
- FCoE in software
- LinkSec Layer 2 data protection

Adapter	Ports	Interface	Speed
X520-SR1	1	10GBASE-SR	1GbE/10GbE
X520-SR2	2	10GBASE-SR	1GbE/10GbE
X520-LR1	1	10GBASE-LR	1GbE/10GbE
X520-DA1	2	10GSFP+ Cu (Direct Attach)	10GbE



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-19

Intel Ethernet X520 Server Adapters are flexible and scalable Ethernet adapters for demanding data center environments. Data center networks are expanding rapidly. The escalating deployments of servers with multicore processors and demanding applications such as HPC, database clusters, and video-on-demand are increasing the need for 10 Gigabit connections.

Customers require flexible and scalable I/O solutions to meet the rigorous requirements of running mission-critical applications in virtualized and unified storage environments. Powered by Intel's third-generation 10 Gigabit Ethernet network controller, the Intel Ethernet 82599 10 Gigabit Ethernet Controller, the X520 server adapter family addresses the demanding needs of the next-generation data center by providing unmatched features for virtualization, flexibility for LAN and SAN networking, and proven, reliable performance.

The family of Intel Ethernet X520 server adapters lowers data center total cost of ownership by providing the ability to route LAN and SAN traffic over a single fabric.

The adapters have the following common characteristics:

- IEEE 802.3ad link aggregation
- 802.1Q VLAN tagging with up to 4096 VLANs
- PXE and iSCSI remote boot support
- VMDq
- LinkSec Layer 2 data protection


Emulex Adapters

LightPulse LPe11000/LPe11002 4G Fibre Channel HBA

- Single vs. dual port 62.5- or 50-micrometer MMF
- NPIV
- FC-SP

OneConnect OCe10102-FX-C CNA

- Dual-port 10GbE with FCoE
- FCoE and TCP offload
- PXE and SAN boot
- 802.1Q VLAN tagging
- 802.3ad link aggregation
- NPIV
- SFP+ Twinax direct attached copper cable or 10GBASE-SR LC fiber optics



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-20

LightPulse 4G Fibre Channel HBA Adapters

A single-channel Emulex LightPulse LPe11000 and dual-channel LPe11002 are ideal solutions for enterprise and mixed-operating system SAN environments. With powerful management tools and broad platform support, they deliver maximum performance in the broadest range of applications and environments.

The LightPulse HBA highly integrated single-chip design minimizes onboard components, while advanced error-checking and correcting methods assure robust data integrity. The Emulex firmware-based architecture enables feature and performance upgrades without costly hardware changes.

OneConnect 10 Gigabit Ethernet CNA Adapter

CNA from Emulex presents both Ethernet NICs and Fibre Channel HBAs to the host operating system, consolidating traffic over a 10 Gigabit unified fabric.

OCe10102-FX-C is a high-performance 10 Gigabit FCoE adapter that consolidates traffic for networking, Fibre Channel, and FCoE storage for Cisco UCS C-Series rack-mount servers.

The OCe10102-FX-C FCoE adapter supports a common 10 Gigabit Ethernet infrastructure for unified networking and storage, extending the Cisco standards-based unified network fabric.

Fibre Channel SAN and PXE boot support make the OCe10102-FX-C an ideal solution for Cisco UCS C-Series rack-mount servers.

QLogic Adapters

QLE2462 4G Fibre Channel HBA

- Dual port 62.5- or 50-micrometer MMF
- NPIV
- VSAN support

QLE8152 CNA

- Dual-port 10 Gigabit Ethernet with FCoE
- TCP offload
- PXE and SAN boot
- 802.1Q VLAN tagging
- 802.3ad link aggregation
- NPIV
- SFP+ Twinax direct attached copper cable or 10GBASE-SR LC fiber optics



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-21

By using standard adapter vendor ASICs to encapsulate multiple traffic streams on the unified network fabric, these CNAs offer complete compatibility with existing data center best practices that are based on the use of Emulex or QLogic HBAs. This approach helps increase compatibility with target storage systems and may reduce the effort that is needed for IT departments to qualify the CNAs.

QLE2462 4G Fibre Channel HBA

The QLE2462 is the enterprise class, 4 Gb/s-to-PCI Express x4 adapter. The QLE2462 delivers unprecedented levels of performance and availability, as well as intelligent networking features specific to enterprise class data centers.

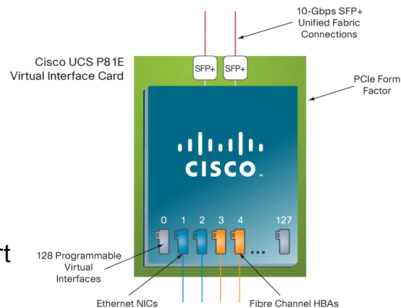
QLE8152 CNA

CNA from QLogic presents both Ethernet NICs and Fibre Channel HBAs to the host operating system, consolidating traffic over a 10 Gigabit unified fabric.

FCoE provides an opportunity to reduce data center costs by converging data and storage networking. Standard TCP/IP and Fibre Channel traffic can both run on the same high-speed 10 Gigabit Ethernet wire, resulting in cost savings through reduced adapter, switch, cabling, power, cooling, and management requirements. The QLE8152 even supports iSCSI storage protocol using iSCSI software initiators, which are available with all major operating systems. The QLE8152 boosts system performance with 10 Gigabit speed and full hardware offload for FCoE protocol processing.

Cisco UCS P81E Virtual Interface Cards

- **Dual-port 10 Gigabit Ethernet PCIe adapter**
 - Unified I/O—virtualization-optimized FCoE PCIe 2.0 x8 10 Gb adapter
 - Lossless Ethernet support
 - Cisco VN-Link virtualization support
 - Hypervisor bypass
- **Up to 2 FC HBAs and 16 Ethernet NICs**
 - Dynamically configured NIC and/or HBA via Cisco IMC
 - Virtual interfaces appear as regular PCIe devices
 - Up to 128 PCIe standards-compliant virtual interfaces (future)
- **Up to 100 adapters** when integrated with Cisco UCS Manager (1.4 software)



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-3-22

The Cisco UCS P81E Virtual Interface Card delivers the full power of Cisco UCS by providing up to 128 Ethernet or Fibre Channel virtual interfaces that are programmed on demand to meet the needs of virtualized and nonvirtualized environments alike. The dual-port card interfaces with a 10 Gigabit unified fabric.

The Cisco UCS P81E can present up to 128 virtual interfaces on a given server. The 128 virtual interfaces can be dynamically configured by Cisco UCS Manager as either Fibre Channel or Ethernet devices. Initially, the Cisco UCS P81E supports up to 2 Fibre Channel and 16 Ethernet devices.

To an operating system or a hypervisor running on a Cisco UCS C-Series rack-mount server, the virtual interfaces appear as regular PCIe devices. In a virtualized environment, Cisco VN-Link technology allows virtual links to be centrally configured and managed without the complexity that traditional approaches interpose with multiple switching layers in virtualized environments. I/O configurations and network profiles move along with virtual machines, helping increase security and efficiency while reducing complexity. As a result of close cooperation between Cisco and VMware, network policies and virtual interfaces can be applied to virtual machines in VMware vCenter. The partnership also enables pass-through switching in the virtual switch, improving hypervisor performance.

Cisco UCS Manager and C-Series Integration

C-Series servers connectivity


- Management connectivity through FEX to 6100
- Data connectivity directly to 6100 via UCS P81eVIC
 - Fabric based failover for vNICs

Stateless computing

- Service profiles extended to C-Series
- Migration among compatible B-Series and C-Series servers

All Cisco UCS Manager services

- Automated discovery
- Fault and monitoring
- KVM access
- Firmware updates



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-23

With version 1.4 of the Cisco UCS Manager, the C-Series can be integrated with both the Cisco UCS Manager and the B-Series system.

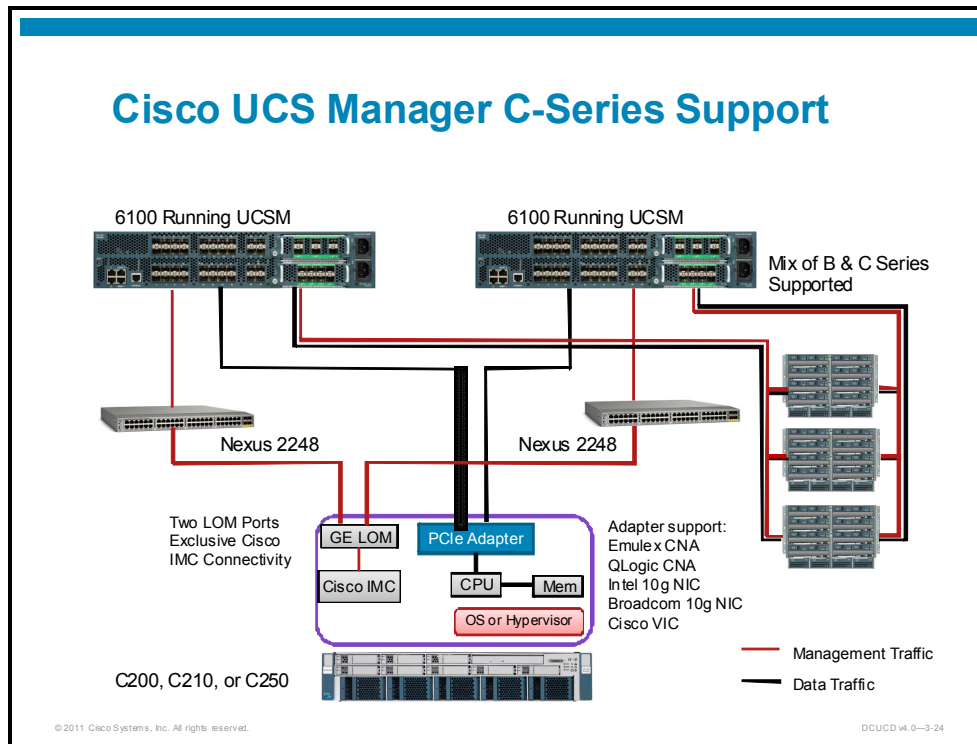
Cisco UCS C-Series servers can also be managed centrally by Cisco UCS Manager, as part of Cisco UCS. Cisco UCS Manager serves as the central nervous system of Cisco UCS, managing the system from end to end, across multiple chassis, as a single logical entity through an intuitive GUI, a CLI, or an XML API. Cisco UCS blade and rack-mount servers can be brought into a common Cisco UCS Manager domain, allowing IT managers to take advantage of the Cisco UCS unique operational efficiency and agility regardless of form factor.

Cisco UCS Manager implements role- and policy-based management using service profiles and templates. This construct improves IT productivity and business agility.

Using Cisco UCS Manager, IT infrastructure can be provisioned in minutes instead of days, shifting IT focus from maintenance to strategic initiatives.

Integrating the Cisco UCS C-Series with Cisco UCS Manager provides all of the functionality and benefits that integrated B-Series blades already possess.

Cisco UCS Manager C-Series Support



The 1.4 version of the software supports integration of the following C-Series hardware:

- C-Series rack servers C200, C210, C250
- Adapters:
 - Cisco UCS P81E Virtual Interface Card
 - Emulex OneConnect Universal Converged Network Adapter
 - QLogic QLE8152 Dual Port 10-Gigabit Ethernet to PCIe Converged Network Adapter
 - Broadcom NetXtreme II 57711 Dual Port 10-Gigabit Ethernet PCIe Adapter Card
 - Intel 82599 (Niantic) Dual port 10-Gigabit Ethernet Adapter

The management ports of the C-Series have to be connected to the Nexus 2248 FEX connected to the Cisco UCS 6100XP Fabric Interconnects—this way, the C-Series can be managed via Cisco UCS Manager.

For the data traffic, the P810 10 Gigabit Ethernet interface is connected directly to the Cisco UCS 6100 Fabric Interconnects.

The number of C-Series servers that are supported by Cisco UCS Manager depends on the maximum number of available server ports in the fabric interconnect.

Supported configurations are C-Series only, and mixed B-Series and C-Series configurations.

Evaluating C-Series Local Storage

This topic identifies and describes C-Series RAID and disk options.

HDD Controller Option—Entry-Level

- Lowest cost
 - Integrated ICH based SATA controller
 - RAID 0 and 1
 - Four-drive maximum
- Low cost SAS and SATA
 - LSI 1068 controller-based mezzanine card
 - RAID 0 and 1
 - Maximum of eight SAS/SATA ports



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-28

C200 M1 and M2 Servers

RAID 0 and 1 support for up to four SATA drives using the built-in RAID controller; RAID 0 and 1 support for up to four SAS or SATA drives with the optional LSI 1064 Controller-Based Mezzanine Card; and RAID.

C210 M1 and M2 Servers

RAID 0 and 1 support for up to four SATA drives using the built-in RAID controller; RAID 0 and 1 support for up to four SAS or SATA drives with the optional LSI 1064 controller-based mezzanine card (four ports).

HDD Controller Options—Performance

- **2U**
 - LSI 9261-8i MegaRAID Controller **PCIe add-on card**
 - Gen 2, 6-Gb SAS controller
 - 512 MB write cache with battery backup
 - RAID 0, 1, 5, 6, 10, 50, 60
 - SAS expander on 2U HDD backplane
- **1U**
 - LSI 8708EM2 MegaRAID Controller **PCIe add-on card**
 - Gen 1, 3-Gb SAS controller
 - 256 MB write cache with battery backup
 - RAID 0, 1, 5, 6, 10, 50, 60

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-27

C200 M1 and M2 Servers

RAID 0, 1, 5, 6, and 10 support through an optional LSI MegaRAID SAS 8708EM2 PCIe RAID Controller that features 256 MB of write cache and battery backup.

C210 M1 and M2 Servers

RAID controllers for RAID 0, 1, 5, 6, 10, 50, and 60 support through a PCIe and mezzanine card form factors.

C250 M1 and M2 Servers

RAID 0 and 1 support for up to eight SAS or SATA drives with the optional LSI SAS30813E-R PCIe RAID Controller; and RAID 0, 1, 5, 6, 50, and 60 support for up to eight SAS or SATA drives with the optional LSI MegaRAID Controller

You can choose to support two RAID controllers through a PCIe and mezzanine card form factors.

One of the five slots on C250 M1 and M2 servers is available to configure RAID support through the optional LSI MegaRAID controller.

C460 M1 Server

RAID 0, 1, and 10 support for up to 12 SAS or SATA drives with the optional LSI SAS9240-8i PCIe RAID controller; and RAID 0, 1, 5, 6, 10, 50, and 60 support for up to 12 SAS or SATA drives with the optional LSI MegaRAID controller.

The 11th slot on C460M1 server available to configure RAID support through the optional LSI MegaRAID controller.

Supported HDDs

- 3.5-inch HDDs supported on C-200
 - 146-GB 10,000 rpm SAS
 - 250-GB 7200 rpm SATA
 - 300-GB 10,000 or 15,000 rpm SAS
 - 500-GB 7200 rpm SATA
 - 450-GB 15,000 rpm SAS
 - 1-TB 7200 rpm SAS
 - 2-TB 7200 rpm SAS
- 2.5-inch HDDs supported on C-210, C-250, and C-460
 - 73-GB 15,000 rpm SAS
 - 146-GB 10,000 rpm SAS
 - 300-GB 15,000 rpm SAS
 - 500-GB 7200 rpm SATA
 - 100-GB SSD SATA (C250 only)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-28

C200 M1 Server

There are up to four front-accessible, hot-swappable, 3.5-inch SAS or SATA drives:

- 250-GB SATA; 7200 rpm
- 500-GB SATA; 7200 rpm
- 146-GB SAS; 15,000 rpm
- 300-GB SAS; 15,000 rpm
- 1-TB SAS; 7200 rpm

C200 M2 Server

There are up to four front-accessible, hot-swappable, 3.5-inch SAS or SATA drives:

- Gen 2 500-GB SATA; 7200 rpm
- Gen 2 1-TB SAS; 7200 rpm
- 2-TB SAS; 7200 RPM, 3.5-inch HDD
- Gen 2 300-GB SAS; 15,000 rpm
- 450-GB SAS; 15,000 rpm

C210 M1 and M2 Servers

There are up to 16 front-accessible, hot-swappable, 2.5-inch SAS or SATA drives:

- 73-GB SAS; 6G, 15,000 rpm
- 146-GB SAS; 6G, 10,000 rpm
- 300-GB SAS; 6G, 10,000 rpm
- 500-GB SATA; 7200 rpm

C250 M1 and M2 Servers

There are up to eight front-accessible, hot-swappable, 2.5-inch SAS or SATA drives:

- 73-GB SAS; 6G, 15,000 rpm
- 146-GB SAS; 6G, 10,000 rpm
- 300-GB SAS; 6G, 10,000 rpm
- 500-GB SATA; 7200 rpm

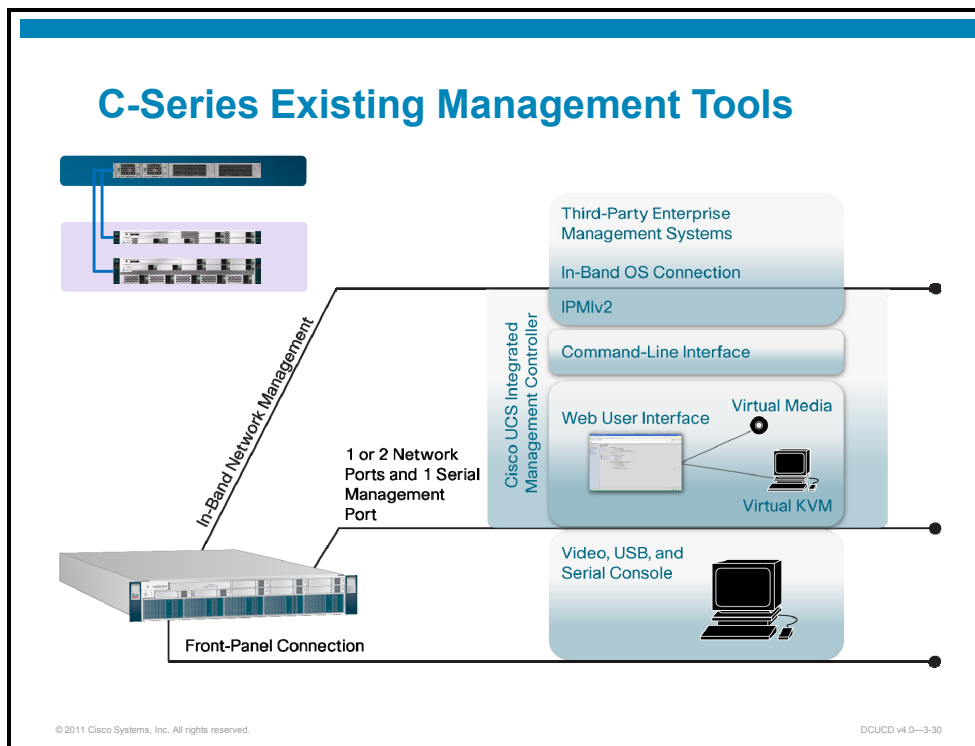
C460 M1 Server

There are up to 12 front-accessible, hot-swappable, 2.5-inch SAS or SATA drives:

- 73-GB SAS; 6G, 15,000 rpm
- 146-GB SAS; 6G, 10,000 rpm
- 300-GB SAS; 6G, 10,000 rpm
- 500-GB SATA; 7200 rpm

Evaluating C-Series Management

This topic identifies and describes Cisco IMC.



Cisco UCS Integrated Management Controller

The Cisco UCS IMC runs in the BMC of the system, and can be accessed through the network management ports of the server. It provides out-of-band management that can be accessed through standard management protocols, CLIs, and web-based interfaces.

The Cisco IMC is the management service for the Cisco UCS C-Series rack-mount server. Cisco IMC runs within the server.

IPMI v2.0

IPMI v2.0 supports out-of-band management through third-party tools including commercial enterprise management systems and open-source tools such as IPMITool. IPMI allows these tools to manage server power states and monitor operational parameters available through temperature, fan-speed, power-supply voltage, and power sensors. Through IPMI, platform event traps can be configured to send standard SNMP v1.0 traps to the management tools.

CLI

The CLI of Cisco UCS IMC can be accessed through an SSH connection to the Cisco UCS IMC. Through this interface, administrators can perform server control and administration tasks.

Enterprise Management Tools

Third-party management tools typically use a combination of in-band and out-of-band management techniques, both of which are supported by Cisco UCS C-Series servers.

- **In-band management** is performed through the data network connection of the server. Different tools use different techniques, including interaction with the host operating system with and without the use of agents. In-band management can interact with operating system-based management tools to accomplish tasks including inventory, performance management, troubleshooting, and operating system and interface provisioning.
- **Out-of-band management** uses features of the Cisco UCS IMC through the serial or network management port. Enterprise management tools typically interact with servers through IPMI and the platform event traps, which appear as SNMP v1.0 traps.

Cisco IMC Software Functions

- Server management and monitoring with IPMI v2.0
 - KCS and IPMI over LAN supported
- Server management and monitoring with WebGUI and CLI
 - WebGUI access over dedicated Cisco IMC Ethernet port or shared LOM
 - CLI access using SSH to Cisco IMC (Cisco UCS syntax, scope, commit)
 - x86 host power control, x86 boot order control, analog sensor monitoring, rack server inventory, and state monitoring
 - Remote KVM
 - Virtual media
 - COM1 serial over support using SSH (via CLI)
 - LED control
 - Cisco IMC firmware update
- Integration with Cisco UCS Manager—requires 1.4 software

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-31

The web user interface supports out-of-band management through a standard web browser. It includes server management, remote KVM, virtual media, and administration capabilities:

- **Server management** includes power management, server reset, component inventory, and event logging.
- **Virtual media** enables an administrator laptop or desktop workstation peripherals such as CD and DVD drives to appear as if they were connected directly to the server, facilitating remote operating system and application software installation.
- **Remote KVM** gives remote administrators the same level of control, including console video, as if they were connected directly to the front-panel USB connection of the server.
- **Administration features** include management of role-based access on a per-user basis, integration with external authentication and authorization systems, certificate management, firmware updating, network configuration, and specification of events to trigger platform event traps.

KVM Features

- Server provides a physical local KVM connection through the front panel of the rack server
 - All C-Series servers can have up to four active KVM over IP sessions in addition to the local connection at front panel.
 - All active sessions have full control of the console.
 - KVM over IP supports text and graphics modes of the graphics controller and needs no manual setting to view data.
 - Graphics modes are supported up to the maximum capabilities of the chip, 1600 x 1200 with 16-bit color.
 - KVM sessions over Linux may, depending on the Linux distribution, require mouse configuration.
- vMedia to connect remote device to the server as if attached to local USB interface
 - Three supported devices—floppy (read/write), hard drive or USB key (read/write), CD ROM drive (read-only)
 - Can use ISO images to attach as CD-ROMs

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-32

Cisco IMC is a separate management module that is built into the motherboard. It includes its own ARM-based processor that runs the Cisco IMC software. The Cisco IMC is shipped with a running version of the firmware. Users can update the Cisco IMC firmware through the Firmware Update Management page. You do not need to install the initial firmware for the Cisco IMC.

You can use the Cisco IMC to perform these server management tasks:

- Power on, power off, power cycle, reset, and shut down the server
- Toggle the locator LED
- View server properties and sensors
- Manage remote presence
- Create and manage local user accounts, and enable remote user authentication through Active Directory
- Configure network-related settings (including NIC properties, IPv4, VLANs, and network security)
- Configure communication services (including HTTP, SSH, and IPMI Over LAN)
- Manage certificates
- Configure platform event filters
- Update Cisco IMC firmware
- Monitor faults, alarms, and server status

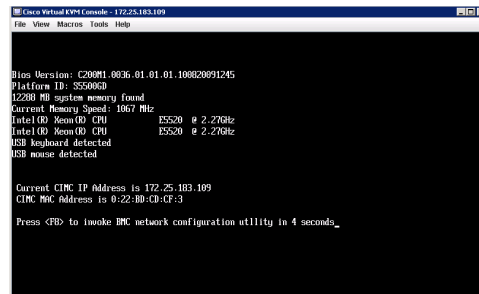
Out-of-the-Box Configuration

Single administrator user

- Username: admin
- Password: password

Initial setup network port configuration

- Cisco IMC dedicated Ethernet port (Cisco IMC eth1) uses DHCP
- Cisco IMC network settings can be overridden from x86 host EFI shell



```

Cisco Virtual BIOS Console - 172.25.183.189
File View Macros Tools Help

BIOS Version: C200M1.0636.01.01.01.10882091245
Platform ID: S250000
12288 MB system memory found
Current Memory Speed: 1067 MHz
Intel(R) Atom(TM) CPU          E5200  @ 2.270GHz
Intel(R) Atom(TM) CPU          E5200  @ 2.270GHz
USB keyboard detected
USB mouse detected

Current CIMC IP Address is 172.25.183.189
CIMC BMC Address is 0:22:80:00:00:03

Press <PB> to invoke BMC network configuration utility in 4 seconds_
```

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-33

The initial out-of-the-box configuration of Cisco C-Series servers (Cisco IMC management) is as follows:

- Single administrator user
 - Username: admin
 - Password: password
- Network port configuration—Cisco IMC dedicated Ethernet port (Cisco IMC eth1) uses DHCP

Cisco IMC network settings can be overridden from the x86 host EFI shell.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco UCS C-Series servers are standalone rack-mount servers with different hardware characteristics.
- To achieve optimal memory configuration, the DIMMs need to be properly populated.
- The Cisco UCS C-Series has LOM, but also supports a variety of Ethernet, Fibre Channel, CNA, and Cisco VIC adapters.
- Local disk drives can be configured for RAID 0,1,10, 5, 6,10, 50, or 60 depending on the RAID controller and the Cisco UCS C-Series type.
- Cisco UCS C-Series servers can be individually managed with Cisco IMC.

Lesson 2

Sizing the Cisco UCS C-Series Solution

Overview

This lesson identifies how to assemble a Cisco UCS C-Series solution for a given set of requirements.

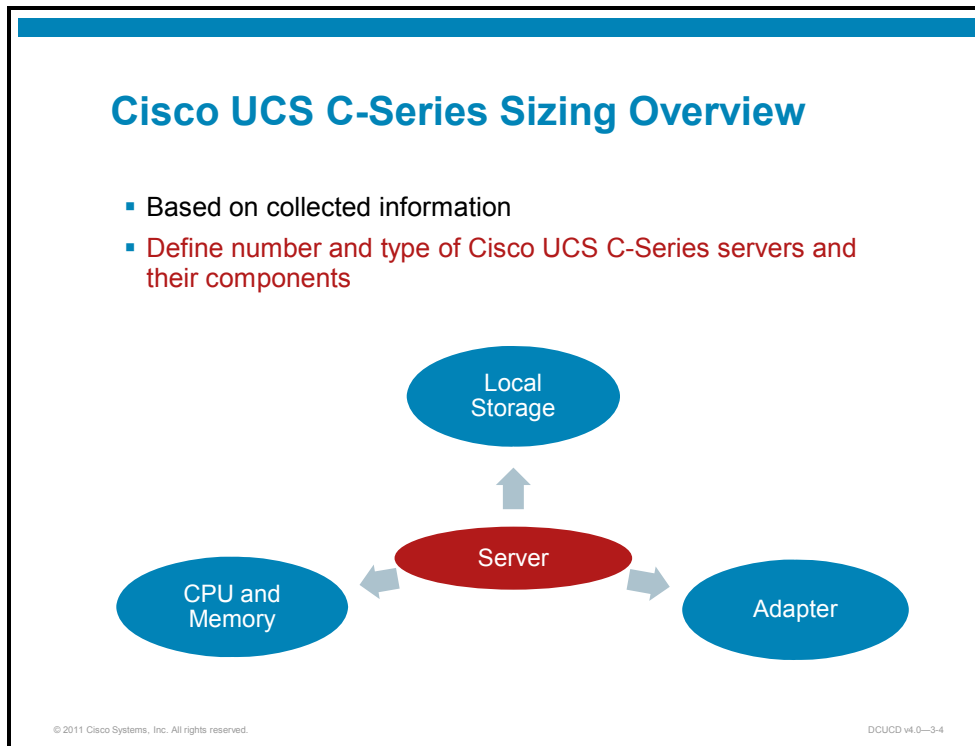
Objectives

Upon completing this lesson, you will be able to assemble Cisco UCS C-Series hardware components and a Bill of Materials (BOM). This includes the ability to meet these objectives:

- Identify and describe solution requirements
- Identify how to select C-Series hardware
- Identify and describe how to create a C-Series solution BOM

Analyzing Requirements

This topic identifies and describes solution requirements.



The sizing process for the Cisco UCS C-Series is used to determine the appropriate components that are needed to build the Cisco UCS per the requirements that were gathered in a design workshop.

The sizing process can be divided into two major categories:

- **Identify Cisco UCS C-Series server types:** With this step, the individual server type is identified and quantities are defined.
- **Size Cisco UCS C-Series servers:** With this step, the components of individual C-Series servers are selected.

Analysis Output

Basis for solution sizing and server deployment

Parameter	Server Type 1	Server Type 2	Server Type 3
CPU	1x 2-core at 1.2 GHz	1x 4-core at 2.4 GHz	2x 4-core at 2.4 GHz
Memory	8 GB	16 GB	24 GB
LAN Connectivity	2x 1 Gigabit Ethernet	6x 1 Gigabit Ethernet	6x 1 Gigabit Ethernet
SAN Connectivity	n/a	2x 2 Gb Fibre Channel	2x 4 Gb Fibre Channel
Boot Media	Local disk	Local disk	SAN
LAN Throughput	0.5 Gb/s	0.8 Gb/s	1.5 Gb/s
SAN Throughput	n/a	1 Gb/s	2 Gb/s

System Name	Make/Model	Configuration										Performance									
		Processors		Memory		Disk		Network		Physical		Processor		Memory		Disk		Netw			
		Count	Speed (MHz)	Size (MB)	Size (GB)	Count	Speed (MHz/sec)	Rack Units	Weight (lbs)	Power (W)	Thermal (BTU/hr)	% Used	Queue per CPU	% Used	Cache (MB)	Page File % (Pg/sec)	Paging (Trans/sec)	I/O (MB/sec)	I/O (MB)	S	
Installable Systems																					
UCS-B1000	UCS-B1000-01000-001	2	2,790	2,560	145.60	2	2,000	2	47.18	400	1,475	1.51	0.00	33.07	233.40	0.29	22.02	10.69	0.25		
UCS-B1000	UCS-B1000-01000-001	2	2,786	4,600	145.60	2	2,000	2	47.18	400	1,475	3.28	0.00	36.24	214.01	0.31	88.66	27.32	0.49		
UCS-B1000	UCS-B1000-01000-001	2	2,784	4,600	145.60	2	2,000	2	47.18	400	1,475	3.66	0.01	36.56	196.41	0.31	75.51	25.51	0.38		
UCS-B1000	UCS-B1000-01000-001	2	2,790	2,560	145.60	2	2,000	2	47.18	400	1,475	3.55	0.01	31.40	178.26	0.38	31.87	17.79	0.52		
UCS-B1000	UCS-B1000-01000-001	2	2,786	2,560	145.60	2	2,000	2	47.18	400	1,475	7.03	0.01	34.06	304.23	0.40	654.75	171.79	3.02		
UCS-B1000	UCS-B1000-01000-001	2	2,785	4,600	145.60	2	2,000	2	47.18	400	1,475	1.72	0.00	40.58	271.07	0.31	5.40	8.52	0.09		
UCS-B1000	UCS-B1000-01000-001	2	2,785	4,600	145.60	2	2,000	2	47.18	400	1,475	10.85	0.01	54.62	185.37	0.50	102.79	100.84	0.07		

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-5

When designing a Cisco UCS C-Series solution for an existing environment, it is vital to examine the analysis output to determine proper requirements for the processing, memory, connectivity, and storage requirements.


Sizing Cisco UCS C-Series— Connectivity

Connectivity aspects

- LAN—1 Gigabit Ethernet vs. 10 Gigabit Ethernet
- OOB management for Cisco IMC
- SAN
 - Fibre Channel with HBA
 - iSCSI with TOE/iSCSI offload
 - FCoE with CNA or VIC
- UCS P81E VIC
 - Max 2 HBAs, 16 NICs

Design the following parameters

- Throughput and oversubscription
- Redundancy—multifabric design



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0–3-6

The UCS C-Series sizing is primarily influenced by two factors—the required throughput and allowed oversubscription.

Both factors should be observed from different perspectives:

- LAN connectivity
- SAN connectivity

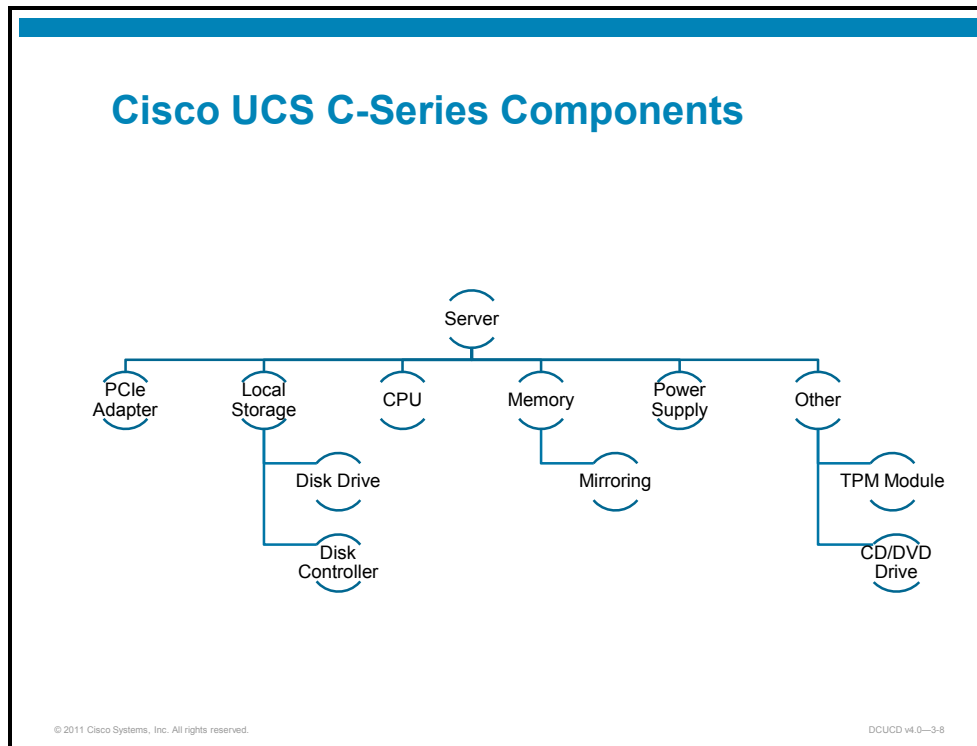
The Cisco UCS C-Series connectivity architecture follows multifabric design. Such a design is used to achieve high availability (with failover on the Cisco UCS level with P81E VIC or on the operating system level for other adapters), to achieve more throughput, or to achieve a combination of redundancy and higher throughput.

When determining the number of logical adapters that need to be presented to the operating system or hypervisor (for example, VMware ESX/ESXi or Microsoft Hyper-V host), the Cisco UCS P81E VIC adapter can be selected for consolidated LAN and SAN connectivity when up to two HBAs and up to 16 NICs are required. When connectivity consolidation is not required, the server can also be equipped with a multiport Ethernet adapter and a multiport Fibre Channel host bus adapter (HBA).

The maximum number of physical adapters is governed by the number of PCI Express (PCIe) slots on the server motherboard.

Sizing the C-Series Solution

This topic identifies how to suggest C-Series hardware.



When sizing the Cisco UCS C-Series server, the designer must select the following:

- Type and quantity of CPUs
- Type, quantity, and population scheme for memory
 - Mirroring if data coherency is required
- Local storage components
 - RAID controller (built-in or additional)
 - Disk drives
- PCIe adapters for LAN and SAN connectivity
 - Ethernet adapters for LAN connectivity
 - Fibre Channel HBAs for SAN connectivity
 - Ethernet adapters with TCP/IP Offload Engine (TOE) or Internet Small Computer Systems Interface (iSCSI) offload for iSCSI-based SAN
 - Converged network adapter (CNA) for consolidated LAN and SAN environment
 - Virtualized adapter for consolidated LAN and SAN and multiple logical adapters
- Power supplies
- Optional elements like Trusted Platform Module (TPM), operating system, and so on

Select Base Server



UCS C250 M2
Extended Memory Rack Server
PID: R250-2480805W



UCS C210 M2
General-Purpose Rack Server
PID: R210-2121605W



UCS C200 M2
High-Density Rack-Mount Server
PID: R200-112040W

Server Type	Target Workload
Performance optimized	High-performance, memory-intensive server for virtualized and large data-set workloads
Economical, high capacity	General-purpose server for workloads requiring economical, high-capacity, internal storage
High density	High-density server with balanced compute performance and I/O flexibility

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-9

First, the Cisco UCS C-Series requirements are identified. The selection depends on the target workload since different C-Series servers are tailored for different workloads:

- Cisco UCS C250 M1 and M2 are performance optimized and positioned as high-performance, memory-intensive servers for virtualized and large data-set workloads.
- Cisco UCS 210 M1 and M2 are economical, high capacity servers that are positioned as general-purpose servers for workloads requiring economical, high-capacity, internal storage.
- Cisco UCS C200 M1 and M2 are positioned as high-density servers with balanced compute performance and I/O flexibility.

Configure Selected Servers—Step 1

- Select CPU type (minimum 1, maximum 4 depending on server type)

CPU Description	PID	Notes
Intel Xeon X5680 Six Core Processor (3.33GHz CPU/12MB cache, 1333MHz)	A01-X0100	C250 support only
Intel Xeon X5670 Six Core Processor (2.93GHz, 95W CPU/12MB cache, 1333MHz)	A01-X0102	
Intel Xeon X5650 Six Core Processor (2.66GHz, 95W CPU/12MB cache, 1066MHz)	A01-X0105	
Intel Xeon E5640 Six Core Processor (2.66GHz, 80W CPU/12MB cache, 1066MHz)	A01-X0109	
Intel Xeon E5620 Six Core Processor (2.40GHz, 80W CPU/12MB cache, 1066MHz)	A01-X0111	
Intel Xeon L5640 Six Core Processor (2.26GHz, 60W CPU/12MB cache, 1066MHz)	A01-X0106	C250 not supported
Intel Xeon E5506 Four Core Processor (2.13GHz, 95W CPU/4MB cache, 800MHz)	A01-X0113	C250 not supported

- C210 M1 and M2 servers optional
 - Select DVD-RW
 - Select SAS expander (default support up to eight drives)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-10

Once the basic server types are selected, the designer can select the CPU or CPUs for individual servers. The selection of the CPU depends on the computing power requirements and should be based on analysis results as well as power consumption and number of cores.

Configure Selected Servers—Step 2

- Select memory
 - C200 and C210—minimum 1, maximum 12 DIMMs (6 per CPU slot)
 - C250—minimum 1, maximum 24 (24 per CPU slot)

C200 M2 and C210 C2 Memory	PID
2-GB DDR3-1333MHz RDIMM/PC3-10600	N01-M302GB1
4-GB DDR3-1333MHz RDIMM/PC3-10600	N01-M304GB1
8-GB DDR3-1333MHz RDIMM/PC3-10600	N01-M308GB2
2-GB DDR3-1333MHz RDIMM LV	N01-M302GB1-L
4-GB DDR3-1333MHz RDIMM LV	N01-M304GB1-L
8-GB DDR3-1333MHz RDIMM LV	N01-M308GB2-L
1-GB DDR3-1066MHz Unbuffered DIMM/PC3-10600/single rank 1-Gb DRAMs	A02-U301GB1
2-GB DDR3-1066MHz Unbuffered DIMM/PC3-10600/dual rank 1-Gb DRAMs	A02-U302GB1

C250 M2 Memory	PID
8-GB DDR3-1333MHz RDIMM/PC3-10600/2x4GB Kit	A02-M308GB1-2
16-GB DDR3-1333MHz RDIMM/PC3-10600/2x8GB Kit	A02-M316GB1-2
8-GB DDR3-1333MHz RDIM/PC3-10600/2x4GB Kit – LV	A02-M308GB1-2-L
16-GB DDR3-1333MHz RDIM/PC3-10600/2x8GB Kit – LV	A02-M316GB1-2-L

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-11

Equally important is the selection of memory—the total capacity governs the size of individual DIMMs. The selection of specific DIMMs is influenced by the power consumption and memory access speed—this is also dependent on the DIMM slots population scheme.

Configure Selected Servers—Step 3

- Select RAID controller (optional)
 - C200 and C210 default SATA RAID 0/1 controller (integrated on motherboard)
 - C250 no default option, must be ordered

Support	RAID Controller Option	PID	Notes
C200 C210	LSI 1064E Controller-Based Mezzanine Card	R2X0-ML002	- Takes up one Mezzanine slot - Supports up to 4 SAS or SATA hard drives, RAID 0, 1, 1E
C200	LSI MegaRAID SAS 8708EM2 PCIe card	R2XX-PL002	- Supports up to 8 SAS or SATA drives, RAID 0, 1, 5, 6, 10, 50, 60 - Takes up 1 PCIe slot - Includes 256MB of Write Cache - Battery Back-Up Option Available, R2X-BATT-BU
C200	LSI 6G MegaRAID 9260-4i PCIe card	R200-PL004	
C200	Battery Back-up for 9260-4i PCIe card	R2XX-LBBU3	
C210 C250	LSI 6G MegaRAID 9261-8 PCIe card	R2XX-PL003	- Takes up 1 PCIe slot - Supports up to 16 SAS or SATA drives, RAID 0, 1, 5, 6, 10, 60 - Includes 512MB of Write Cache - Battery Back-Up Option Available, :R2XX-LBBU2
C250	LSI SAS3081E-R PCIe Card (PID)	:R250-PL003	- Takes up 1 PCIe Slot, - Supports up to 4 SAS or SATA hard drives, RAID 0, 1, 1E, 10

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-12

Rack servers are more often used in environments demanding local storage. Here, the designer has to select two things:

- RAID controller, which governs the RAID level and disk performance
- Disk drivers—type based on capacity, EPM, and media (solid-state drive versus common spindle-based).

The speed of accessing and writing the data to the disk drives that are local to the server depends on the RAID level (which governs number of concurrent disks accessed), disk cache size, rpm, and disk size.

The redundancy is also important and proper RAID levels can help with redundancy, which again is dependent on the RAID controller.

To achieve the proper redundancy level and maximum performance, RAID 10 is typically used.

Configure Selected Servers—Step 4

- Select disk drives
 - 3.5-inch disk drives for C200 (minimum 0, maximum 4)
 - 2.5-inch disk drives for C210, C250, C460 (minimum 0, maximum 16/8/12)

3.5" Drives		PID
250-GB SATA 7.2K RPM 3.5" HDD/hot plug/C200 drive sled		R200-D250GCSATA
500-GB SATA 7.2K RPM 3.5" HDD/hot plug/C200 drive sled		R200-D500GCSATA
146-GB SAS 15K RPM 3.5" HDD/hot plug/C200 drive sled		R200-D146GB
300-GB SAS 15K RPM 3.5" HDD/hot plug/C200 drive sled		R200-D300GB
1-TB SAS 7.2K RPM 3.5" HDD/hot plug/C200 drive sled		R200-D1TC
Gen 2 500-GB SATA 7.2K rpm 3.5" HDD/hot plug/C200		R200-D500GCSATA03
Gen 2 1-TB SAS 7.2K rpm 3.5" HDD/hot plug/C200		R200-D1TC03
2-TB SAS 7,2K rpm 3.5" HDD/hot plug/C200		
	2.5" Drives	PID
Gen 2 300-GB SAS 15K rpm 3.5" HDD/hot plu	500-GB SATA 7.2K rpm SFF HDD/hot plug/C-Series drive sled	R2XX-D500GCSATA
450-GB SAS 15K rpm 3.5" HDD/hot plug/C200	73-GB 6-Gb SAS 15K rpm SFF HDD/hot plug/drive sled mounted	A03-D073GC2
	146-GB 6-Gb SAS 10K rpm SFF HDD/hot plug/drive sled mounted	A03-D146GA2
	300-GB 6-Gb SAS 10K rpm SFF HDD/hot plug/drive sled mounted	A03-D300GA2
	100-GB SATA SSD SFF HDD (C250 only)	A03-D100SSD

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-13

When the RAID controller and RAID level have been determined, the disk drives need to be selected. The designer has to observe the expected I/O intensity to select the proper drives—the cache size, rpm, and size determine the performance.

Second, the selected RAID level for redundancy might govern the minimum number of disks—for example, at least two disks for RAID 1, and more for RAID 5 or 6. RAID 10 requires at least four disk drives.

Configure Selected Servers—Step 5

- Select PCIe adapter options
 - All cards are low profile or half-length cards
 - C200 server supports a maximum of two PCIe cards
 - C210 and C250 servers support a maximum of five PCIe cards
 - Available slots reduced by one if choose RAID option (Step 3)

PCIe Card Description	PID
QLogic QLE8152 Dual Port 10-Gb Ethernet-to-PCIe Converged Network Adapter (CNA) (two ports, copper)	N2XX-AQPCI01
Broadcom NetXtreme II 5709 Quad Port Ethernet PCIe Adapter Card with TOE and iSCSI HBA	N2XX-ABPCI03
Broadcom NetXtreme II 57711 Dual Port 10 Gigabit Ethernet PCIe Adapter Card with TOE and iSCSI HBA	N2XX-ABPCI02
Emulex LightPulse LPe11002 4-Gb Fibre Channel PCI Express Dual Channel HBA	N2XX-AEPCI03
QLogic SANblade QLE2462 Dual Port 4-Gb Fibre Channel-to-PCI Express HBA	N2XX-AQPCI03
Emulex -Tiger Shark Converged Network Adapter	N2XX-AEPCI01
Intel 10 Gigabit Ethernet two-port Niantec Controller	N2XX-AIPCI01

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-14

The connectivity is the vital part of the server configuration. The selection of the number and type of PCIe adapters is governed by the following:

- Connectivity requirement—LAN, SAN (Fibre Channel versus iSCSI)
- Throughput requirement (1 Gigabit Ethernet, 10 Gigabit Ethernet, 4 Gigabit Fibre Channel, and so on)
- Number of segments to connect to (for example, VMware ESX/ESXi—management, vMotion, FT logging, data segments, and so on)

Note that different server types have different numbers of free PCIe slots that can be used for PCIe network adapters. In cases when an advanced RAID controller is used, one of these slots might already be populated with a RAID controller.

Configure Selected Server—Step 6

- Select power supply options
 - Redundant power supply (minimum 0, maximum 1)
 - One power supply always comes standard with the base chassis, redundant can be ordered
 - 650-W PSU for UCS C200 and C210 servers
 - 750-W PSU for UCS C250 servers
- Select power cords (minimum 0, maximum 2)

Country Type	Description	PID
North America	N5000 AC Power Cable, 6-A, 250-V, 2.5-m	CAB-N5K6A-NA
North America	N5000 AC Power Cable, 13-A, 250-V, 2.5-m	CAB-AC-250V/13A
North America	N5000 AC Power Cable, 6-A, 250-V, Power Strip Type	CAB-C13-C14-JMPR
North America	N5000 Power Cord, 125-VAC 15-A NEMA 5 – 15 Plug, 2.5-m	CAB-9K12A-NA
Argentina	N5000 AC Power Cable, 10-A, 250-V, 2.5-m	SFS-250V-10A-AR
Australia	N5000 AC Power Cable, 10-A, 250-V, 2.5-m	CAB-9K10A-AU
China	N5000 AC Power Cable, 10-A, 250-V, 2.5-m	SFS-250V-10A-CN
Europe	N5000 AC Power Cable, 10-A, 250-V, 2.5-m	CAB-9K10A-EU
...

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-15

Individual Cisco UCS C-Series servers also require power supplies to operate. Each server comes with one power supply, which is sufficient for operation. When redundancy from a power supply perspective is required, the second power supply should be selected.

It is also important to select the proper power supply cables.

Configure Selected Servers—Step 7

- Select optional components
 - Rail kit (minimum 0, maximum 1)
 - Cable management arm (minimum 0, maximum 1)
 - TPM (minimum 0, maximum 1)
- Select operating system or hypervisor

PID	Description
RHEL-1A	RHEL/1yr subscription/svcs required/0 media
RHEL-3A	RHEL/3yr subscription/svcs required/0 media
RHEL-AP-1A	RHELAP/1yr subscription/svcs required/0 media
RHEL-AP-3A	RHELAP/3yr subscription/svcs required/0 media
SLES-1A	SLES/1yr subscription/svcs required/0 media
SLES-3A	SLES/3yr subscription/svcs required/0 media
MSWS-03R2-ST	Windows Svr 2008 ST w/ 2003 R2 downgrade, 0 media
MSWS-03R2-ST64	Windows Svr 2008 ST w/ 2003 R2 x64 downgrade, 0 media
MSWS-03R2-EN	Windows Svr 2008 EN w/ 2003 R2 downgrade, 0 media
MSWS-03R2-EN64	Windows Svr 2008 EN w/ 2003 R2 x64 downgrade, 0 media
VMW-VS-ADV-1A	VMware vSphere Advanced (1 CPU), 1yr 24x7 support
VMW-VS-ADV-3A	VMware vSphere Advanced (1 CPU), 3yr 24x7 support
...	...

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-16

With the Cisco UCS C-Series server, some optional components can be selected:

- Rail mount kit for easier installation of servers into the cabinets
- Cable management arm
- TPM for sensitive data that needs to be securely stored (for example, keys and certificates)
- Operating system or hypervisor

Creating the C-Series BOM

This topic identifies and describes how to create a C-Series solution BOM.

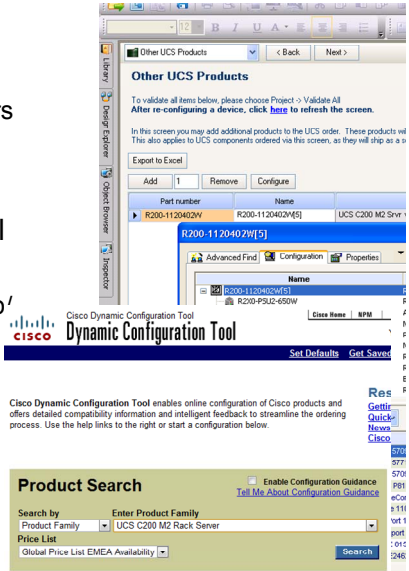
Create BOM

- NetformX DesignXpert
 - Available to Cisco UCS ATP partners
 - Acquire credentials
 - Not a sizing tool
- Cisco Dynamic Configuration Tool
 - Available to CCO account holders
 - <https://tools.cisco.com/qtc/config/jsp/configureHome.jsp>

BOM creation with DesignXpert

- Configuration
- Service selection
- Upload

<http://www.ciscoprc.com>



© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-18

The BOM for the Cisco UCS C-Series can be created using one of the following two tools:

- NetformX DesignXpert
- Cisco Dynamic Configuration Tool

The NetformX DesignXpert tool is available at <http://design.netformx.com/cisco-navigate-to-accelerate/>. In order to be able to use the tool to compile the BOM, proper credentials must be acquired. The credentials are available to Cisco UCS ATP partners.

Once the tool is installed with proper credentials, the BOM can be created. This process has three steps:

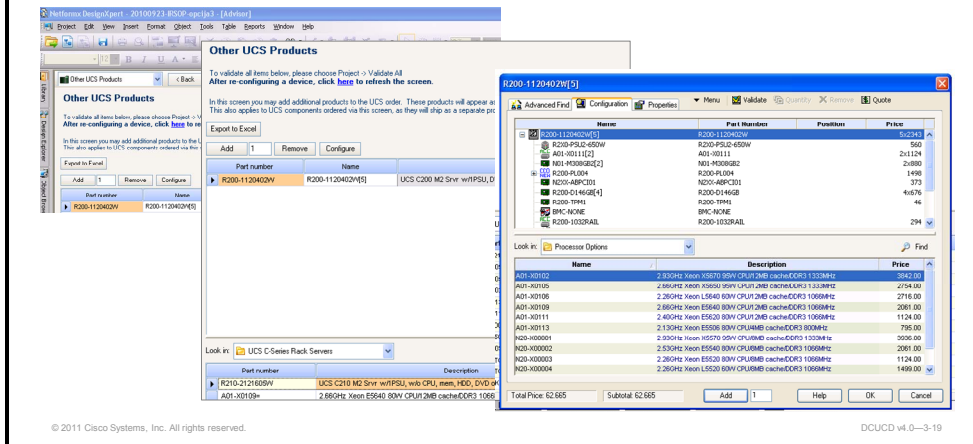
- Configuration
- Service selection
- Upload

More information about the NetformX DesignXpert tool is available at the Cisco Partner Resources Center at <http://www.ciscoprc.com>.

The Cisco Dynamic Configuration tool is available at <https://tools.cisco.com/qtc/config/jsp/configureHome.jsp> and requires a Cisco Connection Online account for access.

Configuration—Other UCS Products

- Pick Cisco UCS C-Series
 - Configure server components (CPU, memory, adapter, disks, controller, etc.)

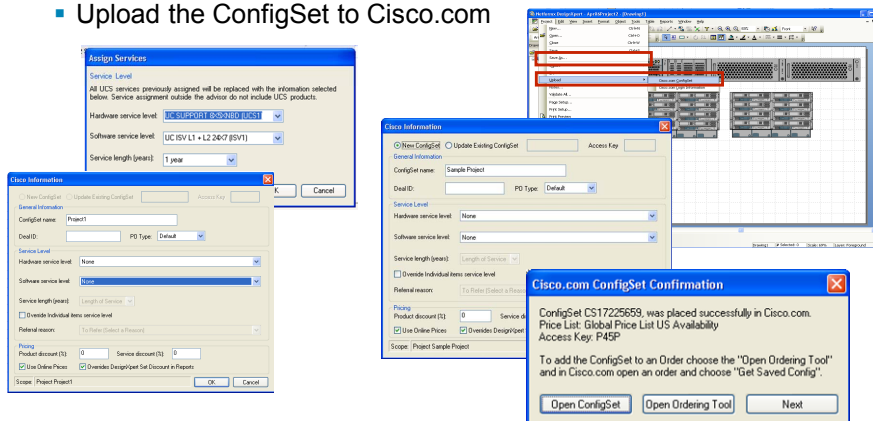


The configuration process employs the Cisco UCS Advisor. The last of the tabs available is “Other UCS Products.”

On this tab, Cisco UCS C-Series server types can be selected with the proper quantity and then configured with the proper components like CPU, memory DIMMs, PCIe adapters, RAID controller and disk drives, and so on.

Service Selection and ConfigSet Upload

- Apply services individually or to the entire system
- Review BOM by different categories
- Upload the ConfigSet to Cisco.com



The service selection process is used to define the type of services that are to be applied to the selected Cisco UCS C-Series components; the services can be selected for the whole BOM or part of the BOM. The selection includes definition of the hardware and software service level, length of the service, and so on.

Before proceeding to the final configuration step, which is uploading, the BOM should be reviewed to verify that it is correct.

Once the hardware and services have been selected, the BOM ConfigSet can be uploaded to Cisco.com.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Important aspects of server sizing are LAN and SAN connectivity, memory, and local storage configuration.
- Cisco UCS C-Series sizing process encompasses selection of server type and components (CPU, memory, PCIe adapter, RAID controller, disk drives, optional components).
- Sizing process defines the BOM, which can be created with NetformX DesignXpert or Cisco Dynamic Configuration tool.

Evaluating Cisco UCS B-Series Architecture

Overview

This lesson assesses the Cisco Unified Computing System (UCS) B-Series hardware components, connectivity, high-availability options, and management.

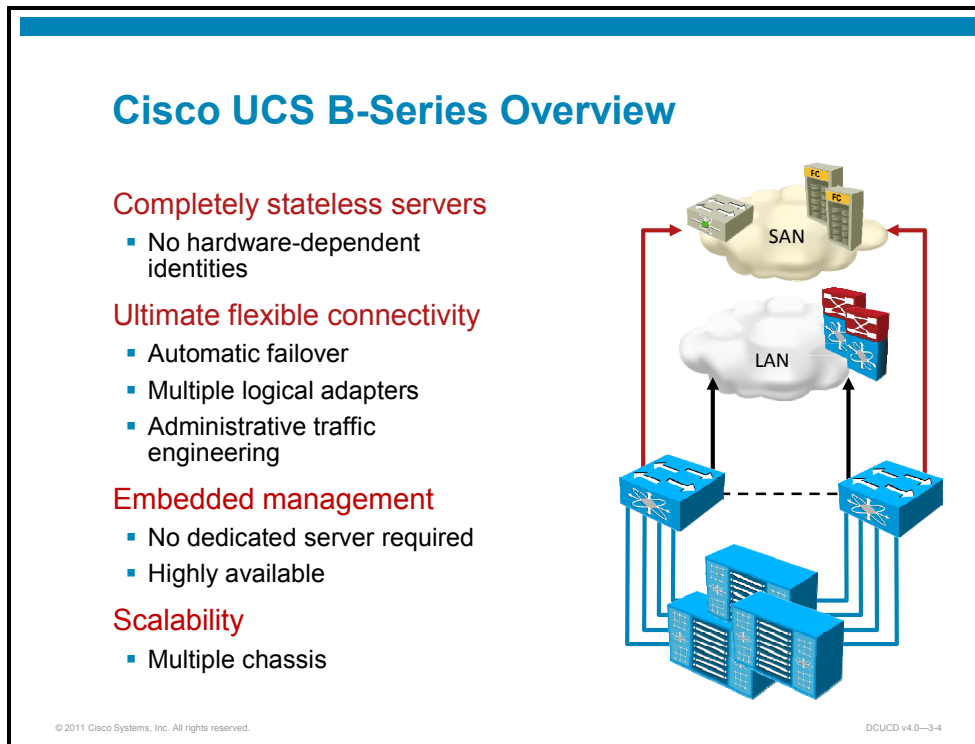
Objectives

Upon completing this lesson, you will be able to identify and describe Cisco UCS B-Series architecture. This includes the ability to meet these objectives:

- Identify and describe Cisco UCS B-Series overall architecture
- Identify and describe Cisco UCS 6100 Fabric Interconnect switches
- Identify and describe Cisco UCS 5108 chassis
- Identify and describe Cisco UCS B-Series server blades
- Identify and describe Cisco UCS B-Series memory architecture
- Identify and describe Cisco UCS B-Series mezzanine options
- Identify and describe Cisco UCS Manager and options
- Identify and describe Cisco UCS B-Series LAN connectivity
- Identify and describe Cisco UCS B-Series SAN connectivity

Evaluating Cisco UCS B-Series

This topic identifies and describes UCS B-Series overall architecture.



Cisco UCS is a data center platform that represents a pool of compute resources that are connected to existing LAN and SAN core infrastructures. It is designed to provide the following:

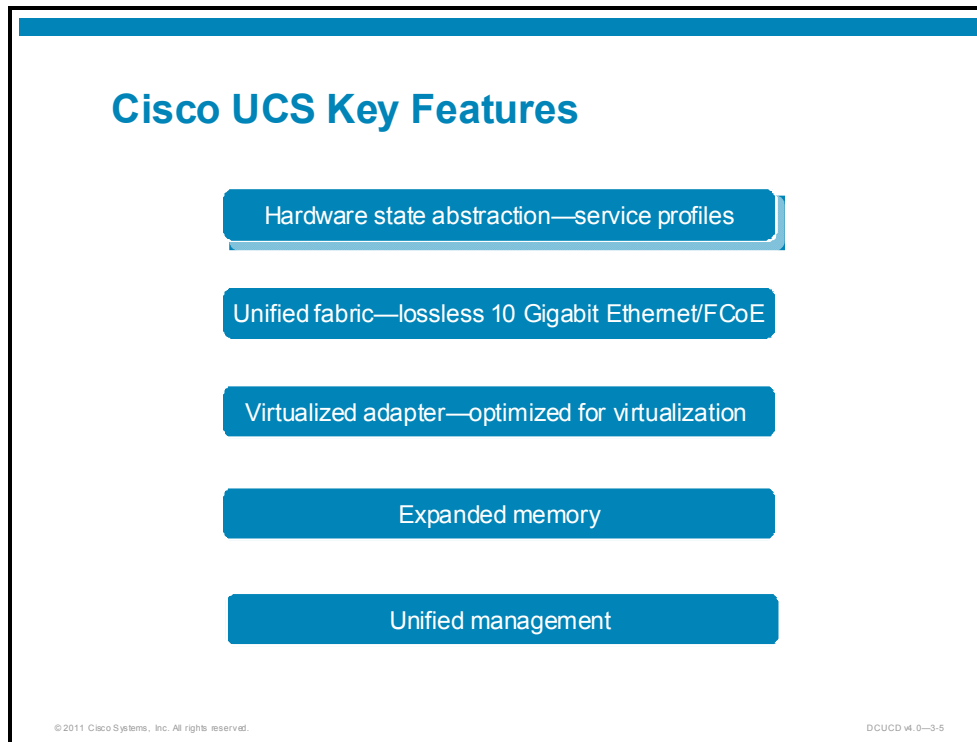
- Improve responsiveness to new and changing business demands
- Ease and accelerate design and deployment of new applications and services
- Provide simple, reliable, and secure infrastructure at the same time

From the perspective of server deployment, Cisco UCS represents a dynamic, “cable once” environment that enables rapid provisioning of new services. The system offers great scalability since a single Cisco UCS system can consist of up to 40 chassis, each hosting up to 8 server blades. This scalability means that the administrator has a single management and configuration point for up to 320 server blades.

Since the unified fabric is an integral part of Cisco UCS, fewer cables are required to connect the system components and fewer adapters need to be installed in the servers.

The network part of the system is realized with the fabric extender concept, which results in fewer switch devices.

Fewer system components result also in lower power consumption, which makes the Cisco UCS solution greener—that is, better power consumption ratio per computing resource is achieved.



Hardware State Abstraction

Using service profiles, the Cisco UCS solution is able to abstract the items that make a server unique. This ability allows the system to see the blades as being agnostic. As a result, moving, repurposing, upgrading, and making servers highly available is very easy in Cisco UCS.

Unified Fabric (10 Gigabit Ethernet/Fibre Channel over Ethernet)

Unified fabric consolidates the different network types that are used for various communications on to a single 10 Gigabit Ethernet. The Cisco UCS solution uses a single link for all communications (data and storage), with the fabric provided through the system controlling device. Therefore, no matter what blade you move a server to, the fabric and communications remain the same.

Virtualized Adapter

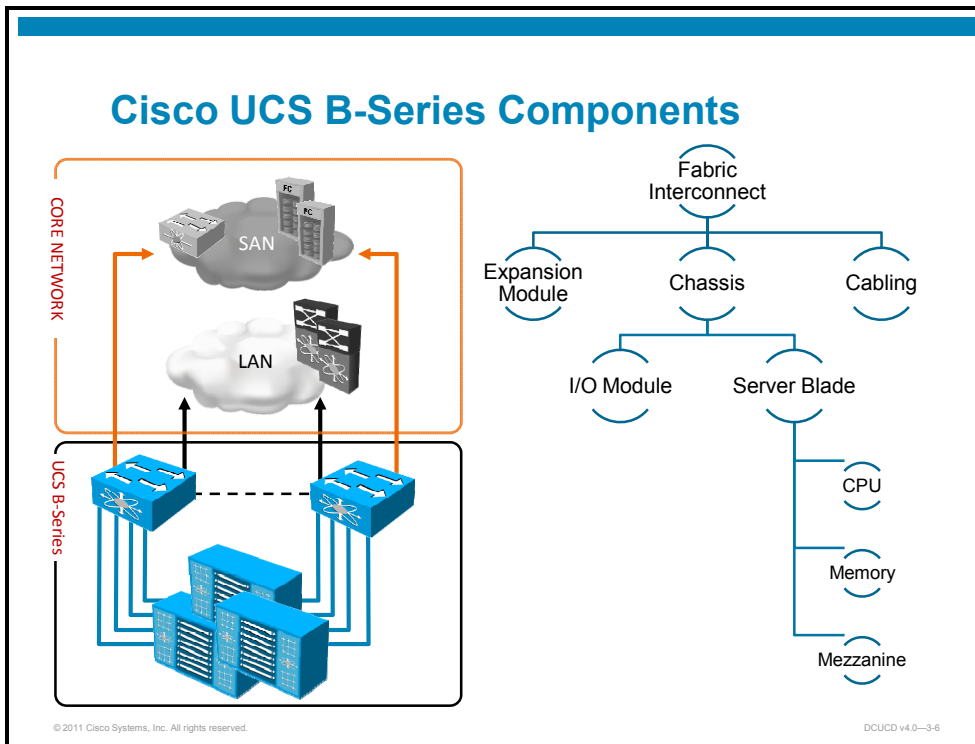
As part of the statelessness, Cisco UCS is designed to provide visibility and control to the networking adapters within the Cisco UCS solution. This visibility and control is achieved with the software running on Cisco UCS and the implementation of the virtual interface card adapters, which, in addition to allowing the creation of virtualized adapters, also alleviates the management overhead that is normally handled by the hypervisor, thus increasing performance.

Expanded Memory

The larger memory footprint of the Cisco UCS B250-M1 2-Socket Extended Memory Blade Server offers a number of advantages to applications that require a larger memory space. One of those advantages is the ability to provide the large memory footprint, using standard-sized and lower-cost DIMMs.

Unified Management

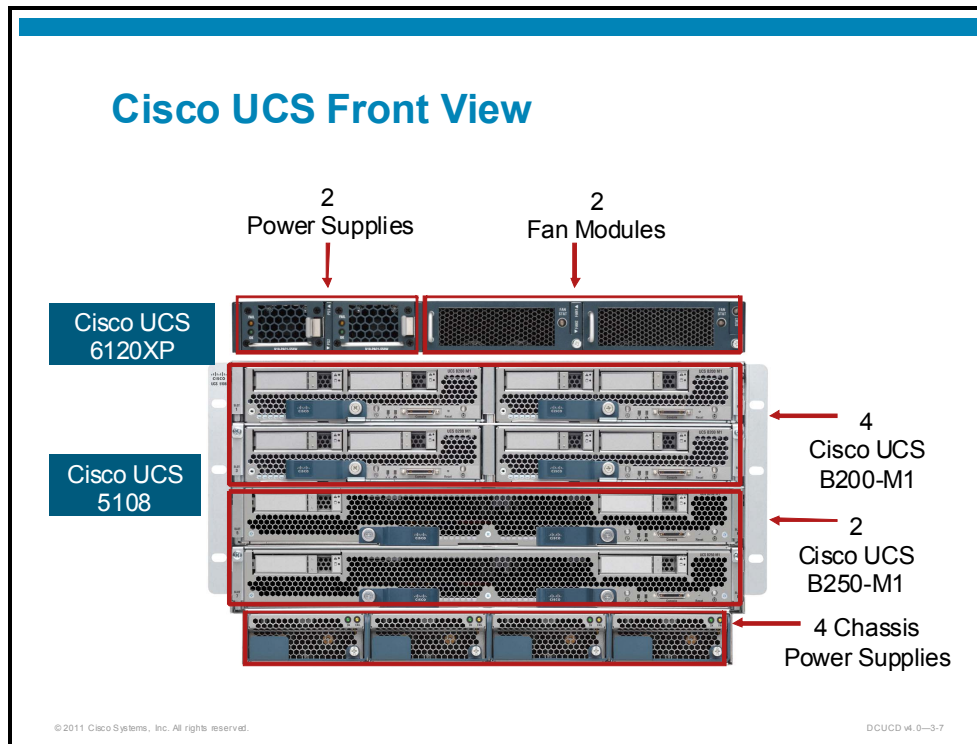
The Cisco UCS Manager is embedded device-management software that manages the system from end to end as a single logical entity through either an intuitive GUI, a command-line interface (CLI), or an XML application programming interface (API).



The Cisco UCS comprises the following components:

- Two Cisco UCS 6100XP Fabric Interconnect switches in a cluster deployment:
 - Provides a single point of management by running the Cisco UCS Manager application
 - Uses a single physical media for LAN and SAN connectivity (consolidates the I/O with Fibre Channel over Ethernet [FCoE])
 - Provides simultaneous traffic switching in high-availability configuration
- Up to 40 Cisco UCS 5108 Chassis per system if Cisco UCS 6140 Fabric Interconnect switches are used:
 - Two Cisco UCS 2104XP IO Modules (IOMs) or Fabric Extenders (FEX) per chassis.
 - Each chassis can host up to 8 server blades (for a total of up to 320 server blades).
 - Server equipped with one or two network adapters.

Note It is possible to use a single Cisco UCS 6100XP Fabric Interconnect in a system, although this configuration is not encouraged since it does not provide adequate redundancy.



The front view shows the Cisco UCS 6100XP Fabric Interconnect switch and the Cisco UCS 5108 Chassis.

The Cisco UCS 6100XP Fabric Interconnect switch front view shows the following:


- Power supplies
- Fan trays

The Cisco UCS 5108 Chassis front view shows the following:

- Power supplies
- Server blades

As indicated on the figure, a single chassis can support half-width and full-width server blades at the same time. Note that a half-width blade consumes one slot, whereas a full-width blade consumes two chassis slots.

In the recommended Cisco UCS setup, two Cisco UCS 6100XP Fabric Interconnect switches are used to achieve proper redundancy and high-availability level. Depending on the fabric interconnect type (20- or 40-port), up to 40 chassis can be present in the system.



The rear view of the Cisco UCS systems shows all the cabling locations that the system needs. Cabling for the switches, chassis power, console, cluster interconnect, Ethernet, and Fibre Channel is attached in the rear.

The Cisco UCS 6100XP Fabric Interconnect rear view shows the following:

- Two cluster ports used for the high-availability fabric interconnect setup
- Out-of-band Ethernet management port
- Console port
- Expansion module bay (or two in the case of Cisco UCS 6140XP)
- Two power entries

Note The fabric interconnect switch shown in this slide is the Cisco UCS 6120XP 20-Port Fabric Interconnect.

The Cisco UCS 5108 Chassis rear view shows the following:

- Two Cisco UCS 2104XP IOM slots (and thus two I/O modules when installed)
- Eight fan modules
- Power entry module with four power entries

A fully redundant Cisco UCS setup includes two Cisco UCS 6100 Fabric Interconnects as well as two Cisco 2104XP IOMs.

In a typical high-availability setup, the left Cisco UCS 2104XP IOM connects to the left Cisco UCS 6100XP Fabric Interconnect switch, and the right connects to the right Cisco UCS Fabric Interconnect. The left pair is Fabric A and the right is Fabric B.


The described setup is recommended for easier configuration and troubleshooting.

Cisco UCS 6100 Fabric Interconnects

This topic identifies and describes UCS 6100 Fabric Interconnect switches.

Cisco UCS 6120XP Fabric Interconnect

- Ports (disabled by default)
 - 20 fixed SFP+ 10 Gigabit Ethernet/FCoE ports
 - 1 Gigabit Ethernet on first eight SFP-compatible ports
 - One expansion module bay for 10 Gigabit Ethernet/FCoE or Fibre Channel-only ports
- Layer 2 characteristics
 - 520 Gb/s or 386.9 Mp/s forwarding performance
 - 16,000 MAC address table entries



A photograph of the Cisco UCS 6120XP Fabric Interconnect switch, a 1RU rack-mountable device. The front panel features a variety of ports including SFP+ ports, a Gigabit Ethernet port, and an expansion module bay. A vertical double-headed arrow on the right side of the device is labeled '1 RU', indicating its height. The device is dark grey with a blue accent on the left side.

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-10

The Cisco UCS 6100XP series Fabric Interconnect switches are 20- or 40-port switches and represent an integral part of the Cisco UCS system. They provide both enterprise Ethernet LAN and Fibre Channel SAN network connectivity to the server blades that are housed in a Cisco UCS 5108 chassis.

The Cisco UCS 6100XP Fabric Interconnect uses a cut-through architecture to support deterministic, low latency, line-rate 10 Gigabit Ethernet and FCoE on all ports, independent of packet size and enabled services.

From the management perspective, the fabric interconnects offer a single high-availability management domain to all the server blades in the domain by embedding the Cisco UCS Manager application.

The switches also enable enhanced virtualization in VMware environments with Cisco Virtual Network Link (Cisco VN-Link) support by providing policy-based virtual machine connectivity, mobility of network properties with the virtual machine, and a consistent operational model for both physical and virtual environments.

Cisco UCS 6120XP Fabric Interconnect

The Cisco UCS 6120XP is a fabric interconnect switch with the following characteristics:

- 1 rack unit (RU) high
- 20 fixed small form-factor pluggable plus (SFP+) 10 Gigabit Ethernet/FCoE port
- First 8 ports supporting 1 Gigabit Ethernet upstream connectivity with proper SFP
- One expansion module bay where 10 Gigabit Ethernet-only, 10 Gigabit Ethernet, and Fibre Channel or Fibre Channel-only expansion modules can be inserted. Ports on these modules are typically used for LAN and SAN connectivity.

- Layer 2 forwarding performance of 520 Gb/s (full-duplex) or 386.9 million packets per second (Mp/s)
- Up to 16,000 MAC address table entries
- Up to 512 VLANs and virtual storage area networks (VSANs) per fabric interconnect switch
- Rapid Per VLAN Spanning Tree Plus (PVRST+)
- Cisco and Link Aggregation Control Protocol (LACP) EtherChannel support with Layer 2, 3, or 4 hashing support
- Jumbo frames
- Internet Group Management Protocol (IGMP) version 1, 2, or 3 snooping

All ports are disabled by default and unconfigured.

Cisco UCS 6140XP Fabric Interconnect

- Ports (disabled by default)
 - 40 fixed SFP+ 10 Gigabit Ethernet/FCoE ports
 - 1 Gigabit Ethernet on first 16 SFP-compatible ports
 - Two expansion module bays for 10 Gigabit Ethernet/FCoE or Fibre Channel-only ports
- Layer 2 characteristics
 - 1.04 Tb/s or 773.8 Mp/s forwarding performance
 - 16,000 MAC address table entries



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-11


Cisco UCS 6140XP Fabric Interconnect

The Cisco UCS 6140XP is a fabric interconnect switch with the following characteristics:

- 2 RU high
- 40 fixed SFP+ 10 Gigabit Ethernet/FCoE ports
- First 16 ports supporting 1 Gigabit Ethernet upstream connectivity with proper SFP
- Two expansion module bays where 10 Gigabit Ethernet-only, 10 Gigabit Ethernet and Fibre Channel or Fibre Channel-only expansion modules can be inserted. Ports on these modules are typically used for LAN and SAN connectivity.
- Layer 2 forwarding performance of 1.04 terabits per second (Tb/s) (full-duplex) or 773.9 Mp/s
- Up to 16,000 MAC address table entries
- Up to 512 VLANs and VSANs per fabric interconnect switch
- PVRST+
- Cisco and LACP EtherChannel support with Layer 2, 3, or 4 hashing support
- Jumbo frames
- IGMP version 1, 2, or 3 snooping

All ports are disabled by default and unconfigured.

Cisco UCS Expansion Modules



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-12

The expansion modules are installed into the fabric interconnect expansion module bay. They are used to provide the connectivity to the enterprise Ethernet LAN and Fibre Channel SAN networks, and as such provide additional scalability to the UCS.

There are three expansion module types: Fibre Channel-only, Combo, and Ethernet-only.

Fibre Channel-Only Expansion Module – N10-E0080

This module offers eight 1/2/4 Gb/s Fibre Channel SFP ports for transparent connectivity with existing Fibre Channel networks.

Combo Expansion Module – N10-E0440

This module offers a combination of four 10 Gigabit Ethernet/FCoE SFP+ ports and four 1/2/4 Gb/s Fibre Channel SFP ports.

Ethernet-Only Expansion Module – N10-E0600

This module offers six 10 Gigabit Ethernet/FCoE SFP+ ports.

This Ethernet expansion module contains six SFP+ ports that support 10 Gigabit Ethernet.

Fibre Channel-Only Expansion Module – N10-E0060

This module offers six 1/2/4/8 Gb/s Fibre Channel SFP ports for transparent connectivity with existing Fibre Channel networks.

Cisco UCS Port Options

- SFP+
 - 10 Gigabit Ethernet
 - Hot-swappable
 - Optical interoperability with 10GBASE XENPAK, X2, and XFP interfaces on the same link

Type	Cable	Distance
CU SFP+	Twinax	1, 3, or 5 m
SR SFP+	MMF OM2	82 m
SR SFP+	MMF OM3	300 m
LR SFP+	SMF	Up to 10 km



- SFP 1 Gigabit Ethernet in SFP-compatible ports
 - GLC-T and SFP-GE-T copper based
 - GLC-SX-MM, GLC-SX-SM, SFP-GE-S, SFP-GE-L fiber optics

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-13

The Cisco UCS 10 Gigabit Ethernet/FCoE ports can house various SFP+ modules:


- Copper-based (CU) using Twinax cable—1, 3, or 5 m.
- Short reach (SR) multimode fiber-based (MMF):
 - Optical Multimode 2 (OM2) cable type—up to 82 m
 - Optical Multimode 3 (OM3) cable type—up to 300 m
- Long reach (LR) single-mode fiber-based (SMF)—up to 10 km

The CU SFP+ presents a cost-effective solution for short distances (for example, physical connectivity for Cisco UCS 5108 Chassis to Cisco UCS 6100XP Fabric Interconnect switch).

Apart from 10 Gigabit, the first 8 or 16 ports support 1 Gigabit Ethernet connectivity with standard Cisco SFP modules.

Fabric Interconnect Cluster Connectivity

- Two Cisco UCS 6100XP Fabric Interconnects in redundant configuration
- Cluster link between two fabric interconnects
 - Carries cluster heartbeat messages and high-level messages for the Cisco UCS Manager elements
- L1 and L2 part of the 802.3ad channel
 - 10/100/1000 Ethernet
 - RJ-45 form factor
 - Must follow L1 to L1, L2 to L2 cabling



Cisco UCS is usually deployed in a clustered fashion, that is, with two Cisco UCS 6100 switches connected over Layer 1 and Layer 2 cluster links. This setup provides redundancy for management as well as switching functionality. The Cisco UCS Fabric Interconnects that are in a cluster are either of the following:

- Primary node: the fabric interconnect that is active
- Subordinate node: the fabric interconnect that is on standby

The standby node runs a Cisco UCS Manager instance with reduced functionality.

The following connectivity requirements must be met for successful deployment of a highly available Cisco UCS cluster:

- Connect L1 of fabric interconnect A to L1 of fabric interconnect B
- Connect L2 of fabric interconnect A to L2 of fabric interconnect B.
- Connect fabric interconnect A to IOM A of each chassis, using one to four uplinks.
- Connect Fabric Interconnect B to IOM B of each chassis, using one to four uplinks.


Cluster interfaces provide a cluster link between two Cisco UCS 6100 Series Fabric Interconnects. They carry the cluster heartbeat messages between the two Cisco UCS Fabric Interconnects as well as high-level messages between Cisco UCS Manager elements. These links are part of an IEEE 802.3ad bond managed by the underlying operating system. The bond is configured to run LACP, which brings up the bond link only when there is either a single link between two LACP-enabled nodes, or when both links are between LACP-enabled peers. The IP addresses on these links are fixed.

The management port (mgmt0) of each fabric interconnect should be connected to the same Layer 2 network to facilitate failover and failback of the management IP address. Each fabric interconnect should connect to only one side of each chassis.

Cisco UCS B-Series Connectivity Overview

IOM Fabric Port Pinning

No. Of Fabric Ports	Fabric Port	Server Blade
1	Port 1	1, 2, 3, 4, 5, 6, 7, 8
2	Port 1	1, 3, 5, 7
	Port 2	2, 4, 6, 8
4	Port 1	1, 5
	Port 2	2, 6
	Port 3	3, 7
	Port 4	4, 8



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-15

Cisco UCS components are interconnected using physical and logical interfaces or ports.

The fabric interconnects are connected in the following manner:

- Via physical uplink ports to external LAN or SAN network
- Via physical server ports to IOMs

Each IOM provides the following:

- Eight internal backplane ports
- Four external fabric ports
- One internal management network

An individual IOM is connected in the following manner:


- Via physical fabric port to fabric interconnects server port
- Via logical backplane port through chassis midplane to each individual server blade

An individual server blade is connected to both IOMs. Depending on the number of IOM fabric ports, the server blades ports are pinned in the following manner:

- All server blades are pinned to IOM port 1 when a single fabric port is used.
- Server blades 1, 3, 5, 7 are pinned to IOM port 1, server blades 2, 4, 6, 8 to IOM port 2, when two fabric ports are used.
- Server blades 1 and 5 are pinned to IOM port 1, server blades 2 and 6 to port 2, server blades 3 and 7 to port 3, and server blades 4 and 8 to port 4, when four fabric ports are used.

Cisco UCS 5108 Chassis

This topic identifies and describes the Cisco UCS 5108 chassis.



The Cisco UCS 5108 Blade Server Chassis is 32 inches deep and 6 RU high. The blade servers and power supplies are installed from the front and are locked into place using thumbscrews. The same chassis may be used for the following:

- Up to eight half-width blade servers—Cisco UCS B200-M1 2-Socket Blade Servers
- Up to four full-width blade servers—Cisco UCS B250-M1 2-Socket Extended Memory Blade Servers
- Any physically possible combination of half- and full-width blade servers


Note For the half-width blade servers, a vertical divider is used to separate horizontally adjacent blade servers. To install full-width blade servers, the divider must be removed.

A Cisco UCS system can scale up to as many as 20 chassis (using Cisco UCS 6120XP) or 40 chassis (using Cisco UCS 6140XP).

Chassis Power

There are a maximum of four 2500 W hot-pluggable power supplies per chassis. The actual number of power supplies required depends on the system's hardware configuration and desired power redundancy.

Cisco UCS 5108 Chassis Rear View



From the rear of the Cisco UCS 5108 chassis, the two Cisco UCS 2104XP IOMs are installed. These provide connection between the blade server mezzanine cards (network adapters) and the fabric interconnect switches.


For the front-side power supply units, four 220 V power entry ports at the rear of the chassis provide power.

Chassis Cooling

The chassis cools from front to rear with a high-efficiency, high-reliability flow-through airflow design. The fan modules are sized to support all blade servers running at full power budget. There are eight hot-pluggable dual-fan modules per chassis with status indicators on each module. The fan modules provide cooling to the blade servers and chassis components.

Cisco UCS 5108 Chassis Power Supply

- Redundancy modes
 - Nonredundant mode
 - N+1 redundancy
 - Grid redundancy
- Input power lead failure—power loss on associated power supply
- All active power supplies available for all components



Power supplies are single-phase 220 V, IEC320-C 19. The power supplies provide a 550 W budget per slot for up to eight half-width blade servers and an 1100 W budget per slot for up to four full-width blade servers. This design supports future growth and power budget requirements with no service disruption. The efficiency of the power supplies reaches 92 percent.

Chassis Power Supply Modes

The Cisco UCS 5108 chassis supports multiple modes of power supply redundancy:

- A nonredundant mode provides the power capacity to support the running of Cisco UCS with all installed power supplies active. If any power supply fails, some or all components may be affected.
- An N+1 redundancy mode enables Cisco UCS to tolerate a failure of any single power supply. This configuration requires one more power supply than the nonredundant configuration.
- A grid redundancy mode enables Cisco UCS to tolerate a loss of an entire power grid within the data center if properly cabled. The mode does require twice the nonredundant configuration.


Chassis Power Connectivity

The power inlet ports on the Cisco UCS 5108 chassis have no crossover connections to the chassis power supplies (that is, they are not cross-connected). Thus if any component in the path to the power supply fails, the power supply experiences power loss.

The power supply outputs have crossover connections to the chassis components (that is, they are cross-connected). This means that a failure of one power supply does not affect any other component in a valid redundant power supply configuration.

Cisco UCS 2100 IOM

- Four external fixed SFP+ 10 Gigabit Ethernet/FCoE ports
- Eight internal 10 Gigabit Ethernet/FCoE ports
 - Connected through the midplane to each half-width slot
- No local switching
- Up to 8:1 oversubscription rate supported
- Chassis Management Controller (CMC)



SFP+ 10 Gigabit Ethernet Port

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-20

The Cisco UCS 2104XP IOM, or FEX, logically extends the fabric from the fabric interconnect switch to the blade server. Its primary function is to connect the blade servers to the fabric interconnect, which in turn handles all switching for blade-to-blade and blade-to-fabric communications.

Since the IOM is similar to a distributed line card, no local switching occurs in the IOM and the management is done from fabric interconnect switches.

The IOM has the following ports:

- Four fixed SFP+ 10 Gigabit Ethernet/FCoE external ports for fabric interconnect connectivity
- Eight internal 10 Gigabit Ethernet/FCoE ports connected through the midplane to each half-width slot in the chassis

The IOM is installed into the back of the Cisco UCS 5108 chassis and is typically used in pairs to provide redundant fabric connectivity as well as increased capacity. When both are used, they are hot-swappable and fully redundant from an Ethernet perspective. (Fibre Channel redundancy is discussed later.) Each blade server connects to both IOMs in the chassis via the mezzanine card.

Note The chassis can be equipped with one IOM only. In that case, the IOM must be placed into the left bay (as viewed from the rear of the chassis).

The IOM supports up to an 8:1 oversubscription ratio, which would be the case for a fully populated chassis with half-width blades, using a single 10 Gigabit Ethernet uplink from the IOM to the fabric interconnect.


Chassis Management Controller

The Cisco UCS 2104XP IOM also manages the chassis environment—the power supply and fans as well as the blades—in conjunction with the fabric interconnect, which eliminates the need for separate chassis management modules.

This management is performed by a chassis management controller (CMC). The CMC collects status data from the IOM using the Intelligent Platform Management Interface (IPMI) protocol over the inter-integrated circuit (I2C) serial bus. This information is then communicated to the management node using the Ethernet server link. The CMC also serves as a proxy for the Cisco UCS Manager to the blade servers for certain functionality, and has a role in the high-availability protocols.

Cisco UCS 5108 Chassis Midplane

- Redundant power and Ethernet connectivity
 - For IOM and server blades
- Redundant management path
- Dedicated Ethernet management
- Any component auto-discovery
- Up to 20 Gb/s I/O bandwidth per half-width slot



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-21

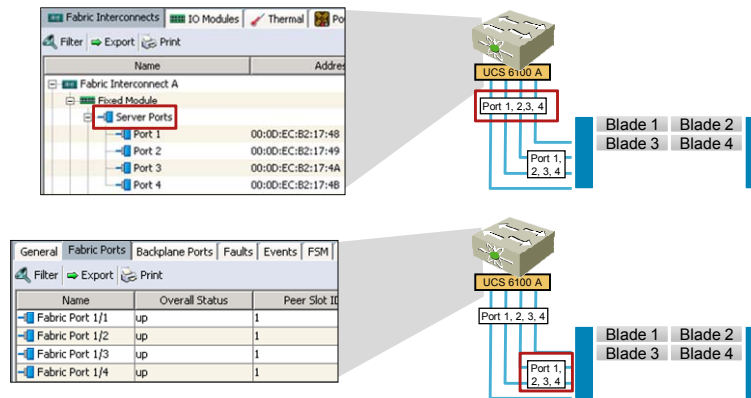
The Cisco UCS midplane provides the following features:

- Redundant power to the IOM and thus the CMC and blade servers that plug into it
- Redundant data network (Ethernet) connectivity between the IOM and the blade servers
- Redundant I2C management paths
- Dedicated management network (Ethernet) connectivity
- Auto-discovery of all components plugging into it

A passive midplane provides up to 20 Gb/s of I/O bandwidth per half-width slot and up to 40 Gb/s of I/O bandwidth per full-width slot. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

Chassis Connectivity

- Server ports—configured per fabric interconnect
- Fabric ports—fixed per IOM



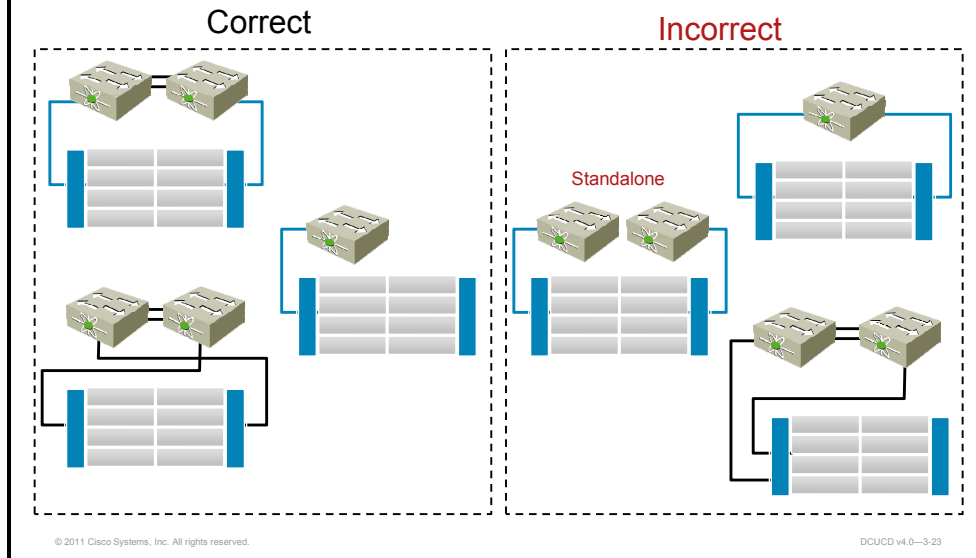
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-22

The fabric interconnect to IOM connectivity has the following characteristics:

- Fabric interconnect server ports are configurable—they depend on the actual fabric interconnect to IOM connectivity.
- IOM fabric ports are fixed—depending on the actual fabric interconnect to IOM connectivity, they may or may not be used.

IOM Connectivity Options




An individual IOM can be connected to only one fabric at a time. Connecting the IOM and chassis in a non-supported way will result in broken connectivity.

Cisco UCS B-Series Server Blades

This topic identifies and describes Cisco UCS B-Series server blades.

Cisco UCS B-Series Blade Servers Overview

- CPU sockets for Intel Xeon architecture
 - Must be of the same type per blade
- Internal SAS disk drive bays
- 2-, 4-, 8-, 16-GB DDR3 memory modules
- Support for statelessness (MAC, WWN, UUID, BIOS, boot order)
- Cisco IMC
 - Remote KVM access
 - Uses dedicated Ethernet link



Property	Value
Operating temperature	10 to 35 °C
Operating Humidity	5 to 93% noncondensing
Operating Altitude	0 to 3000 m

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-25

There are two types of blade servers available:

- Cisco UCS B200-M1, B200-M2, and B230-M1 half-width
- Cisco UCS B250-M1, B250-M2, and B440-M1 full-width

Three factors distinguish the blade servers:

- The form factor
- The number of DIMM slots
- The number of mezzanine adapter slots

The blade servers have a common (x86 Intel Xeon) architecture and share many features, such as being hot-swappable in the chassis, and having two or four CPU sockets, two internal Serial Attached SCSI (SAS) disk drive bays, and a Cisco Integrated Media Controller (IMC).

Note Only one CPU is required for normal system operation. If only one CPU is installed, then it must go in the first socket. The CPUs must be identical in the same blade server, but can be mixed between blade servers in the same chassis.

Cisco IMC

The baseboard management controller (BMC) is a microcontroller on the motherboard that provides “lights out” hardware status and configurability to the system management software over dedicated Ethernet links between the BMC and each IOM.

The BMC uses the IPMI protocol over the I2C serial bus to manage devices on the baseboard. The BMC is also responsible for providing remote keyboard, video, mouse (KVM) access to the end user through the switch over the dedicated Ethernet links.

Advanced Functionality

The choice of blade servers provides increased scalability:

- Number of CPUs to populate (1–4)
- Number of disk drives to populate (1–2)
- Amount of memory to install (depending on blade server, up to 384 GB)

Another Cisco UCS feature supported for the blade servers is statelessness, where the following server personality elements are able to be relocated between blade servers:

- MAC addresses
- Worldwide names (WWNs)
- Universally Unique Identifiers (UUIDs)
- BIOS version
- SAN boot configuration

B200 M1/M2 Blade

CPU	2x Intel Xeon processor architecture 5500/5600 (4-, 6-core)
Memory	12 DIMM slots—up to 96GB RAM
Disk	2x 2.5-in. hot-swappable disks (RAID 0, 1)
Mezzanine	1x 10 GE mezzanine
Form Factor	Half-width
I/O Throughput	20 Gb/s



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-26

B200-M1 Blade

The Cisco UCS B200-M1 2-Socket Blade Server (half-width blade server) contains two SFF drive bays that support SAS disks in a hardware RAID 0 or RAID 1 configuration.

The blade server has a 550 W power and cooling budget. The chassis can hold up to eight half-width blade servers.

The motherboard contains two CPU sockets and 12 DDR3 SDRAM DIMM memory sockets. There is a single mezzanine adapter slot.

The supported CPU option is the Intel Xeon 5500 series.

B200-M2 Blade

The Cisco UCS B200-M2 2-Socket Blade Server (half-width blade server) contains two SFF drive bays that support SAS disks in a hardware RAID 0 or RAID 1 configuration.

The blade server has a 550 W power and cooling budget. The chassis can hold up to eight half-width blade servers.

The motherboard contains two CPU sockets and 12 DDR3 SDRAM DIMM memory sockets. There is a single mezzanine adapter slot.

The supported CPU options are the Intel Xeon 5500 or Intel Xeon 5600 series.

B230 M1 Blade

CPU	2x Intel Xeon processor architecture 6500/7500 (4-, 6-, 8-core)
Memory	32 DIMM slots—up to 256 GB RAM
Disk	2x 2.5-in. hot-swappable disks (RAID 0, 1)
Mezzanine	1x 10 GE mezzanine
Form Factor	Half-width
I/O Throughput	20 Gb/s



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-27

B230-M1 Blade

The Cisco UCS B230-M1 2-Socket Blade Server (half-width blade server) contains two SFF drive bays that support SAS disks in a hardware RAID 0 or RAID 1 configuration.

The blade server has a 550 W power and cooling budget. The chassis can hold up to eight half-width blade servers.

The motherboard contains two CPU sockets and 32 DDR3 SDRAM DIMM memory sockets. There is a single mezzanine adapter slot.

The supported CPU options are the Intel Xeon 6500 or Intel Xeon 7500 series.

B250 M1/M2 Blade

CPU	2x Intel Xeon processor architecture 5500/5600 (4-, 6-core)
Memory	48 DIMM slots—up to 384 GB RAM
Disk	2x 2.5-in. hot-swappable disks (RAID 0, 1)
Mezzanine	2x 10 GE mezzanine adapters
Form Factor	Full-width
I/O Throughput	40 Gb/s



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-28

B250-M1 Blade

The Cisco UCS B250-M1 2-Socket, Extended Memory Blade Server (full-width blade server) has two SFF drive bays.

The blade has an 1100 W power and cooling budget. The chassis can hold up to four full-width blades.

The motherboard contains two CPU sockets and 48 DDR3 DIMM memory sockets, four times the memory capacity of the half-width blade. There are two mezzanine adapter slots, doubling the I/O bandwidth of the half-width blade.

The supported CPU option is the Intel Xeon 5500 series.

B250-M2 Blade

The Cisco UCS B250-M2 2-Socket, Extended Memory Blade Server (full-width blade server) has two SFF drive bays.

The blade has an 1100 W power and cooling budget. The chassis can hold up to four full-width blades.

The motherboard contains two CPU sockets and 48 DDR3 DIMM memory sockets, four times the memory capacity of the half-width blade. There are two mezzanine adapter slots, doubling the I/O bandwidth of the half-width blade.

The supported CPU options are the Intel Xeon 5500 and 5600 series.

B440 M1 Blade

CPU	4x Intel Xeon processor architecture 7500 (4-, 6-, 8-core)
Memory	32 DIMM slots—up to 256 GB RAM
Disk	4x 2.5-in. hot-swappable disks (RAID 0, 1, 5, 6, 10)
Mezzanine	2x 10 GE mezzanine adapters
Form Factor	Full-width
I/O Throughput	40 Gb/s



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-29

B440 M1 Blade

The Cisco UCS B440-M1 Four-Socket, Extended Memory Blade Server (full-width blade server) has four SFF drive bays.

The blade has an 1100 W power and cooling budget. The chassis can hold up to four full-width blades.

The motherboard contains four CPU sockets and 32 DDR3 DIMM memory sockets. There are two mezzanine adapter slots, doubling the I/O bandwidth of the half-width blades.

The supported CPU option is the Intel Xeon 7500 series.

Hard Disk Drives

- 2/4 SFF drive bays
- SAS backplane
- Drive options
 - 73 GB 15,000 rpm
 - 146 GB or 300 GB 10,000 rpm
 - 100 GB SSD
- RAID 0 or RAID 1
- RAID 0, 1, 5, 6, 10 and battery backup with M1-440 blade
- Diskless configuration supported (requires SAN/LAN boot)



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-30

Both half- and full-width blade servers have drive bays that plug into an SAS backplane. There are two supported drives:

- 73 GB 15,000 rpm SAS SFF drives
- 146 or 300 GB 10,000 rpm SAS SFF drives

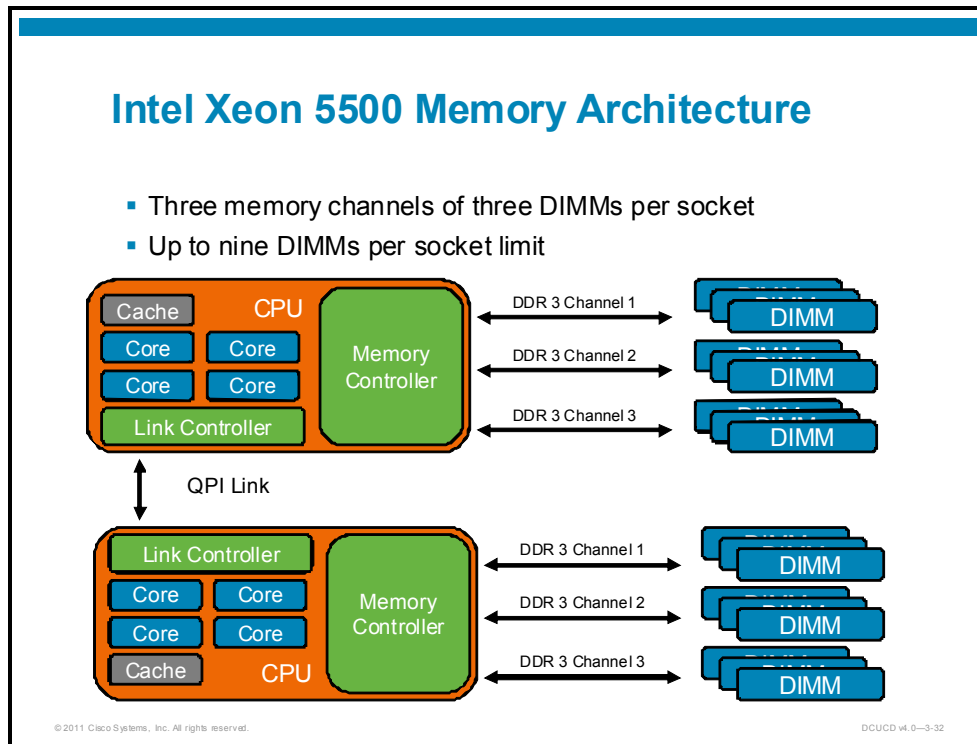
While no local hard drive is required, if only one is installed, then it must be placed in the first drive bay. If two hard drives are installed, they must both be of the same size and speed and are hot-swappable. The SATA controller supports RAID 0 and RAID 1.

The Cisco UCS disk drives include Cisco UCS feature support for statelessness by scrubbing any data on a disk drive during a service profile migration. The local disk policies (for example, RAID 0 or RAID 1) may be configured and can migrate with logical service profiles between blade servers.

The B440-M1 blade supports four disk drives with multiple RAID levels (0, 1, 5, 6, 10) with battery backup.

Cisco UCS B-Series Memory

This topic identifies and describes UCS B-Series memory architecture.



Intel Xeon 5500 series processors support three memory channels with three DIMMs per channel, which are governed mainly by the physical (electrical) characteristics of the memory bus. Note that installing multiple DIMMs per memory channel decreases the bus speed.

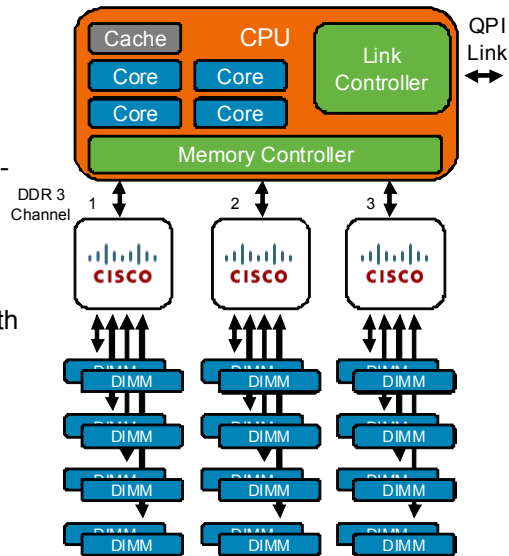
Keeping that in mind, many servers are deployed using no more than two DIMMs per channel to keep the bus speed high enough. Although the number of DIMMs per socket is nine, observing the bus speed limitation, the actual number of DIMMs per socket decreases to six.

Currently, the servers are typically limited to 12 DIMMs. If the servers use two sockets and 8-GB DIMMs, then the servers will use 96 GB of memory.

- The Cisco UCS B200-M1 server blade uses the same approach.
- If a vendor wishes to increase the amount of memory per server, additional sockets and thus processors need to be available and installed in the server. Of course, this change does not scale the amount of memory per socket.

Cisco Extended Memory Architecture

- Cisco ASIC
- Virtualizes 8 DIMMs per channel
 - Supported by B/C 250-M1/M2 servers
 - 192 GB per socket
 - 384 GB per blade
- Reduces memory cost with smaller DIMMs
- Operating speed varies



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-33

The Cisco Extended Memory Architecture employs a custom memory controller chipset that virtualizes DIMMs as seen by the CPU memory controller. With the technology, up to eight DIMMs can be installed per CPU memory channel without reducing the bus speed.




The Cisco UCS B250-M1 server equipped with two sockets uses a Cisco Extended Memory Technology that allows up to eight DIMMs per channel, resulting in 24 DIMMs per socket. In a two-socket server, this counts towards 48 DIMMs, which if 8-GB-sized DIMMs are used, results in 384 GB of memory per server.

The benefit of a large number of DIMMs is also per-server memory cost reduction. Note that with 48 DIMMs, 2-GB-sized DIMMs can be used to deploy 96 GB of memory per server. The cost of the 2-GB DIMM compared to the 8-GB DIMM is much smaller.

Cisco UCS B-Series Adapters

This topic identifies and describes UCS B-Series mezzanine options.

UCS B-Series Mezzanine Adapters

<div style="background-color: #0070C0; color: white; padding: 5px; border: 1px solid black; margin-bottom: 5px;">Virtualization</div> <p>VM I/O Virtualization and Consolidation</p>  <p>Cisco UCS VIC M81KR Up to 128 vNICs Support for Cisco VN-Link</p>	<div style="background-color: #70AD47; color: white; padding: 5px; border: 1px solid black; margin-bottom: 5px;">Compatibility</div> <p>Minimal Disruption Using Existing Driver Stacks</p>  <p>EMULEX QLOGIC</p> <p>Cisco UCS CNA M71KR (Failover) CNA M72KR (No Failover)</p>	<div style="background-color: #E69A00; color: white; padding: 5px; border: 1px solid black; margin-bottom: 5px;">Cost</div> <p>High-Speed Ethernet Connectivity</p>  <p>Cisco UCS 82598KR-CI No Failover Support</p>
--	--	--

© 2011 Cisco Systems, Inc. All rights reserved.DCUCB v1.0-3-05

A Cisco UCS B-Series blade server can be equipped with different types of mezzanine adapters (network adapters).

Cisco UCS NIC 82598KR-CI

The Cisco UCS 82598KR-CI Network Adapter is made by Intel. It has two 10 Gigabit Ethernet ports to the backplane, which runs native Ethernet. To provide “free SAN access” to the host the operating system implemented, FCoE may be run over the adapter.

Note This card does not provide failover functionality.

The adapter may be used for standard 10 Gigabit Ethernet connectivity to servers that do not need Fibre Channel protocol support, or servers that provide Fibre Channel protocol and FCoE encapsulation in software.

Cisco UCS NIC M51KR-B Broadcom BCM57711

The Cisco UCS NIC M51KR-B Broadcom BCM57711 Network Adapter is a dual-port 10-Gb/s KR Ethernet network adapter mezzanine card with an x8 PCI Express (PCIe) host interface. The KR interface and PCIe host interface interconnect on a single mezzanine connector. Designed specifically for the Cisco UCS blades, the adapter combines offload technology with standard Ethernet functions. Together, these features provide the performance and bandwidth critical to I/O-intensive applications such as virtualization and high-performance computing (HPC).

The TCP Offload Engine (TOE) reduces server CPU utilization and improves application performance; TOE offloads the TCP protocol processing from the server CPU onto the server adapter.

Cisco UCS CNA M61KR-I Intel

The Cisco UCS CNA M61KR-I Intel Converged Network Adapter is a mezzanine card based on the Intel 82599 10 Gigabit Ethernet controller, which is designed to meet the demanding needs of the next-generation data center by providing outstanding features for virtualization; flexibility for LAN, FCoE, and Small Computer System Interface over IP (iSCSI) SAN networking; and proven, reliable performance.

The adapter provides intelligent, hardware-based acceleration that supports native iSCSI initiators and accelerates iSCSI traffic while improving data processing on multicore processor-based servers.

Cisco UCS CNA M71KR

The Cisco UCS CNA M71KR is a Converged Network Adapter (CNA) with two host-side 10 Gigabit Ethernet ports and two Fibre Channel ports to the backplane.

The two network ports can run either native Ethernet and/or FCoE protocols and can be configured for failover. This fabric failover is performed by the adapter ASIC and does not require multipathing software on the host operating system.

The adapter provides increased scalability since the FCoE protocol encapsulation is offloaded from the host and performed in hardware on the Cisco UCS CNA M71KR. For the purpose, the adapter utilizes the Cisco UCS CNA M71KR ASIC, which is a multiplexer designed by Cisco and FCoE protocol offload engine with a 350 MHz 24 K MIPS processor.

There are two versions of this card:

- Cisco UCS CNA M71KR-E Emulex Converged Network Adapter
- Cisco UCS CNA M71KR-Q QLogic Converged Network Adapter

Both adapters support existing proven driver stacks.

Cisco UCS CNA M72KR

The Cisco UCS CNA M72KR is a CNA with two host-side 10 Gigabit Ethernet ports and two Fibre Channel ports to the backplane. The two network ports can run either native Ethernet and/or FCoE protocols.

There are two versions of this card:

- Cisco UCS CNA M72KR-E Emulex Converged Network Adapter
- Cisco UCS CNA M72KR-Q QLogic Converged Network Adapter

Cisco UCS VIC M81KR

The Cisco UCS VIC M81KR is a CNA with dual 10 Gigabit Ethernet ports and dual Fibre Channel ports to the backplane.

The adapter supports up to 128 Ethernet virtual network interface cards (vNICs), or Fibre Channel virtual host bus adapters (vHBAs), running at 600,000 I/O operations per second (IOPs) in both initiator and target mode.

These vNICs are administratively defined, and are instantiated on the server at time of service profile assignment. The BIOS, operating system, and hypervisors see the vNICs as regular PCI Express (PCIe) devices. The architecture of the adapter allows for the presentation of the vNIC both inward toward the CPU and outward toward the network from the perspective of the physical connection to the fabric. Because the existence, identity, and policy are applied from the service profile to the adapter, these devices are known to the BIOS prior to the operating system boot.

This card provides failover between the redundant Ethernet fabrics with no multipathing software required on the host operating system for this functionality.

UCS VIC M81KR Adapter Advantages

- Hypervisor bypass to offload ESX host
- VN-Link technology
- Integration with VMware vSphere vCenter (configuration)
- Use cases
 - VMware server virtualization
 - Database deployment
 - IaaS



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-36

A typical VMware server virtualization environment connects to multiple distinct LAN and SAN interfaces to provide separate connectivity for the VM kernel, service console, and virtual machine (VM) data traffic and shared SAN storage. Therefore, the use of four to eight NICs and two or more HBAs is very common.

The Cisco UCS VIC M81KR adapter allows a user to create these vNICs from a single dual-port 10-Gb/s adapter and apply network policies to each. The vNICs are presented to the VMware ESX server, which can use them as physical NICs for VMs. The personality of a VM can migrate (with vMotion) using the Cisco VN-Link functionality of the Cisco UCS VIC M81KR and Cisco UCS 6100XP Fabric Interconnect.

Cisco VN-Link Technology

Cisco VN-Link technology allows the unique identification of the vNICs and presents them as logical interfaces (which are logically tied directly to the virtual adapters) on the Cisco UCS 6100XP Series Fabric Interconnects. The VN-Link technology also allows administrators to configure policy groups and to include the logical interfaces in this grouping.

VMware Integration

In virtualized environments, the adapter offers close integration with VMware vCenter – if added, moved, or deleted from a VMware ESX server, the network policy and management information (for example, counters) can be controlled through VMware vCenter.

Hypervisor Bypass

The adapter has built-in architectural support for each VM to directly access the adapter hardware, bypassing the hypervisor completely. This capability relieves some of the computing burden on the hypervisor and further improves performance, without sacrificing crucial benefits such as VMware vMotion.

Use Cases

The adapter provides the benefits of reducing the number of physical adapters, simplifying management, and scaling performance for various environments:

- VMware server virtualization
- Demanding database deployments
- Infrastructure as a Service (IaaS) environments

Adapter Characteristics

Model	Part No.	Type	vNIC, vHBA	Failover
Broadcom	M51KR-B	NIC	2,0	OS
Cisco	M81KR-C	CNA	Up to 56	Fabric/OS
Emulex	M71KR-E	CNA	2,2	Fabric/OS
Emulex	M72KR-E	CNA	2,2	OS
Intel	M61KR-I	NIC	2,0	OS
QLogic	M71KR-Q	CNA	2,2	Fabric/OS
QLogic	M72KR-Q	CNA	2,2	OS


© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-37

The table summarizes the characteristics of the B-Series adapter options.

Connectivity High-Availability Overview

- Requirements
 - Fabric interconnect cluster
 - Cisco UCS CNA M71KR or Cisco UCS VIC M81KR with fabric failover configuration
- Operating system-level teaming can be used to additionally scale the bandwidth



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-38

To deploy fabric failover for server connectivity, the following requirements must be met:

- Cisco UCS must be deployed in a cluster.
- The server must be equipped with either Cisco UCS CNA M71KR or Cisco UCS VIC M81KR adapters.
- Fabric failover must be configured for LAN adapters.

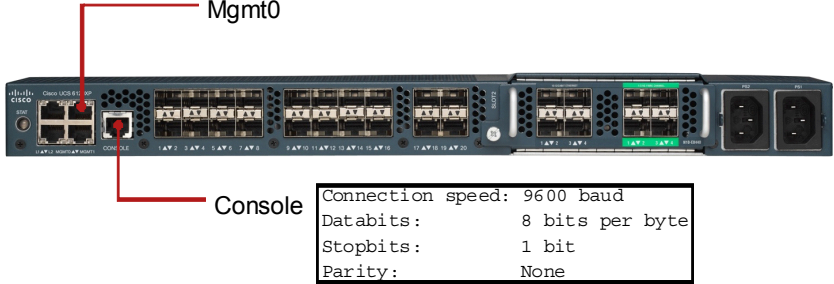
The LAN-based hardware failover eliminates the need for operating system-level teaming for the failover purpose. Still, the operating system-level teaming can be used to scale the bandwidth when at least two LAN adapters are used.

Cisco UCS B-Series Management

This topic identifies and describes Cisco UCS Manager and its options.

Management Connectivity

- Console port
- Out-of-band Ethernet management – Mgmt0
 - Configurable management IP address



Connection speed:	9600 baud
Databits:	8 bits per byte
Stopbits:	1 bit
Parity:	None

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-40


The console terminal of the Cisco UCS 6100 Series Fabric Interconnect can be accessed using the RJ45-based serial console port on the rear of the switch. The console port uses standard settings.

Out-of-band management is performed using the 10/100/1000 Ethernet management port—`mgmt0`—located on the rear of the Cisco UCS 6100 Series Fabric Interconnect. Before using this port for the management access, you must assign an IP address to it. Afterwards, the system can be accessed via assigned IP address using Telnet, Secure Shell (SSH), or Cisco UCS Manager GUI from a remote workstation.

Cisco UCS Management Networks

- All system management traffic carried via reserved VLANs
- System reserved VLAN IDs from 3968 to 4048

Name	VLAN
Data Center Operating System Networks	4042
Adapter Management Networks	4043
Adapter Management Infrastructure Networks	4044
Cisco UCS Manager PXE Networks	4047



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-41

The Cisco UCS has VLANs 3968 to 4048 reserved for management purposes. These VLANs are used in Fabric A as well as in Fabric B.

Data Center Operating System Network—VLAN 4042


Used to provide connectivity between the Data Center Operating System (DCOS) (Cisco UCS Manager) instance in each fabric interconnect, Fabric A and B CMCs, and the management entities on the mezzanine cards. The network traffic includes Layer 2 protocols, such as the ones used to bootstrap the IOM and the virtual interface card protocol used to support Network Interface Virtualization.

Adapter Management Network—VLAN 4043

Used to provide connectivity between the DCOS (Cisco UCS Manager) instance in each fabric interconnect and the management entities on the mezzanine cards for the purposes of allocating resources, defining identities, and monitoring adapter status. Layer 2 protocols are run by both Cisco UCS Manager instances for general adapter functionality, and Layer 3 (IP) is used for remote login to the adapter operating system. IP addresses are algorithmically determined by Cisco UCS Manager and passed to the adapter during bootstrapping.

Cisco UCS Management Architecture

- All management performed via Cisco UCS Manager
 - Runs on Cisco UCS 6100XP Fabric Interconnect
 - Access via fabric interconnect management IP address on Mgmt0 port
 - Management traffic uses dedicated VLANs on Ethernet downlinks



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-42

The Cisco UCS 6100 Fabric Interconnect switches are the single management point in Cisco UCS deployments. Fabric interconnect switches are equipped and running Cisco Nexus Operating System (Cisco NX-OS) software within which the Cisco UCS management software runs.

The Cisco UCS management system is accessed using a management IP address configured on the fabric interconnect management port mgmt0. The management port provides CLI, GUI, and XML access and presents the remote KVM functionality for server blades.

The management traffic between the Cisco UCS 6100XP Fabric Interconnect switches and other system components is exchanged in-band via dedicated VLANs. There is no special management connectivity between the system components.

- The management tasks that are related to the chassis are passed via redundant serial links to the appropriate device.
- IOM-related management tasks are passed via inter-integrated circuit links (I2C) to the appropriate IOM chipset.
- The management tasks that are related to the server blades are passed via dedicated redundant Ethernet links to the BMC of the server blades.

CMC

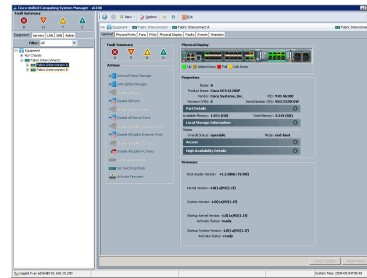
The CMC is a key part of management infrastructure since it collects status data from the IOM using the IPMI protocol over the I2C serial bus and sends that information to the management node using the Ethernet server link. The CMC controls the power supply and fan speeds, and serves as a proxy for the Cisco UCS Manager to the blade servers.

BMC

The BMC communicates “lights-out” status information, KVM, and other status information to the CMC for the individual blade server.

Cisco UCS Manager Overview

- Embedded device manager
 - Single, highly available management domain for all Cisco UCS components
 - Management tasks—discovery, inventory, configuration, monitoring, diagnostics, statistics collection, fault detection, auditing
 - Scalable—single instance manages up to 320 server blades
 - RBAC for administration
 - Service profiles for server provisioning
 - Support for auto-discovery and configuration
- No external management server required



© 2011 Cisco Systems, Inc. All rights reserved.

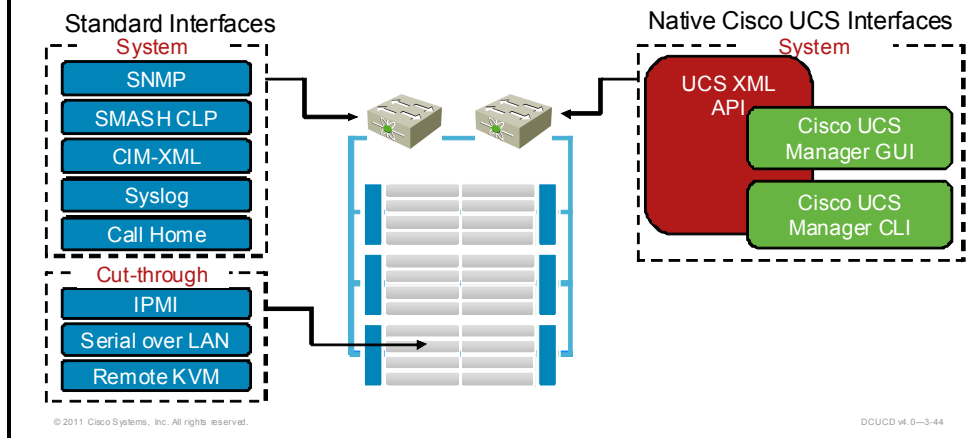
DCUCD v4.0-3-43

Cisco UCS Manager is embedded device-management software that manages the system from end to end as a single logical entity through an intuitive GUI, a CLI, or an XML API.

Cisco UCS Manager implements policy-based management of the server and network resources in Cisco UCS. Network, storage, and server administrators all participate in creating policies in their areas of domain expertise. Policies are used in service profiles, allowing Cisco UCS Manager to fully configure the servers, adapters, and fabric extenders and appropriate isolation, quality of service (QoS), and uplink connectivity on the Cisco UCS 6100 Series Fabric Interconnects.

Management Protocol Support

- Support for multiple management interface types
- APIs for integration with standard data center management protocols



Cisco UCS Management Model

Cisco UCS uses a Management Information Model (MIM) where each real-world resource—either physical or logical like fabric interconnects, chassis, blade servers, and so on—is a managed object, that is, an abstract of a real-world object. These managed objects are automatically created for each real-world resource in order for the Cisco UCS management interfaces to be able to monitor and manage them.

Cisco UCS Management Interfaces

There are multiple protocols and management interfaces supported by Cisco UCS, and they fall into two categories:

- Systemwide management interfaces, which are used for systemwide Cisco UCS management and monitoring (from management of blade servers and adapters to the management of the fabric interconnect)
- Cut-through management interfaces, which are used to provide direct blade server access and thus management and monitoring of the individual blade servers

Systemwide Native Management Interfaces

The native Cisco UCS management interfaces terminate in the Cisco UCS Manager. Using this type of interface, all system management access is handled by the Cisco UCS Manager, that is, each request is queued, interpreted, checked against the user's privileges and executed by the Cisco UCS Manager.

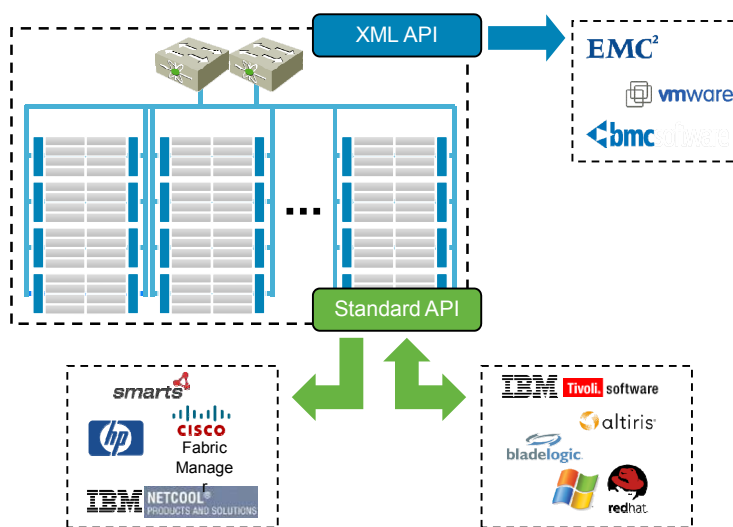
Systemwide Standard Management Interfaces

The systemwide management interfaces can use standard or native Cisco UCS protocols. The standard interfaces typically cover only the subset of the Cisco UCS management model.

Cut-Through Management Interfaces

The cut-through standard management interfaces and protocols are used to access the server blades directly via a unique external management IP address omitting the Cisco UCS Manager application.

Management Interface Integration



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-45

The available standard API and the XML API show Cisco UCS as a flexible resource pool with management that can be integrated with the existing management and diagnostic tools.

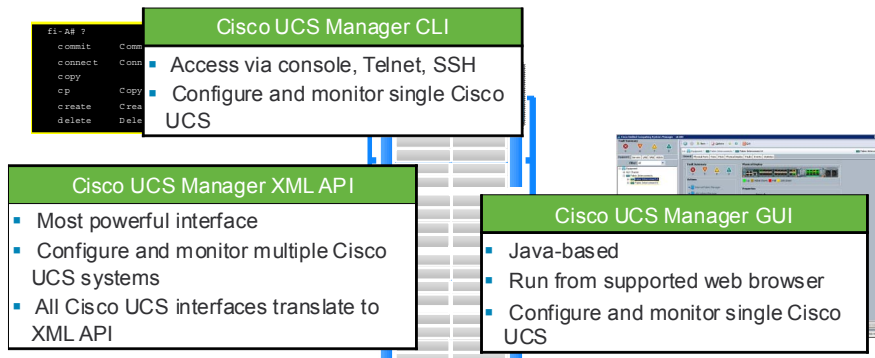
The management applications from EMC, BMC Software, or VMware can utilize the XML API to query anything in the system—either a physical state or logical definition. Apart from that, the application could also modify the Cisco UCS parameters and configuration. Cisco UCS, on the other hand, does not provide the application or operating system provisioning.

The standard API enables management applications to execute limited amount of operation. Some applications can only read the Cisco UCS parameters.

All the management interfaces (except cut-through) are internally translated into native XML API. The vital internal component of the embedded Cisco UCS manager is the Data Management Engine (DME), which in essence handles all the requests. Such centralized requests processing guarantees uniform enforcement of RBAC and other Cisco UCS Manager mechanisms.

Native Cisco UCS Management Interfaces

- Support for multiple management interface types
- APIs for integration with standard data center management protocols



Cisco UCS native management interfaces provide management access to a whole Cisco UCS solution—from adapters to blade servers to fabric interconnects. The Cisco UCS management is thus centralized for any of the resources, and includes the following devices:

- Fabric interconnects
- Software switches for virtual servers
- Power and environmental management for blade enclosure and blade servers
- Configuration and firmware updates for Ethernet and Fibre Channel adapters
- Firmware and BIOS settings for servers

The management interfaces provide full read and write capabilities, that is, the system can be configured and not only monitored. The result of tasks performed in one interface also reflects changes in another interface.

Since Cisco UCS abstracts the physical server hardware, the management interfaces support virtual and physical servers also. The state information is virtualized in service profiles that can be applied to any blade server in the system. The native Cisco UCS management interfaces are:

- Cisco UCS Manager CLI
- GUI
- XML API

The access is provided via the Cisco UCS system management IP address. Native UCS interface options enable the administrator to perform almost all tasks from any interface.

Cisco UCS Manager CLI

The CLI management interface allows the administrator to access, configure, and monitor the Cisco UCS system. The CLI is available from the console port or through remote access using the Telnet or SSH protocols. The CLI management interface uses the management object model and is transactional.

Cisco UCS Manager GUI

The GUI application is Java-based and thus can be started and run from any supported web browser with proper Java Runtime Environment (JRE) installed. The interface is a full-featured user interface for Cisco UCS management and monitoring.

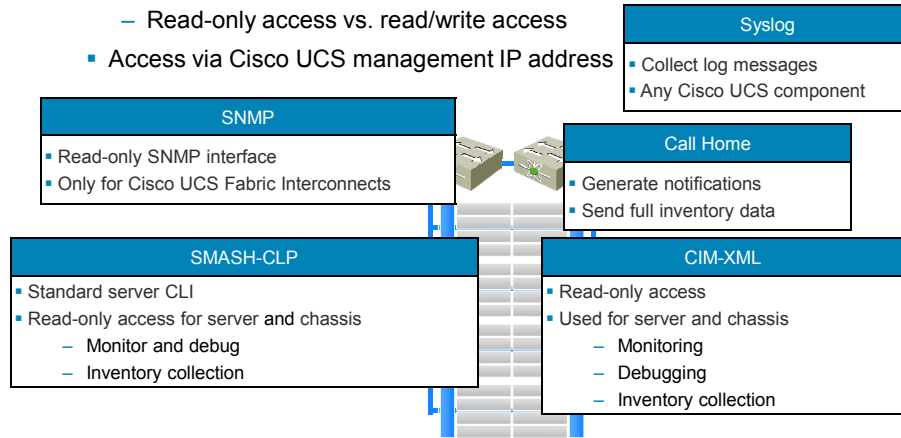
Cisco UCS Manager XML API

The XML API interface offers the highest level of integration and interaction with the Cisco UCS due to its nature—it is a generic, content-driven, hierarchical interface.

The XML API supports event subscription, which enables a subscribing client, that is, monitoring application, to receive all Cisco UCS events or state changes (the application—subscriber—must have the right privileges).

Standard Cisco UCS Interface Options

- Provide access to a whole Cisco UCS solution
 - Standardized protocols
 - Read-only access vs. read/write access
- Access via Cisco UCS management IP address



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-47

The standard Cisco UCS management interfaces utilize standardized protocols. The access to Cisco UCS is provided via the Cisco UCS management IP address in the same manner as the native Cisco UCS management interfaces. The interfaces can be used for read-only or read/write access.

Systemwide standard management interfaces that are supported by Cisco UCS are as follows:

- Simple Network Management Protocol (SNMP)
- Systems Management Architecture for Server Hardware Command Line Protocol (SMASH-CLP)
- Common Information Model XML (CIM-XML)
- Syslog
- Call Home

SNMP

Cisco UCS supports SNMP v3 read-only monitoring of the Cisco UCS Fabric Interconnect as defined in the MIBs for the Cisco NX-OS 4.0.(0)(N1)(1a). There are approximately 60 different MIBs defined (for example, authentication, authorization, accounting [AAA], Call Home, STP, VLAN, and Fibre Channel device aliases).

SMASH-CLP

SMASH is a suite of specifications that deliver industry-standard protocols to manage elements of a data center. The SMASH-CLP specification provides an interface to heterogeneous servers independent of server hardware, operating system, or network protocol method. It is a standard method for local and remote management of server hardware using out-of-band communication. SMASH is developed by the Desktop Management Task Force Server Management Working Group (DMTF SMWG).

With SMASH-CLP-enabled products, users can execute common operations (such as system power on and off, system log display, boot order configuration and text-based remote console) using the same commands across different vendor platforms.

CIM-XML

CIM-XML is a Web-based enterprise management (WBEM) protocol that uses an XML encoding encapsulated in HTTP to exchange CIM data and methods. CIM-XML is defined and managed by the DMTF. CIM is structured into these distinct layers:

- Core model: An information model that captures notions that are applicable to all areas of management.
- Common model: An information model that captures notions that are common to particular management areas, but independent of a particular technology or implementation. The common areas are systems, applications, networks and devices. The information model is specific enough to provide a basis for the development of management applications. This schema provides a set of base classes for extension into the area of technology-specific schemas.
- Extension schemas: Represent technology-specific extensions of the common model. These schemas are specific to environments, such as operating systems and servers.

Syslog

Syslog is a traditional UNIX logging facility that can be configured to send messages from different programs (or “facilities”) of varying severity (that is, info, warn, crit, emerg) to the console, to a file, or to an external host (also known as “loghost”). Syslog contains several built-in facilities such as kern (or kernel), print (that is, print server), and mail (that is, SMTP server). Users can also configure their own custom facilities for applications that do not fall neatly into one of the predefined facilities. These facilities are named local0, local1, and so on.

In Cisco UCS, you can configure syslog to report to a local destination or a remote destination. Local destination configurations include console, monitor, and file. Remote destination configurations (for up to three log hosts) include level, hostname, and facility.

Call Home

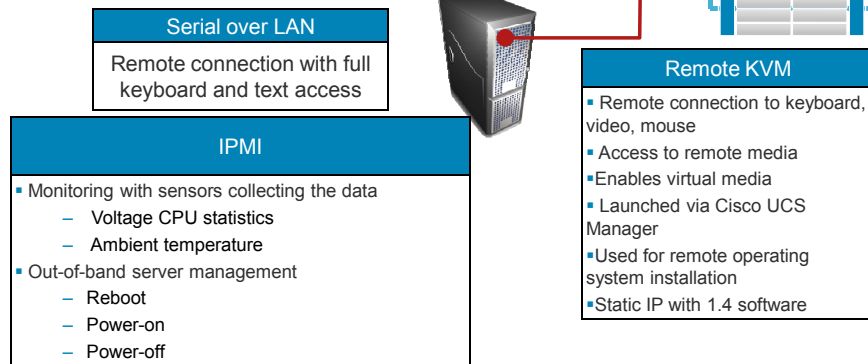
Cisco UCS Manager uses the Call Home functionality to create email messages in either XML or standard text format and sends the messages to predefined recipients when certain events occur. Configurable profiles define the following:

- Criticality levels (for example, debug, warning, and fatal)
- Email format (XML or text)
- Email recipients

In Cisco UCS, you can enable Call Home policies to generate notifications on thermal, voltage, power, identity, Field Replaceable Unit (FRU), and equipment events. You can also send full inventory data by pushing a button or schedule it to be sent on a periodic basis.

Cut-Through Interface Options

- Provide access to a single server
 - Standardized protocols
- Access via unique external management IP address
 - Assigned by the Cisco UCS Manager



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-48

The blade servers present the computing resources in Cisco UCS. These resources sometimes need to be managed individually. The management access to the blade server is provided with the dedicated unique external management IP address—this IP address is given per blade server by the Cisco UCS Manager from the management IP address pool.

Cisco UCS supported cut-through management interfaces are as follows:

- IPMI
- Serial over LAN (SoL)
- Remote KVM

IPMI

IPMI operates independently of the operating system and allows administrators to manage a system remotely even in the absence of an operating system or even if the monitored system is powered off, but connected to a power source. IPMI can also function after the operating system has started. IPMI prescribes only the structure and format of the interfaces as a standard, while detailed implementations may vary.

IPMI can send out alerts via a direct serial connection, a LAN or a SoL connection to a remote client. System administrators can then use IPMI messaging to query platform status, to review hardware logs, or to issue other requests from a remote console through the same connections.

The IPMI consists of the BMC and other satellite controllers. The satellite controllers within the same chassis connect to the BMC via the system interface called Intelligent Platform Management Bus/Bridge (IPMB) — an enhanced implementation of I2C. The BMC connects to satellite controllers or another BMC in another chassis via Intelligent Platform Management Chassis (IPMC) bus/bridge. It may be managed with the Remote Management Control Protocol (RMCP), a specialized wire protocol defined by this specification.

A FRU holds the inventory (such as vendor ID, manufacturer, and so on) of potentially replaceable devices. A Sensor Data Records (SDR) repository provides the properties of the individual sensors present on the board. For example, the board may contain sensors for temperature, fan speed, and voltage.

SoL

SoL is a mechanism that enables the input and output of the serial port of a managed system to be redirected via an IPMI session over IP (that is, using Telnet or SSH).

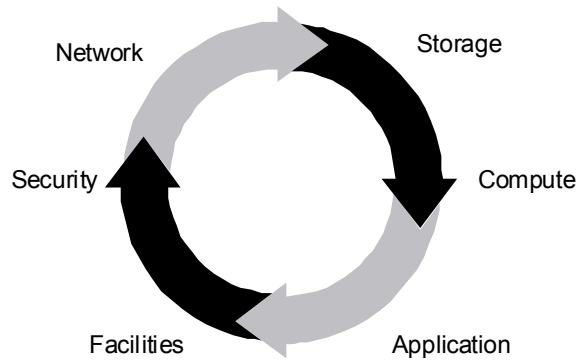
In Cisco UCS, the serial ports on the blades are not connected to a traditional serial port interface. To allow users to access applications on the servers via the serial port, I/O of the serial port is redirected to the BMC and made available to external users on the management network. For example a user wishing to access a blade server that is running Linux via the serial port can use SSH to go to the management network IP address associated with that blade and log in. On the blade server, the login will be seen as coming through the serial port.

In Cisco UCS, this is configured through a SoL policy, which defines the following:

- Service enabled or disabled
- Baud rate for terminal sessions

Cisco UCS Management Benefits

- Cisco UCS represents end-to-end architecture
- Enables better collaboration among administrative roles



© 2011 Cisco Systems, Inc. All rights reserved.


DCUCD v4.0-3-49

Implementing the Cisco Unified Computing solution has implications in the organizational structure and data center management. If the management roles were less interconnected before the Cisco UCS deployment, the collaboration is elevated once Cisco UCS is implemented.

The data center management benefits come from the Cisco UCS end-to-end architecture since the system is used to configure and provision compute, LAN, and SAN data center aspects.

Cisco UCS Manager High Availability

- Cisco UCS Manager accessible via cluster IP address
 - Floating IP address for automatic failover
 - Management port in both fabric interconnects must be connected
- Cisco UCS Manager runs in two instances
 - Active instance on primary fabric interconnect
 - Standby instance on secondary fabric interconnect (subordinate)
 - Database and state information replicated over cluster links
 - Split-brain scenarios prevented by architecture itself
 - Automatic process restart upon failure



The Cisco UCS Manager application is accessible via the Cisco UCS cluster IP address. Thus the management port (mgmt0) of each fabric interconnect should be connected to the same Layer 2 network to facilitate failover and failback of the management IP address. Each fabric interconnect should connect to only one side of each chassis.

Cisco UCS Manager Controller

The Cisco UCS Manager controller is a distributed application running on both primary and subordinate Cisco UCS Managers. Each instance is represented by a unique ID (the same as the node ID). The Cisco UCS Manager controller is implemented as a separate process. The Cisco UCS Manager controller decides which Cisco UCS Manager components should run in primary or subordinate mode. All the Cisco UCS Manager processes are always started on both nodes.

Cisco NX-OS System Manager

All Cisco UCS Manager processes including the Cisco UCS Manager controller, on both primary and secondary nodes, are started and managed by the Cisco NX-OS system manager. This fact implies that all the processes on both nodes are always started. This approach enables control of the Cisco UCS Manager processes from the Cisco UCS Manager controller while exploiting the process-monitoring features that are provided by the Cisco NX-OS system controller.

Cisco UCS Manager Disaster Recovery

- Cisco UCS Manager configuration can be saved:
 - Full database backup—includes also physical hardware state (service profile association)
 - All configuration backup
 - System configuration backup only
 - Logical configuration backup only
- Backup operation can be scheduled and automated (FTP, SCP, FTP server)
- Configuration backups produce XML-format backups
- UCS Manager recovery process can incorporate different backup types

© 2011 Cisco Systems, Inc. All rights reserved.

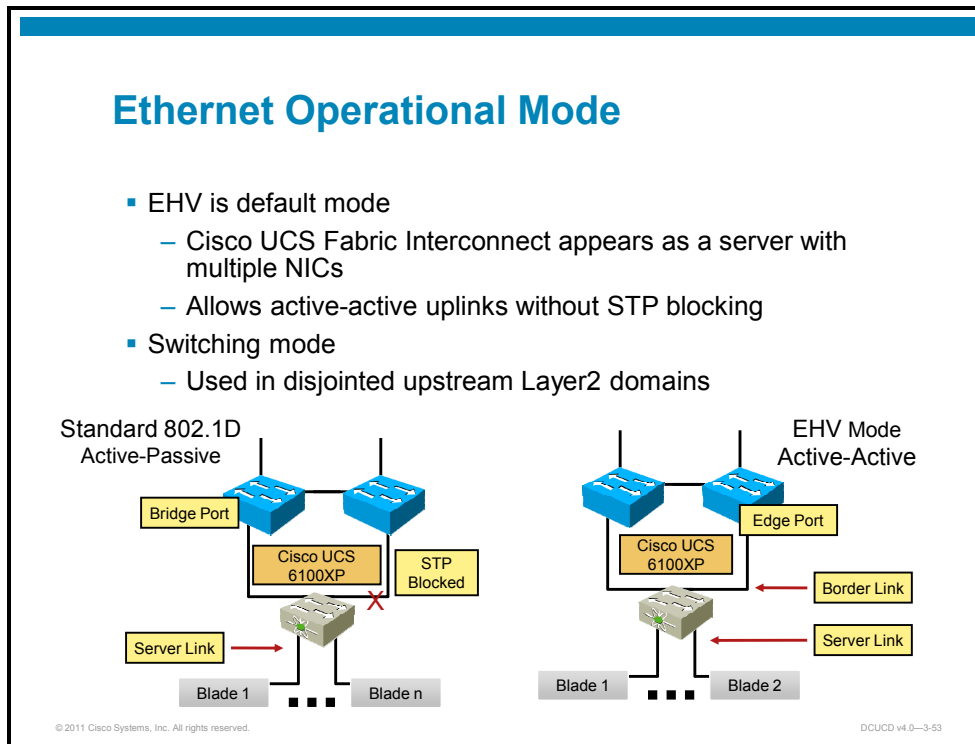
DCUCD v4.0—3-51

All configuration within the Cisco UCS environment is stored within the Cisco UCS Manager database. Administrators will want to ensure that proper care is taken to maintain backups of this database. Backups can be made to local storage on the fabric interconnect and/or to a remote location, and can be scheduled to reduce administrative overhead.

An administrator can perform full state backup, which includes hardware state in a form of service profile association, or various configuration backups.

Cisco UCS B-Series LAN Connectivity

This topic identifies and describes Cisco UCS B-Series LAN connectivity.



EHV Mode

The Cisco UCS Fabric Interconnect operates in an EHV mode by default (also known as end host virtualization). In the EHV mode, Cisco UCS appears to the external LAN as an end station with multiple adapters.

There are two types of links in the EHV operational mode:

- Server links
- Border links

Border links are Cisco UCS uplinks and can be in a form of a single link or aggregated in a channel. When operating in the EHV operational mode, Cisco UCS Fabric Interconnect does not participate in a Spanning Tree Protocol (STP) topology. Instead, the following is used to achieve loop free topology:

- Border links must connect to Layer 2 network
- Traffic forwarding between border links is denied

The benefit of running a Cisco UCS fabric interconnect in the EHV mode is that the LAN STP topology is simplified and the STP domain size is reduced. Second, since no links are blocked by the STP, the active-active approach utilizes all redundant links to a Layer 2 network.

Switching Mode

In a normal LAN topology STP protocols takes care of the loops in the topology. It does so by disabling some of the links. Thus, the underlying network infrastructure is not fully utilized.


If desired, the fabric interconnects can be set to operate in traditional Ethernet switching mode.

This method introduces the need for STP to avoid broadcast loops. Despite operating in this manner, not all typical Cisco switch options are available. Thus, this deployment method is typically not recommended.

The switching mode is used in case of disjointed upstream Layer 2 domains.

EHV

- Server interface pinned to border interface
 - Server-to-network follows pinned uplink
 - Server-to-server traffic locally switched
 - Network-to-server forwarded to server if arrives on pinned uplink
 - Server traffic on any uplink except pinned is dropped



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-54

In the EHV mode, each server link is pinned to exactly one border link. The pinning logic equally distributes server links to various border links.

Note that the server-to-server traffic is locally switched, while server-to-network traffic goes out on the pinned border link. To achieve local switching the MAC addresses within the chassis are learned.

Network-to-server unicast traffic is forwarded to the server only if it arrives on a pinned border link. The Reverse Path Forwarding (RPF) check is performed to verify that. Server traffic that is received on any border link except the pinned border link is dropped (déjà vu check).

EHV (Cont.)

- MAC address learning
 - Internal MAC addresses are learned
 - Learning is disabled on border interfaces
- Traffic forwarding
 - To server based on destination MAC
 - Learned MAC addresses age out when server link goes down or server is moved (repinned)

Traffic	Received on	Sent to
Learned Unicast	Server or uplink port	Server
Unknown Unicast	Server port	Pinned uplink and all active server ports
Unknown Unicast	Uplink port	Dropped
Broadcast/Unknown Multicast	Uplink port	All server ports
Broadcast/Unknown Multicast	Server port	Pinned uplink and all active server ports
Known Multicast	Server or uplink port	As per IGMP snooping

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-55

The MAC address learning in EHV mode is as follows:


- Learning is disabled on border links—network MAC addresses are never learned.
- Learning is enabled on server links—traffic to server is forwarded based on destination MAC address.

Learned MAC addresses never age unless the server link goes down or is deleted; in that case, server MACs can move (in the event of repinning).

The table in the preceding figure explains where certain traffic is sent when received.

LAN Pinning

- To identify specific uplink port or port channel on one or both fabric interconnects
- Administrative pinning—using pin groups
 - Pin traffic from specific blade network adapter to specific uplinks with service profile
- Dynamic pinning (default)—not using pin groups
 - Cisco UCS automatically chooses uplink for adapter
 - Could choose incorrectly if LAN uplinks are on different Layer 2 domains



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-56

Cisco UCS allows the use of a mechanism called a pin group.


With pin groups, traffic from a specific blade server adapter is tied to a particular uplink port or port channel on each fabric interconnect. This configuration is achieved via the service profile configuration.

If pin groups are not used in service profiles, Cisco UCS automatically chooses an uplink port or port channel for the adapter on the blade associated with each profile.

If you need to force that specific blade to use a certain uplink port, the pin group can be applied via the service profile to achieve the forced use. This might be the case when the uplinks are connected to different Layer 2 domains (VLANs). If left to default, the Cisco UCS Manager does not know what the correct uplink for that specific blade is.

LAN Connectivity—VLANs

- Initially configured with VLAN 1 (default) only
- VLANs configured globally per Cisco UCS cluster
 - Can be created per Fabric A or B or both
- No participation in VLAN Trunking Protocol



Cisco UCS LAN connectivity consists of these elements:

- VLANs: Defined to separate traffic from different segments
- Uplink connectivity: Connects the Cisco UCS system to the external LAN network
- Server connectivity: Connects the blade servers via uplink ports to the external LAN network

VLANs

VLANs in a Cisco UCS system are defined to connect the Cisco UCS with the external LAN networks. The VLANs defined in Cisco UCS should match the VLANs defined on the other side of the uplink ports since they are trunked using IEEE v802.1q.

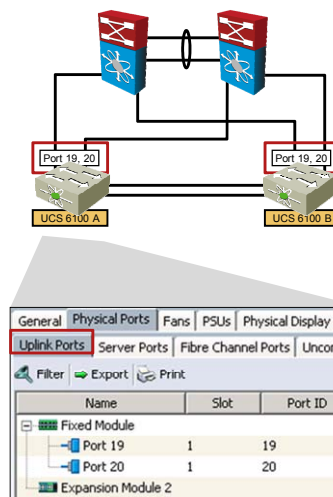
Initially, a new Cisco UCS system is configured with default VLAN 1 only, which is not recommended.

VLAN configuration in Cisco UCS is done globally for the whole system, but can be later assigned to specific service profiles. Since the Cisco UCS in a redundancy setup consists of Fabric A and Fabric B, the VLANs can also be designed per individual fabric or independent of the fabric (which enables the failover option if later used in a service profile).

When created in the Cisco UCS Manager, the VLAN needs the name, which is later used in other configuration parts and the valid IEEE 802.1Q VLAN ID.

LAN Connectivity—Uplink Ports

- Carry Ethernet traffic only
- Uplink port allowed; VLAN list adjusted automatically per configuration
- Uplink switch must trunk all VLANs used in service profiles on that interface
- Port channel can scale bandwidth
 - Must match upstream switch configuration
- Uplink ports in EHV mode
 - Appear to be host with multiple NICs
- Uplink ports in switching mode
 - Appear to be Ethernet switch



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-58

Uplink Ports

Uplink ports in a Cisco UCS are those physical ports on the fabric interconnects that are dedicated for the connectivity to a LAN device external to Cisco UCS (for example, for the connectivity to Cisco Nexus 7000).

These ports can be either from the fixed port range or from the expansion module if present.

The ports carry VLAN traffic only and are configured as 802.1Q trunk ports to carry VLAN traffic for all the VLANs used in the service profiles for a particular fabric. The allowed VLAN list is adjusted automatically based on Cisco UCS VLAN configuration.

Depending on the Cisco UCS configuration, an uplink port carries VLANs that belong to the fabric the uplink port is part of (A or B) and those VLANs that are not fabric-dependent are defined globally in the LAN cloud.

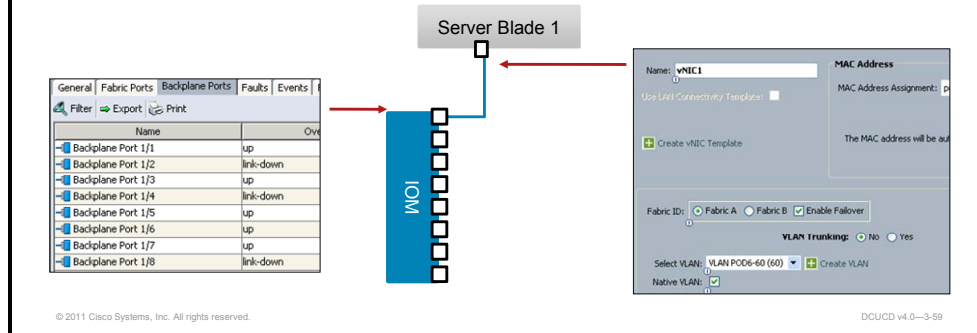
Port Channels

Uplink ports bandwidth can be scaled using port channels, which use the LACP 802.3ad protocol. The configuration of the port channel must match the other side (VLANs being trunked must be the same).

Port channel can be configured using uplink ports from a single Cisco UCS 6100XP Fabric Interconnect in a cluster, that is, from the same fabric.

LAN Connectivity—Server Ports

- Server blade to IOM connectivity—vNIC
 - Configured via associated service profile only (no direct configuration of a physical server port)
 - Attributes—VLAN, trunking characteristics, and redundancy settings
 - Native VLAN—traffic sent untagged



Server Blade to IOM Server Ports—vNIC

The communication between the server blade and the IOM is governed by the service profile configuration. This configuration includes vNICs, which define how a server connects to the LAN network.

A vNIC has a number of parameters associated with it, among which the most important parameters are VLAN, trunking characteristic (trunk vs. nontrunk interface), and redundancy setting, which defines whether upon primary fabric failure the communication fails over to the second (for example, from Fabric A to Fabric B or vice versa).

Server Port—vNIC

- Cisco UCS Ethernet and CNA adapters
 - Two vNIC per physical adapter can be created (one per fabric)
 - Failover only with generation 1 CNA
- Cisco UCS VIC M81KR adapter
 - NIC virtualization supports multiple vNIC creation
 - Fabric A or B with failover
 - Number of vNICs depends on the IOM-fabric interconnect uplinks

Uplinks	vNICs per Adapter
1	13
2	28
4	58

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-60

Cisco UCS CNA M71KR/M72KR, UCS 82598KR-CI Adapter

You can create up to two vNICs using either the Cisco UCS CNA M71KR/M72KR or the Cisco UCS 82598KR-CI adapters.

With all but M71KR, you must match the physical setting (first adapter goes to Fabric A, second to Fabric B), and you cannot choose failover.

With Cisco UCS CNA M71KR, each vNIC is associated with a particular fabric, but you can enable failover.

Cisco UCS VIC M81KR Adapter

The Cisco UCS VIC M81KR supports NIC virtualization either for a single operating system or for VMware vSphere. The number of virtual interfaces supported on an adapter depends on the number of uplinks between the IOM and the fabric interconnect, as well as the number of interfaces in use on other adapters sharing the same uplinks.

LAN Adapter Failover

- IOM fabric port failure
 - MAC address is automatically repinned to the second fabric interconnect
- Fabric interconnect uplink port failure
 - Other uplinks available on the same fabric interconnect
 - Dynamic repinning to the uplink on same fabric interconnect occurs
 - No uplinks available on the same fabric interconnect
 - Switch vNIC with failover to the other fabric interconnect

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-61

Within Cisco UCS, the LAN adapter can handle either IOM fabric port failure or fabric interconnect uplink port failure.

When IOM fabric port failure occurs, the LAN adapter is automatically repinned to the second fabric.

When fabric interconnect uplink port failure occurs, two scenarios are possible:

- When other uplinks on the same fabric interconnect are available, dynamic repinning to the uplink of the same fabric interconnect is performed.
- When no other uplinks on the same fabric interconnect are available, fabric failover to the second fabric occurs.

Cisco UCS 1.4 Network Enhancements

Feature Details

Direct Connect NAS w/o Switch mode

- VLAN and QoS Support

Full virtual machines support for fabric failover

SPAN support


PVLAN Isolated access support

Higher limits

- 1024 VLANs
- 6000 Logical ports
- 2000 Virtual Interfaces

Fabric Extender Transceiver support

Management Interface monitoring and failover



© 2011 Cisco Systems, Inc. All rights reserved.

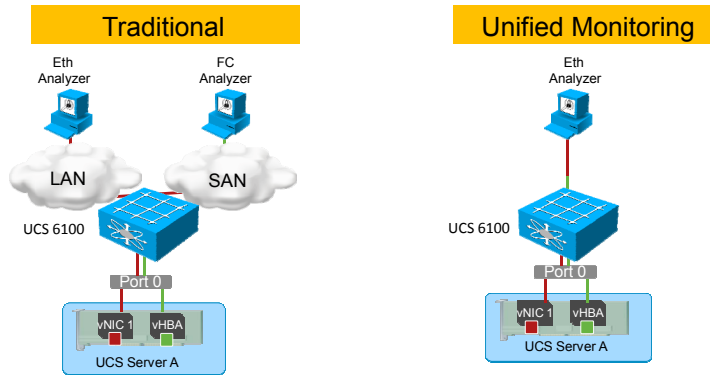
DCUCD v4.0-3-62

The 1.4 version of the software brings several benefits from the network perspective:

- Direct Connect NAS without Switch mode with VLAN and QoS Support
- Full virtual machines support for fabric failover
- Switched Port Analyzer (SPAN) support
- PVLAN Isolated access support
- Higher limits:
 - 1024 VLANs
 - 6000 logical ports
 - 2000 virtual interfaces
- Fabric Extender Transceiver support
- Management Interface monitoring and failover

Unified Monitoring with Cisco UCS Manager 1.4

Monitor Fibre Channel and Ethernet traffic with a single analyzer



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-63

Cisco UCS Manager version 1.4 brings support for unified monitoring.

The traditional monitoring requires the following:

- Dedicated Fibre Channel analyzer
- Separate analyzers for Fibre Channel and Ethernet
- Specialized in-line device (tap) which is typically expensive

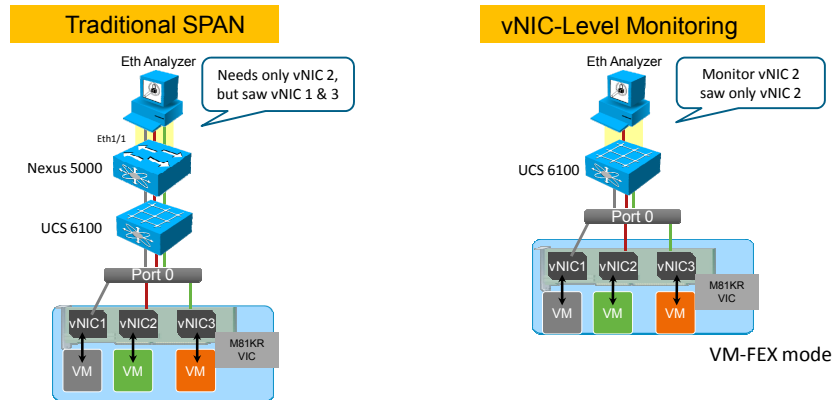
Unified monitoring brings the ability to monitor Fibre Channel and Ethernet traffic using a Ethernet analyzer.

- Single analyzer to monitor Fibre Channel and Ethernet
- Eliminates expensive Fibre Channel analyzer

The SPAN functionality on Cisco UCS is limited to a total of four local SPAN sessions per Cisco UCS system, two on each Cisco UCS 6100.

vNIC and vHBA Level Monitoring

vNIC-level monitoring and troubleshooting capability



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-64

Another improvement is the ability to perform monitoring on a vNIC or vHBA level.

A traditional SPAN port and monitoring has the following drawbacks:

- Complex connectivity
- Inability to monitor specific VM traffic
- Lack of visibility, that is, seeing traffic for all vNICs

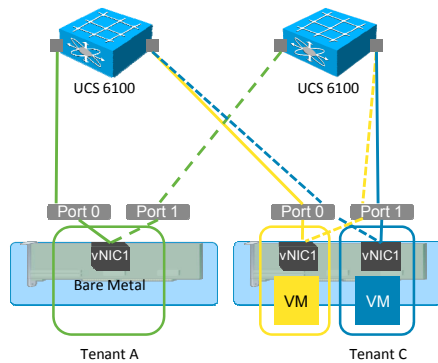
The vNIC- level monitoring provides the following benefits:

- Simple connectivity
- Ability to monitor specific VM traffic
- Granular visibility

PVLAN Support in Cisco UCS Manager 1.4

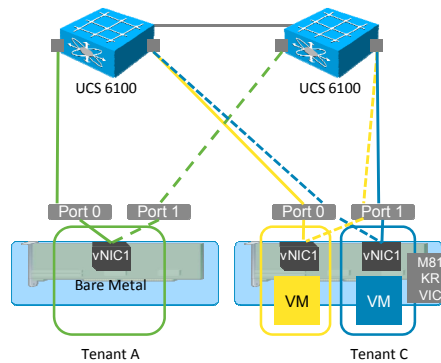
No. of VLANs Used = 3

VLAN Deployment



No. of VLANs Used = 1

PVLAN Deployment



Traditionally, the Layer 2 separation is achieved with VLANs. In such designs, the following applies:

- One VLAN or more is used per tenant/application/function
- VLAN scalability is limited by platforms

With PVLANS, the Layer 2 separation can be achieved using a single VLAN combined with PVLAN functionality. In such designs, the following applies:

- L2 separation is achieved with isolated PVLAN
- Single isolated VLAN can be used for all tenants/applications
- The level of isolation is the same as the one of the standard VLAN
- Maximum number of VLANs is no longer scalability limiting factor

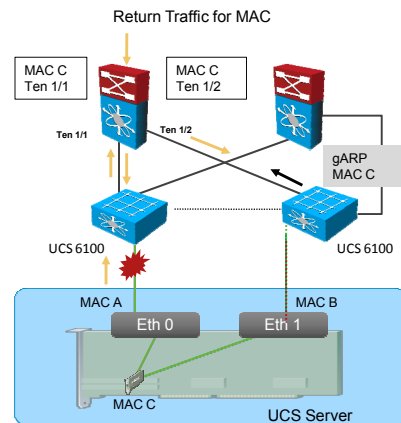
The Cisco UCS implementation of PVLANS has the following characteristics:

- Only isolated access PVLAN is supported
- One isolated VLAN is supported per primary
- The same vNIC cannot carry PVLAN and regular VLANs at the same time
- Community PVLAN is not supported
- Promiscuous port is not supported
- VMware DVS does not support native VLAN on the trunk, thus the isolated PVLAN on Cisco UCS does not work with VMware DVS

Fabric Failover and Synchronization

■ Fabric Failover

- When active path fails, failover to standby fabric occurs
- 6100 updates path changes to upstream switches via gratuitous ARP of vNIC MAC upon active link failure



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-66

Fabric failover is the advanced functionality available with Cisco UCS, with the following characteristics:

- Fabric provides redundant path for each vNIC
- Hardware based active/standby failover mechanism is implemented for Ethernet traffic
- Host operating system is unaware of failure and recovery
- Unlike operating system NIC teaming, redundancy is provided with a single interface
- Fabric failover load-balances traffic on per vNIC basis
- Cisco UCS m81KR VIC provides up to 58 vNICs
- Each vNIC maps to one of the fabric for active data path


When the active path fails, failover to standby fabric occurs and UCS 6100XP updates path changes to the upstream switches via gratuitous ARP of vNIC MAC upon active link failure. This is performed only for the static vNICs with the following pitfalls:

- VM MAC behind a vSwitch are not updated to upstream switches
- For silent VMs (one-way communication), does not provide path update

Fabric Failover and Synchronization (Cont.)

▪ Fabric Synchronization

- Maintains synchronized MAC address tables between 6100XP
- 6100 updates path changes to upstream switches via gratuitous ARP also for VM MACs upon active link failure



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-67

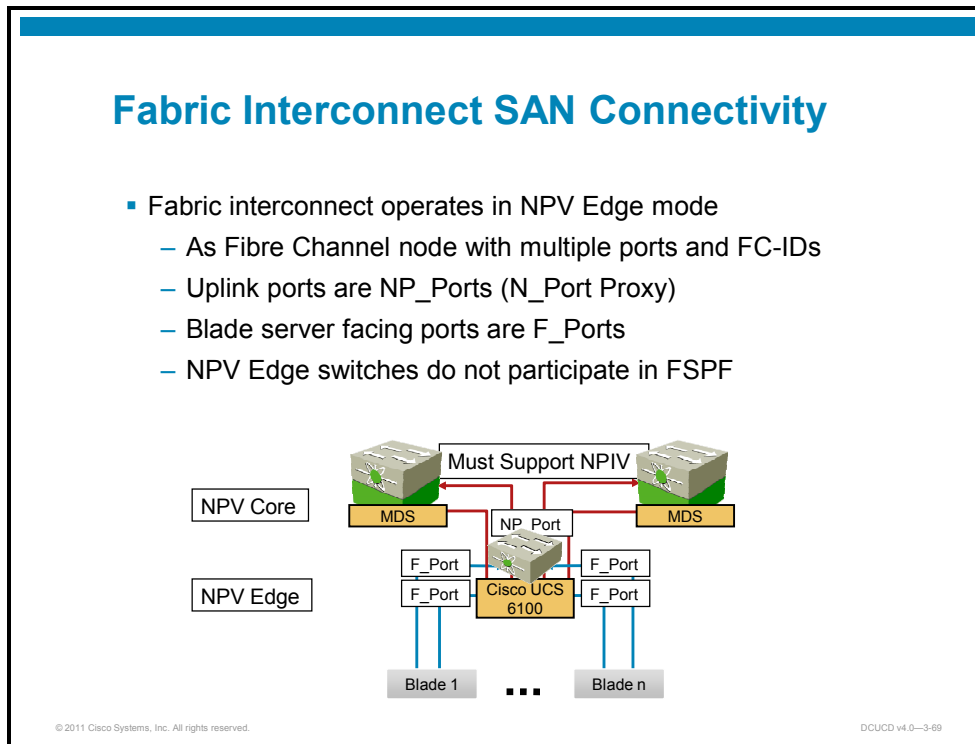
With fabric synchronization, the MAC address tables are synchronized between Cisco UCS 6100 XP. The UCS 6100XP updates path changes to upstream switches via gratuitous ARP not only for vNIC MACs but also VM MACs behind a vSwitch upon active link failure.

The feature is always enabled since version 1.4 of the software.

The feature is specifically useful for Hyper-V deployments, since Hyper-V officially does not support NIC teaming or bonding thus the UCS hardware based failover provides redundancy. With fabric synchronization, the VMs running on the Hyper-V in Cisco UCS benefit from MAC table synchronization.

Cisco UCS B-Series SAN Connectivity

This topic identifies and describes Cisco UCS B-Series SAN connectivity.



The Cisco UCS fabric interconnects operate in the N_Port Virtualizer (NPV) edge mode.

The upstream Fibre Channel switch (for example, MDS) must thus support and be enabled with the N_Port ID Virtualization (NPIV) feature, which allows multiple Fibre Channel IDs (FC-IDs) to be assigned a single node port (N_Port).

In the NPIV topology, there are two types of interfaces:

- **Server interface:** The server-facing interface is either physical Fibre Channels, or virtual Fibre Channel interfaces operating in a fabric port (F_Port) mode.
- **Border interface:** Border interfaces are network-facing and always operate in a proxy N_Port mode.

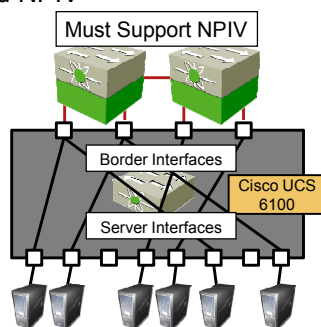
There is no local switching of the Fibre Channel traffic on the Cisco UCS 6100XP; all packets are forwarded to the NPV core switch.

The fabric login (FLOGI)-related processing is relayed in software (FLOGI, fabric discover [FDISC], and corresponding link service accept [LS_ACC], link service reject [LS_RJT], and so on) to the same uplink interface.

Note Every uplink can be connected to different Fibre Channel switches and VSANs.

N_Port Virtualizer

- Each server interface pinned to one border interface
- Pinning logic distributes server interfaces between border interfaces (round-robin)
- All traffic follows the pinned port
- All traffic passed to upstream device for switching
- Supports nested NPIV




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-70

Cisco UCS 6100XP in NPV edge mode pins each server link to exactly one border link. The pinning logic load-balances server links to various border links, while all traffic is forwarded to the upstream SAN device for switching.

Pinning and Load Balancing

- Server interfaces pinned to uplink interfaces (not FLOGI)
- Traffic on server interface
 - Sent to the pinned interface
 - No forwarding lookup
 - NPV switch does not participate in FSPF
 - Binding check performed to verify frame SID is on the right server interface
 - Prevent address spoofing
- Traffic on border interface
 - Forwarding lookup is performed per frame DID
 - DID points to the server interface
 - Packets are discarded on miss



© 2011 Cisco Systems, Inc. All rights reserved.

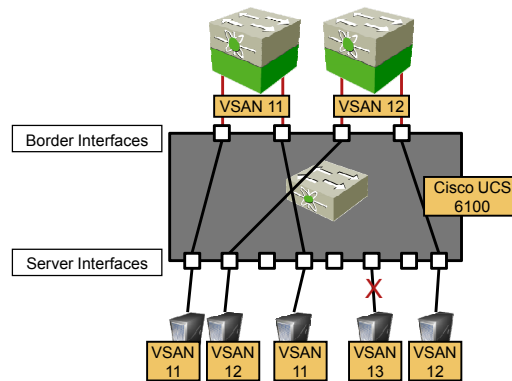
DCUCD v4.0-3-71

With NPV edge mode, each downstream device (server or blade server) is pinned to an uplink port based on a round-robin algorithm.

The Cisco UCS 6100XP switch in NPV edge mode no longer handles FLOGI login requests or makes routing decisions using Fabric Shortest Path First (FSPF). Instead, these operations are passed to the upstream switch known as the NPV core switch. The NPV core switch utilizes NPIV in order to interpret multiple logins from the same port.

Pinning and Load Balancing with VSANs

- Pinning based on VSANs
 - Server interface pinned to border interface with same VSAN
 - Server interface kept down if no interface with same VSAN available



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-72

With VSANs, the NPV edge mode pinning takes into account uplink port VSAN also. The server is pinned to an uplink port based on the uplink interface VSAN membership (still in a round-robin fashion).

SAN Pinning

- To identify a specific Fibre Channel uplink port on one/both fabric interconnects
- Administrative pinning—using pin groups
 - Pin traffic from specific blade SAN adapter to specific uplink using service profile.
- Dynamic pinning (default)—not using pin groups
 - Cisco UCS automatically chooses uplink for SAN adapter based on VSAN

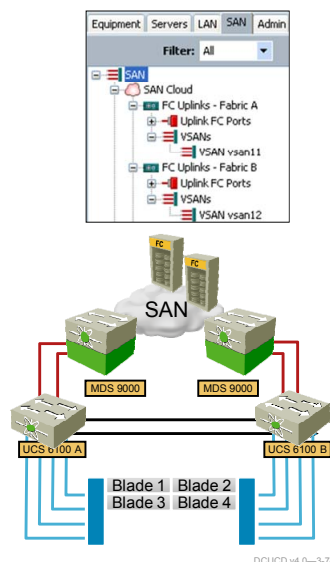
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-73

If required, the SAN pinning can be administratively defined with pin groups. The configuration is applied via a service profile.

SAN Connectivity—VSANs

- Initially configured with VSAN 1 (“default”) only
- Configured globally per Cisco UCS Cluster
 - Can be created per Fabric A or B or both
 - Requires FCoE VLAN ID
 - Associated with the uplink port(s)
 - Single VSAN per uplink and server port (before version 1.4)
- Supported by MDS FC switches
- Used internally by UCS if connected to non-MDS switches



Cisco UCS SAN connectivity consists of these elements:

- VSANs: Defined to separate traffic for different Fibre Channel fabrics
- Uplink connectivity: Connect the Cisco UCS system to the external SAN network
- Server connectivity: Connect the blade servers via uplink ports to the external SAN network

VSANs

VSANs in a Cisco UCS system are defined to connect the Cisco UCS with the external SAN networks. The VSANs defined in Cisco UCS should match the VSANs defined on the other side of the uplink port. Initially, a new Cisco UCS system is configured with default VSAN 1 only, which is not recommended.

VSAN configuration in Cisco UCS is done globally for the whole system, but can be assigned later to specific service profiles. Since the Cisco UCS in a redundancy setup consists of Fabric A and Fabric B, the VSANs can also be designed per individual fabric or independent of the fabric.

When created in the Cisco UCS, the VSAN needs the following:


- Name, which is later used in other configuration parts
- VSAN ID
- FCoE VLAN ID, used to carry the VSAN traffic over from the server blade to the IOM to the fabric interconnect

If you want to support any VSAN, it needs to be configured globally into Cisco UCS Manager, and then it can be associated with a particular vHBA. The default VSAN is preconfigured into Cisco UCS Manager and is automatically chosen as the default connectivity for each vHBA.

While Cisco MDS switches use VSANs, other vendors have not yet implemented VSAN technology within their switching architectures. Cisco UCS continues to use VSANs internally to isolate physical fabrics. Each uplink connects to one physical fabric and is mapped internally to a VSAN number. Through this method, the same Cisco UCS fabric interconnect could be connected via uplink to multiple, physically separate Fibre Channel fabrics without causing those fabrics to merge. All Fibre Channel services would be kept isolated by VSANs and no inter-VSAN routing is possible.

SAN Connectivity—FCoE VLAN

- SAN traffic from servers carried via FCoE in dedicated VLANs
- Must not overlap the LAN VLANs




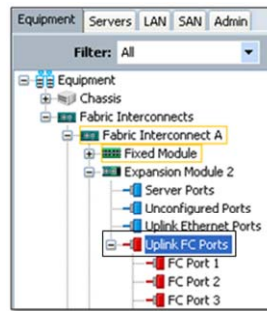
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-75

Within Cisco UCS, an FCoE VLAN ID is used for VSAN. The FCoE VLAN ID should not overlap with regular VLANs used for LAN connectivity.

SAN Connectivity—Uplink Ports

- Connectivity to external SAN devices
 - Carries Fibre Channel traffic only
 - One VSAN per uplink port (before version 1.4)
- Fibre Channel ports on expansion modules



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-76

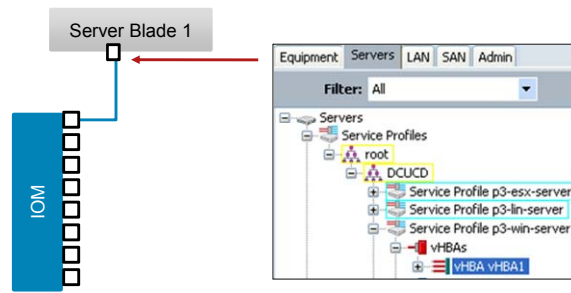
Uplinks in a Cisco UCS are physical Fibre Channel ports on the Cisco UCS Fabric Interconnect expansion modules used for the connectivity to a SAN device external to Cisco UCS (for example, for the connectivity to MDS 9000). Note that a single port carries traffic for one VSAN only.

Depending on the Cisco UCS configuration, an uplink port carries VSANs that belong to the fabric it is part of (A or B).

A single uplink port carries traffic for one VSAN only, that is, is connected to one physical fabric and is mapped internally to a VSAN number. The same Cisco UCS Fabric Interconnect can be connected via uplinks to multiple, physically separate Fibre Channel fabrics without causing those fabrics to merge. All Fibre Channel services are thus kept isolated using VSANs and no inter-VSAN routing is possible.

SAN Connectivity—Server Ports

- Defined in a service profile with vHBA:
 - Assigned to a single VSAN
 - VSAN and properties assigned dynamically via service profile
- VSAN used internally to isolate fabrics even if uplinks connected to non-MDS switches.



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-77

A server SAN port is configured as a vHBA that corresponds to VSAN. The concept is similar to the one used in LAN—VLAN/vNIC.

For the server to be connected to the SAN, the service profile must be configured with the vHBA where a VSAN must be selected. A vHBA configuration is applied to the Fibre Channel interface on the physical blade when the service profile is associated with the blade server.

Before the VSAN is associated with vHBA, it must be configured globally in Cisco UCS Manager.

VSANs are supported on Cisco MDS switches but not by other vendors, even though Cisco UCS internally still uses VSANs to distinguish between and isolate the fabrics.

Server Port—vHBA

- Cisco UCS CNA adapters
 - Two vHBA per physical adapters can be created (one per fabric).
 - No hardware failover support—must use host-based multipathing.
- Cisco UCS VIC M81KR adapter
 - Supports multiple vHBA creation.
 - Fabric A or B without failover.
 - Number of vHBAs depends on the IOM – Fabric Interconnect uplinks.

Uplinks	vNICs per Adapter
1	13
2	28
4	58

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-78

The VSAN/vHBA logic is mostly analogous to the VLAN/vNIC logic discussed in the previous lesson, but they are not identical. You must have a Cisco UCS Manager vHBA object in a profile to specify connectivity for a Fibre Channel adapter on a blade to which the profile will be associated.

If you want to support any VSAN, it needs to be configured globally into Cisco UCS Manager, and then it can be associated with a particular vHBA. The default VSAN is preconfigured into Cisco UCS Manager and is automatically chosen as the default connectivity for each vHBA.

The Cisco UCS CNAM71KR/M72KR mezzanine adapter supports up to two vHBA interfaces connected to either of the fabrics.

The Cisco UCS VIC M81KR mezzanine adapter supports creation of multiple vHBA interfaces that are presented to the host operating system as separate physical interfaces.

In either case, the hardware-based failover, as with vNIC, is not supported for the vHBA. To be able to deploy multiple paths, a host-based multipathing solution must be used.

SAN Adapter Failover

- No hardware-assisted fabric failover
 - Traditional SAN host multipathing must be used for fabric failover.
- Fibre Channel uplink failure
 - Other uplinks in same VSAN available on the same fabric interconnect:
 - Interface goes down/up to trigger server FLOGI
 - Dynamic repinning to the uplink on same fabric interconnect
 - If failed uplink comes back, no repinning to avoid disruption
 - No uplinks available on the same fabric interconnect
 - Server port brought down

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-79

Within the Cisco UCS, SAN-level fabric failover is not available. Thus the traditional SAN host-based multipathing must be used to achieve that level of high availability.


If there are multiple FC uplinks on the same fabric interconnect for the same VSAN, the following occurs:

- Interface goes down and up to trigger server FLOGI.
- Dynamic repinning to the uplink of the same fabric interconnect in the same VSAN occurs.
- Upon the failed uplink restoration, repinning is avoided to prevent disruption.

If there are no Fibre Channel uplinks in the same fabric interconnect for the same VSAN, the server port is brought down.

Direct Connection of Storage (Cisco UCS 1.4)

- Support for NetApp and EMC DAS
- Ability to turn on/off Fibre Channel limited switching
 - Zoning configuration not supported
 - Zoning may be inherited from upstream switch
- Ethernet and Fibre Channel switching modes are independent
- Allow direct-connected NAS devices in EHV mode
 - New port type allows customers with direct connect storage to leave switch mode off
 - New “appliance” port type for Ethernet appliances and direct connect storage.
- Ability to connect FCoE devices (single hop) to Cisco UCS



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-80

Cisco UCS 1.4 supports directly connected storage which can be Fibre Channel, FCoE, or NFS.

The NFS type of appliance is denoted as Ethernet appliance which can be any specialized device for use on a Ethernet network (for example Network Attached Storage [NAS], iSCSI, security appliances, or Nexus 1010) that does not run STP. Note that this type of the port should not be used for switch connectivity to avoid traffic loops.

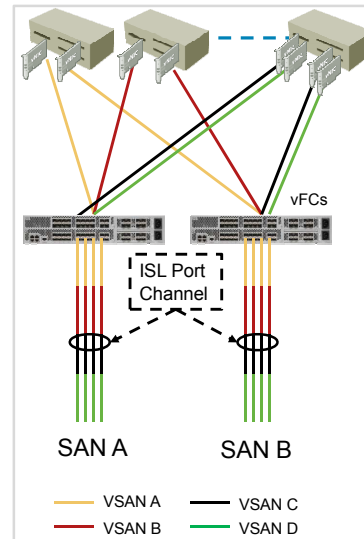
Even though new types of storage attachment are available, the same rules apply—that is, there is no hardware-assisted fabric failover for storage due to the complexity of operation. A traditional SAN host multipathing must be used for fabric failover.

The direct attached storage brings several benefits for the customer:

- Cisco UCS 6100XP switch mode no longer needed for such environments
- External Fibre Channel switch may still be required, but no longer in data path (if zoning has to be implemented)

Fibre Channel PortChannels and Trunking (Cisco UCS 1.4)

- Up to 16 Fibre Channel ports aggregated in a single port channel
 - Different combination of Fibre Channel ports from different expansion modules on the FI can be placed on the same port channel
 - In case of port speed mismatch, port channel forces port speed to highest commonly supported speed
- VSANs can be trunked over the port channel
- VSAN trunking and port channel supported for both NPV and switch mode FI operation



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-81

Cisco UCS version 1.4 also removes the limitation of a single VSAN per single uplink Fibre Channel port.

There are two features that can be used to scale the SAN implementation:

- Fibre Channel port channels
 - Up to 16 Fibre Channel ports can be aggregated together for a single port channel
 - Different combination of Fibre Channel ports from different expansion modules on the UCS 6100XP can be placed on the same port channel
 - In case of port speed mismatch, the port channel forces port speed to highest commonly supported speed
- Fibre Channel port trunking
 - VSANs can be trunked Fibre Channel ports
 - VSAN trunking and port channel supported for both NPV and switch mode FI operation

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco UCS is a system that encompasses compute, storage, and network resources.
- Cisco UCS 6100XP Fabric Interconnect switches with Cisco UCS 2104XP IO Modules provide network and storage connectivity to the blade servers.
- The Cisco UCS 5108 Chassis provides power, cooling, and connectivity to the blade servers.
- The Cisco UCS B-Series Blade Servers architecture encompasses two CPUs, hot-swappable hard disk, memory, and I/O capability.
- The Cisco UCS VIC M81KR Adapter is optimized for virtualization.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-82

Summary (Cont.)

- Cisco UCS uses reserved VLANs for internal management communication.
- Cisco UCS Manager is a single point of management for compute, SAN, and LAN.
- Cisco UCS Manager resides on the Cisco UCS fabric interconnects.
- IOM is connected with fabric ports to fabric interconnect server ports.
- EHV mode allows all uplinks to be active.
- LAN pin groups can be used to administratively select uplink ports for a server.
- VLANs are carried within 802.1Q trunks.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-83

Summary (Cont.)

- VLANs and VSANs can be defined in both fabrics or in a single fabric.
- Port channels aggregate multiple Ethernet links into a single logical link.
- Number of vNICs depends on system connectivity with virtualized adapter.
- Cisco UCS 6100XP operates in NPV Edge mode in SAN.
- Dynamic SAN pinning selects border interface based on VSAN.
- VSAN requires FCoE VLAN ID.
- No hardware failover is available for vHBAs.

Sizing the Cisco UCS B-Series Solution

Overview

This lesson discusses how to assemble a Cisco UCS B-Series solution for a given set of requirements.

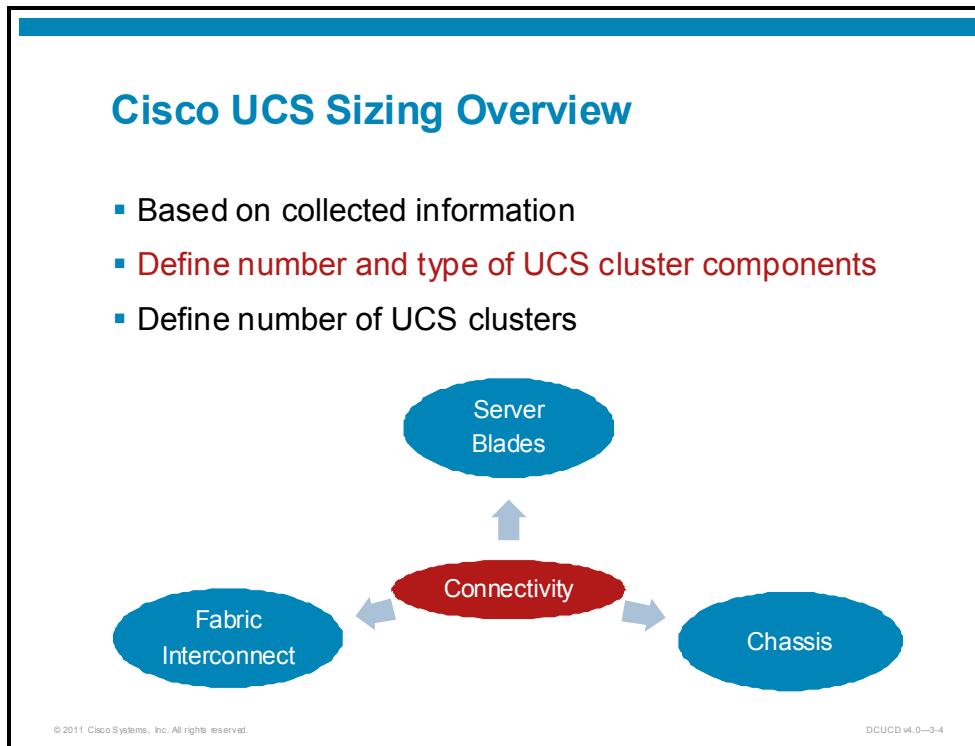
Objectives

Upon completing this lesson, you will be able to assemble a Cisco UCS B-Series solution. This includes the ability to meet these objectives:

- Identify and describe solution requirements
- Describe how to propose B-Series hardware per requirements
- Identify and describe how to create a B-Series solution BOM

Analyzing Requirements

This topic identifies and describes solution requirements.



The sizing process for Cisco UCS deployment is used to determine the appropriate components that are needed to build the Cisco UCS per the requirements that were gathered in a design workshop.

The sizing process can be divided into three major categories:

- **Identify and size Cisco UCS server classes:** With this step, the individual server type is identified and sized per requirements.
- **Identify and define Cisco UCS chassis classes:** With this step, the server blades are put into the chassis—that is, the chassis classes are identified and properly sized—for the purpose of identifying the correct number of server uplinks.
- **Identify and size Cisco UCS fabric interconnect clusters:** With this step, the fabric interconnects (20- versus 40-port switch), the expansion modules, and the number and type of chassis that are connected to the individual fabric interconnect clusters are identified.

Analysis Output

Basis for solution sizing and server deployment

Parameter	Server Type 1	Server Type 2	Server Type 3
CPU	1x 2-core at 1.2 GHz	1x 4-core at 2.4GHz	2x 4-core at 2.4 GHz
Memory	8 GB	16 GB	24 GB
LAN Connectivity	2x 1 Gigabit Ethernet	6x 1 Gigabit Ethernet	6x 1 Gigabit Ethernet
SAN Connectivity	n/a	2x 2 Gb Fibre Channel	2x 4 Gb Fibre Channel
Boot Media	Local disk	Local disk	SAN
LAN Throughput	0.5 Gb/s	0.8 Gb/s	1.5 Gb/s
SAN Throughput	n/a	1 Gb/s	2 Gb/s

System Name	Make/Model	Capacity										Utilization									
		Processors	Memory	Disk	Network	Physical			Processor		Memory			Disk		Network					
		Count	Speed (MHz)	Size (MB)	Size (GB)	Count	Speed (MB/sec)	Back (Units)	Weight (lb.)	Power (W)	Thermal (BTU/hr)	% Used	Queue per CPU	% Used	Cache (MB)	File %	Page %	Paging (Trans/sec)	I/O (MB/sec)	I/O (MB/sec)	
Susable Systems																					
3	SPR0000	1,000	2,790	2,560	146.68	2	2,000	2	47.10	400	1,475	1.51	0.00	33.07	233.40	0.20	22.02	10.69	0.25	0.0	
3	SPR0016	1,000	2,784	4,608	146.68	2	2,000	2	47.10	400	1,475	3.20	0.00	36.24	214.01	0.31	89.66	27.32	0.49	0.0	
3	SPR0017	1,000	2,784	4,608	146.68	2	2,000	2	47.10	400	1,475	3.66	0.01	36.56	196.41	0.31	75.51	25.91	0.39	0.0	
3	SPR0018	1,000	2,790	2,560	146.68	2	2,000	2	47.10	400	1,475	3.56	0.01	31.40	179.26	0.30	31.87	17.73	0.52	0.0	
3	SPR0019	1,000	2,796	2,560	146.68	2	2,000	2	47.10	400	1,475	7.03	0.01	34.06	304.23	0.40	654.75	171.79	3.02	0.0	
3	SPR0017	1,000	2,795	4,608	146.68	2	2,000	2	47.10	400	1,475	1.72	0.00	40.59	271.07	0.31	5.40	8.52	0.09	0.0	
3	SPR0015	1,000	2,795	4,992	159.70	2	2,000	2	47.10	400	1,475	10.85	0.03	54.52	185.97	6.55	102.78	100.84	5.07	0.0	

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-5

When designing a Cisco UCS B-Series solution for an existing environment, it is vital to examine the analysis output to determine proper requirements for the processing, memory, connectivity, and storage requirements.


Sizing Cisco UCS Cluster—Connectivity

Connectivity aspects

- Server blade to IOM
- IOM to fabric interconnect
- Fabric interconnect Ethernet and Fibre Channel uplinks

Design the following parameters

- Throughput and oversubscription
- Redundancy



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-6

The UCS B-Series sizing is influenced by two factors—the required throughput and the allowed oversubscription.

Both factors should be observed from different perspectives

- The individual server blade
- Individual chassis and cumulative requirements of the server blades installed in the chassis
- Individual fabric interconnect and cumulative requirements of the chassis connected to the switch, and the upstream LAN and SAN connectivity requirements.


Redundant Connectivity

Dual fabric design

- Two fabric interconnects and IOMs

Determine number of

- Server downlinks
- LAN uplinks
- SAN uplinks



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-7

The Cisco UCS B-Series connectivity architecture typically follows the dual fabric design. Such design can be used to either achieve high availability (with failover on the Cisco UCS level or on the operating system level) or to achieve more throughput by directing traffic from some servers to fabric A and from other servers to fabric B. It is also possible to direct traffic to both fabrics, combining redundancy and higher throughput.

Logical Adapters—vNICs

Server downlinks options—1, 2, 4

Cisco UCS VIC M81KR adapter

- Number of vNICs depends on server downlinks quantity

1x IOM, 8x Server Blades

Server Downlinks	vNICs per Mezzanine
1	13
2	28
4	58

2x IOM, 8x Server Blades

Server Downlinks	Adapters per Mezzanine
1	26
2	56
4	116

© 2011 Cisco Systems, Inc. All rights reserved.

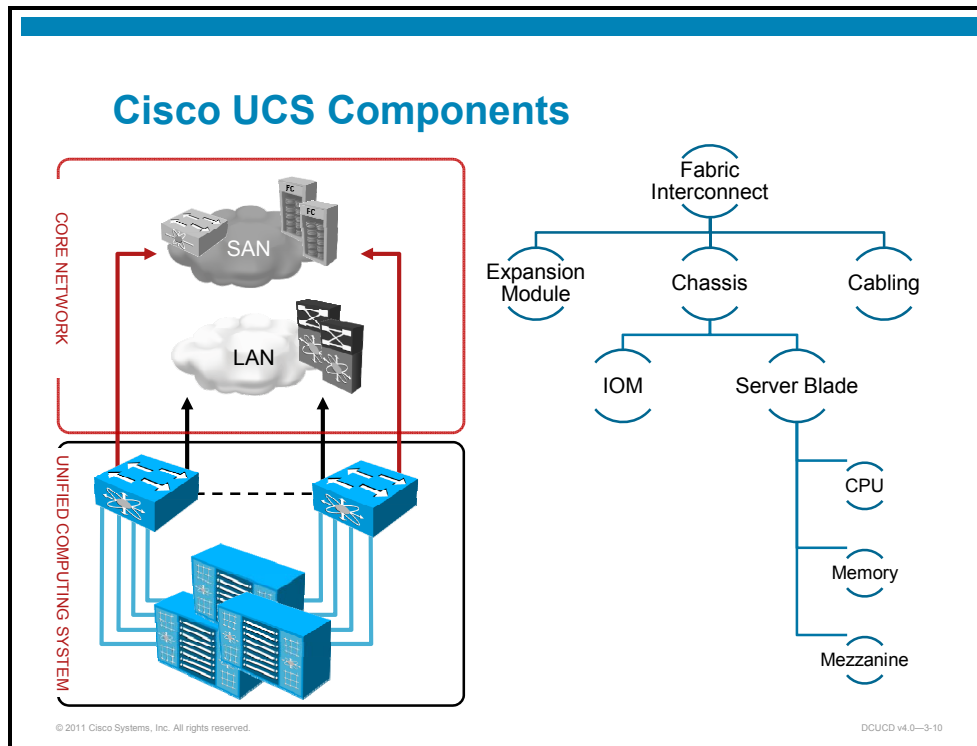
DCUCD v4.0-3-8

When determining the number of uplinks from the I/O modules (IOMs) of the individual chassis, the important factor is the number of adapters (vNIC and/or vHBA) that are required by the operating system or hypervisor level. When more than two vNICs and two vHBAs are required, the Cisco UCS B-Series should be equipped with a Cisco UCS VVIC M81KR mezzanine, which supports multiple vNIC and vHBA adapters.

The maximum number of adapters supported is governed by the number of chassis uplinks. Therefore, even if the throughput requirement could be managed with a single chassis uplink, the requirement of having 14 vmnics in VMware ESX would dictate using two chassis uplinks.

Sizing the B-Series Solution

This topic describes how to propose B-Series hardware per requirements.



When sizing the Cisco UCS B-Series the designer must select the following:


- Type and quantity of Cisco UCS 6100XP fabric interconnects:
 - Number of additional port licenses required
 - Number and type of small form-factor pluggable (SFP) and/or small form-factor pluggable plus (SFP+) modules
 - Cabling
 - Expansion module type and quantity—when connecting to Fibre Channel SAN, the expansion module is required
- Type and quantity of Cisco UCS 5108 chassis:
 - Number of IOMs
 - Number of individual IOM uplinks
 - Server blade population scheme
- Type and quantity of Cisco UCS server blades:
 - CPU, memory, disk drive, disk controller, mezzanine options

The process of selecting the components can be either bottom-to-top (that is, starting with server blades, moving to the chassis, and finally selecting the fabric interconnects) or top-to-bottom (starting by selecting the fabric interconnects, then the chassis, and finally the server blades).

Sizing Cisco UCS Cluster—Server Blades

Design server blades (type and quantity)—select:

- CPU type and quantities
- Memory DIMM quantities and size
- Network adapter—mezzanine
- LAN and SAN connectivity
- Number of interfaces
- Redundancy
- Local disk



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-11

The Cisco UCS server blade sizing design process has two substeps:

- First, the Cisco UCS server blade classes are identified. A server blade class has certain properties such as required resources (CPU type, CPU speed, number of CPUs, memory size, and so on) and connectivity requirements (required LAN and/or SAN bandwidth, redundancy level, number of adapters, and so on).
- Second, the quantity of individual Cisco UCS server blades per server blade class has to be identified.

When the Cisco UCS server blade class is designed, the following sizing criteria are used:

- The type of Cisco UCS server blade, that is, either half- or full-sized. Which one to use is governed by the amount of memory required per blade. One full-sized blade has a memory capacity of up to 384 GB per blade.
- The type, speed, number of cores per processor, and number of CPUs per server blade. Currently both half- and full-sized blades can be equipped with up to two processors.
- The total memory size and individual DIMM type and size. A half-sized blade supports up to 96 GB of memory, whereas a full-size blade supports up to 384 GB of memory. They also support different DIMM sizes from 2 GB onwards. Using smaller DIMMs decreases memory cost per blade.
- The type and number of mezzanine adapters. Cisco UCS server blades support three different types of mezzanine adapters: 10 Gigabit Ethernet, 10 Gigabit Ethernet with Fibre Channel over Ethernet (FCoE) support, and 10 Gigabit Ethernet with FCoE and virtualization support.

Example—Server Blade Sizing

High-Performance Blade Cisco UCS B250-M2

Component	Quantity
Intel Xeon X5680 (130 W)	2
Cisco UCS VIC M81KR	2
Memory (96 GB)	24 x 4 GB DIMM

Standard Performance Blade Cisco UCS B200-M2

Component	Quantity
Intel Xeon X5650 (95 W)	2
Cisco UCS VIC M81KR	1
Memory (24 GB)	12 x 2 GB DIMM

Energy Performance Blade Cisco UCS B200-M2

Component	Quantity
Intel Xeon L5520 (60 W)	2
Cisco UCS M71KR-Q	1
Memory (12 GB)	6 x 2 GB DIMM

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-12

An example result of the Cisco UCS server blade sizing process, seen in the figure, gives three Cisco UCS server blade types with three memory size options.

High-Performance Blade Class

This Cisco UCS blade class provides performance-seeking highest functionality and optimal server return on investment (ROI).

Standard Performance Blade Class

This Cisco UCS blade class provides a mix of performance, values, and advanced features.

Energy Performance Blade Class

This Cisco UCS blade class provides a more energy-efficient server, targeted at applications and environments that are not compute-intensive but still require their own physical box.

Memory Options

When selecting the Cisco UCS blade class, a memory configuration that is best suited for the application must be selected also.

Sizing Cisco UCS Cluster—Chassis

Define Cisco UCS chassis (type and quantity)

- Populate with servers (number of servers per chassis)
- Connectivity requirements
- 10 GE server ports for **chassis connectivity** (IOM uplinks)
- 10 GE uplink ports for **LAN connectivity**
- Fibre Channel uplink ports for **SAN connectivity**
- **Redundancy** requirements



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-13


The second step in the sizing process determines the Cisco UCS chassis classes. Individual chassis can host up to eight half-sized blades or up to four full-sized blades.

The Cisco UCS Chassis sizing determines the following:


- Maximum number of blades a certain Cisco UCS chassis class can host
- Cisco UCS server blade classes and Cisco UCS chassis class hosts
- Required number of uplinks from IOM to fabric interconnect
- Redundancy requirements—for example, whether a single fabric interconnect failure in a Cisco UCS cluster can result in bandwidth being halved or not

Cisco UCS Chassis Options

Cisco UCS 5108 Chassis



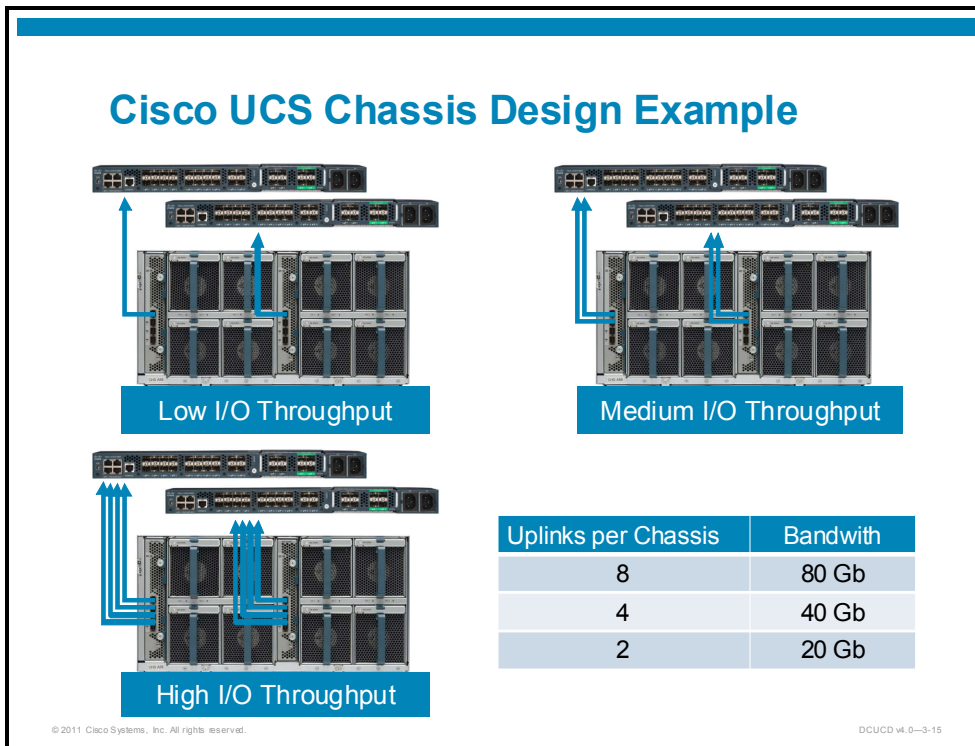
Cisco UCS 2104 XP IO Module



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-14

The fabric redundancy from the chassis perspective is addressed by using two Cisco UCS 2104XP IOM modules.



An example result of a Cisco UCS chassis sizing process, seen in the figure, gives three chassis types with regard to the number of uplinks used. The general requirement for any Cisco UCS chassis class with this sizing example is that the individual chassis must comply to dual-fabric design—that is, connect to two fabric interconnects in a Cisco UCS cluster. In the example, the Cisco UCS 6120XP Fabric Interconnects that are in a cluster setup are used.

High I/O Throughput Chassis Class

Connections from the Cisco UCS 6120XP Switch to Cisco UCS 2104XP IOMs use all four ports per fabric, supporting up to five chassis per switch cluster. This chassis class can be used for server applications requiring high I/O throughput, for example, virtualization solutions, databases, and other I/O-intensive clients. In this setup, the total bandwidth available per chassis is 80 Gb. The actual amount used is governed by the applications that are deployed on this setup.

Medium I/O Throughput Chassis Class

Connections from the Cisco UCS 6120XP Switch to Cisco UCS 2104XP IOMs use 2 ports per fabric, supporting up to 10 chassis per switch cluster. This chassis class can be used for server applications requiring medium I/O throughput, for example, some bare-metal operating system application deployments. In this setup, the total bandwidth available per chassis is 40 Gb.

Low I/O Throughput Chassis Class

Connections from the Cisco UCS 6120XP Switch to Cisco UCS 2104XP IOMs use 1 port per fabric, supporting up to 20 chassis per switch cluster. This chassis class can be used for server applications requiring low I/O throughput, for example, non-I/O-demanding bare-metal operating system application deployments.

Sizing Cisco UCS Fabric Interconnect

Identify fabric interconnect requirements

Select type—20- or 40-port

Set 10GE SFP+ type—twinax vs. FO (distance)

Define 10 GE port quantities (licensing)

- Uplinks to LAN core
- Server downlinks to chassis—number of 10-GE ports to connect chassis



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-16

The last step of the Cisco UCS sizing process must identify fabric interconnect requirements—that is, the number of 10 Gigabit Ethernet server downlinks, 10 Gigabit Ethernet LAN uplinks, and Fibre Channel SAN uplinks.

The number of 10 Gigabit Ethernet server downlinks is determined by the chassis class and quantity of connections to the fabric interconnect.

The number of 10 Gigabit Ethernet LAN uplinks is determined by the summary of LAN throughput requirements of all the blade servers that connect to the fabric interconnect cluster.

The number and speed of Fibre Channel SAN uplinks is determined by the summary of SAN throughput requirements of all the blade servers that connect to the fabric interconnect cluster.

Apart from LAN and SAN throughput requirements of all the blade servers, the redundancy also influences the number of 10 Gigabit Ethernet and Fibre Channel uplinks. A fabric interconnect failure should not result in bandwidth being halved; the summary values for 10 Gigabit Ethernet and Fibre Channel should be doubled.

Sizing Cisco UCS Cluster—Fabric Interconnects

SAN core connectivity—Fibre Channel uplinks

- Define 1/2/4/8 Gb Fibre Channel port quantities
- Dedicated Fibre Channel uplink per VSAN

Expansion module options

- 6x 10 GE ports
- 4x 10 GE and 4x 1/2/4 Gb Fibre Channel ports
- 8x 1/2/4 Gb Fibre Channel ports
- 6x 1/2/4/8 Gb Fibre Channel ports



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-17

When the Cisco UCS B-Series cluster needs to connect to the Fibre Channel SAN, the proper expansion module has to be selected.

There are three options for attaching to the Fibre Channel SAN—the main difference being the number and speed of Fibre Channel ports.

For high throughput Fibre Channel attachment, the 6-port 8 Gb Fibre Channel expansion module is the most appropriate, if the core SAN also supports the 8 Gb Fibre Channel rates.

The second factor that governs the Fibre Channel expansion module selection is the number of VSANs to which Cisco UCS has to be attached. Note that a single Fibre Channel port on the expansion module can carry traffic for a single VSAN or it can trunk multiple VSANs (version 1.4 software required).

In the example, the decision was made to have a dedicated VSAN per uplink Fibre Channel port. Thus, the requirement to attach to eight different VSANs from individual fabric interconnects (that is, to the physical SAN fabric) leaves no choice for the Cisco UCS 6120XP Switch. The 8-port 4 Gb Fibre Channel expansion module should be used in this case.

Example—Design Requirements

Server Type	Qty	Description
Type 1	18	Physical server installation
Type 2	17	Physical server installation
Type 3	15	VMware vSphere environment

Type 1 Requirements	Type 2 Requirements	Type 3 Requirements
Single medium-fast CPU	Single CPU	Dual 4-core fast CPU
8 GB memory	4 GB memory	48 GB memory
2x 1 GE trunk with VLANs	2x 1 GE	6x 1 GE trunk with VLANs
2x HBA	2x HBA	2x HBA
Local disk	SAN boot	Local disk
LAN throughput = 1 Gb/s	LAN throughput = 0.5 Gb/s	LAN throughput = 1.5 Gb/s
SAN throughput = 0.4 Gb/s	SAN throughput = 0.4 Gb/s	SAN throughput = 1 Gb/s

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-18

The design workshop has accessed the following requirements for the servers:

- 18 physical servers with a bare-metal operating system must be deployed. Requirements for such servers are as follows:
 - A single fast CPU should be present.
 - Server must have at least 48 GB of memory.
 - Two 1 Gigabit Ethernet network interface cards (NICs) with 802.1Q trunking support should be available for redundant connectivity.
 - Two Fibre Channel HBAs for redundant SAN connectivity should be available.
 - Local disk is used for operating system installation.
 - Applications will save data on a SAN-attached storage.
- 17 physical servers with a bare-metal operating system must be deployed. Requirements for such servers are as follows:
 - A single medium-fast CPU should be present.
 - Server must have at least 12 GB of memory.
 - Two 1 Gigabit Ethernet NICs for redundant LAN connectivity should be available.
 - Two Fibre Channel HBAs for redundant SAN connectivity should be available.
 - Operating system will be booted from SAN, and applications will store data on SAN attached storage.

- 15 physical servers with VMware ESX host must be deployed. Requirements for such servers are as follows:
 - Two fast multicore CPUs should be present.
 - Server must have at least 96 GB of memory.
 - Four 1 Gigabit Ethernet NICs for redundant LAN connectivity should be available.
 - Two Fibre Channel HBAs for redundant SAN connectivity should be available.
 - VMware vSphere ESX will be booted from local disk.
 - Virtual machines (VMs) will be stored on SAN attached storage.

Example—Sizing Server Blades

Redundancy requirement

- Fabric interconnect failure should not result in throughput halved

	Type 1 Blade	Type 2 Blade	Type 3 Blade
Blade Type	B200-M1	B200-M1	B200-M1
Processor	1x Intel Xeon E5540	1x Intel Xeon L5520	2x Intel Xeon X5570
Memory	4x 2-GB DIMM	2x 2GB DIMM	12x 4-GB DIMM
Mezzanine	M71KR-Q CNA	M71KR-Q CNA	M81KR VIC
Disk	2x 73 GB 15K rpm	None	2x73GB 15K rpm

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-19

The first step of the sizing process determines the blade server classes. From the input information, the following three blade server classes have been identified:

- PS-class 1 that is sized in the following manner:
 - **Blade type:** Cisco UCS B200-M1
 - **Processor:** Intel Xeon 5540
 - **Memory:** 48 GB with 12x 4-GB DIMMs in bank 0
 - **Adapter:** M71KR-Q CNA since the server must be connected to LAN and SAN
- PS-class 2 that is sized in the following manner:
 - **Blade type:** Cisco UCS B200-M1
 - **Processor:** Intel Xeon 5520
 - **Memory:** 12 GB with 6x 2-GB DIMMs in bank 0
 - **Adapter:** M71KR-Q CNA since the server must be connected to LAN and SAN
- VS-class 1 that is sized in the following manner:
 - **Blade type:** Cisco UCS B200-M1
 - **Processor:** Two Intel Xeon 55470
 - **Memory:** 96 GB with 24x 4-GB DIMMs spread between the two memory banks
 - **Adapter:** M81KR VIC since the server must be connected to LAN with multiple NICs and also to SAN

Apart from that, the input information also states that the failure of individual fabric interconnects must not result in throughput degradation. You have also gathered required LAN and SAN throughput for each individual server class from the customer.

Example—Analyze System Requirements

	Type 1	Type 2	Type 3	Summary
Blade quantity	18 half-sized	17 half-sized	15 half-sized	50 half-sized
LAN throughput	18 Gb/s	8.5 Gb/s	22.5 Gb/s	49 Gb/s
SAN throughput	7.2 Gb/s	6.8 Gb/s	15 Gb/s	29 Gb/s
Server throughput	25.2 Gb/s	15.3 Gb/s	37.5 Gb/s	78 Gb/s

	Type 1	Type 2	Type 3	Summary
10 GE LAN uplinks	2	1	3	6
4G Fibre Channel uplinks	2	2	4	8
10 GE server downlinks	3	2	4	9

Minimum of seven chassis

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-20

The system requirements analysis is the first step in evaluating the number of chassis required. The total number of server blades that is required is 50 and all of them are half-sized.

Dividing 50 blades by 8 blades per chassis (which is the maximum number of blades per chassis) shows that 7 chassis should be enough to hold all the blades.

Next, the math for the server side, LAN uplinks, and SAN uplinks is done.

- The total throughput from 18 PS-class 1 servers is 27 Gb/s, which is 18×1.5 Gb/s. 1.5 Gb/s is the summary of the required Ethernet and Fibre Channel throughput ($1 + 0.5$),
- The total throughput from 17 PS-class 2 servers is 17 Gb/s, which is 17×1 Gb/s. 1 Gb/s is the summary of the required Ethernet and Fibre Channel throughput ($0.5 + 0.5$),
- The total throughput from 15 VS-class 1 servers is 90 Gb/s, which is 15×6 Gb/s. 6 Gb/s is the summary of the required Ethernet and Fibre Channel throughput ($3 + 3$),

From these numbers, the server uplinks can be determined:

- To connect 18 PS-class 1 servers without oversubscription, three 10 Gigabit Ethernet IOM fabric ports are required.
- To connect 17 PS-class 2 servers without oversubscription, two 10 Gigabit Ethernet IOM fabric ports are required.
- To connect 15 VS-class 1 servers without oversubscription, nine 10 Gigabit Ethernet IOM fabric ports are required.

The number of LAN uplink ports is determined by calculating the required LAN throughput. Note that oversubscription is avoided.

- $18 \text{ servers} \times 1 \text{ Gb/s} = 18 \text{ Gb/s} \Rightarrow$ two 10 Gigabit Ethernet LAN uplink ports
- $17 \text{ servers} \times 0.5 \text{ Gb/s} = 8.5 \text{ Gb/s} \Rightarrow$ one 10 Gigabit Ethernet LAN uplink port
- $15 \text{ servers} \times 3 \text{ Gb/s} = 45 \text{ Gb/s} \Rightarrow$ five 10 Gigabit Ethernet LAN uplink ports

Similar to the LAN uplink port requirement, the SAN uplink port requirement can be calculated.

Example—Sizing Blade Chassis

Chassis Type 1	Description
Server blade	4x type 1 and 4x type 2
UCS 2104XP IOM	2x
LAN throughput	6 Gb/s (4 + 2)
SAN throughput	3.2 Gb/s (2 x 1.6)
10GE server downlinks	1x

Chassis Type 2 a/b	Description
Server blade	5x type 3, 1x type 1 / type 2
UCS 2104XP IOM	2x
Max LAN throughput	8.5 Gb/s (7.5 + 1)
Max SAN throughput	5.4 Gb/s (5 + 0.4)
10 GE server downlinks	2



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-21

Next, the blade chassis classes are identified. It was decided to go with two classes to keep the Cisco UCS setup simple, which eases the management effort when the system needs to be scaled.

BC-class 1 chassis class will be used to host the PS-class 1 and PS-class 2 server blades. There will be up to eight servers per chassis. A single chassis will be equipped with two Cisco UCS 2104XP IOM, and a single port on an individual IOM will be connected to the upstream fabric interconnect.

BC-class 2 chassis class will be used to host up to five VS-class 1 server blades. A single chassis that is equipped with two Cisco UCS 2104XP IOMs will be connected to the upstream fabric interconnect with four server uplinks.

Example—Populating Blade Chassis

Chassis No.	Chassis Type	Type 1	Type 2	Type 3
1	Chassis 1	4	4	-
2	Chassis 1	4	4	-
3	Chassis 1	4	4	-
4	Chassis 1	4	4	-
5	Chassis 2	1	-	5
6	Chassis 2	1	-	5
7	Chassis 2	-	1	5
Summary		18	17	15

	Fabric A	Fabric B
10 GE server downlinks	8	8
10 GE LAN uplinks	5	5
4G Fibre Channel uplinks	8	8

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-22

Next, you need to determine the number of required blade chassis. You know that the number of servers you must deploy is 18x PS-class 1, 17x PS-class 2, and 15x VS-class 1.

The BC-class 2 chassis is supposed to hold up to five VS-class 1 blades, which results in five such chassis. The limit of five server blades is determined by the available server uplink throughput, which is four 10 Gigabit Ethernet for LAN and SAN traffic (remember that fabric interconnect failure should not result in bandwidth halved).

The BC-class 1 chassis can hold up to eight half-size blades. If a single chassis is populated with eight PS-class 2 blades, the server uplink throughput for LAN and SAN is sufficient. Thus, two such chassis can be deployed.

With a single PS-class 2 blade, you can put an additional six PS-class 1 blades in a BC-class 1 chassis to maintain the sufficient server uplink throughput for LAN and SAN connectivity.

Finally, you are left with 12 PS-class 1 blades, which can be evenly spread between 2 additional BC-class 1 chassis.

Example—Sizing Fabric Interconnect

Cisco UCS 6120XP Fabric interconnect per fabric

- N10-E0080 expansion module with eight Fibre Channel ports
- Five on-demand port activation licenses

	Fabric A	Fabric B
Server downlinks	8x twinax	8x twinax
LAN uplinks	5x USR SPF+	5x USR SFP+
Fibre Channel uplinks	8x 4G Fibre Channel MM SFP	8x 4G Fibre Channel MM SFP

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-23

Here is the result of the chassis design:

- 5x BC-class 1 chassis, which require 5x 10 Gigabit Ethernet server side uplinks
- 3x BC-class 2 chassis, which require 12x 10 Gigabit Ethernet server side uplinks
- 8x 10 Gigabit Ethernet LAN uplinks are required for the LAN connectivity
- 16x 4G Fibre Channel SAN uplinks are required for the SAN connectivity.

The required number Ethernet and Fibre Channel ports calculated is per individual fabric interconnect.

For the fabric interconnect sizing, you operate with the following input information per individual fabric interconnect:

- 25x 10 Gigabit Ethernet ports are required for server side and LAN uplink connectivity
- 16x 4G Fibre Channel ports are required for SAN uplink connectivity

Given these numbers, the Cisco UCS 6140XP Fabric Interconnect in a cluster setup is the choice for the example Cisco UCS systems. The individual fabric interconnect will be equipped with the following:

- 2x N10-E0080 expansion modules, where each module has eight 1/2/4G Fibre Channel ports
- 16x 1/2/4G Fibre Channel MM SFP for SAN connectivity
- 8x 10 Gigabit Ethernet Ultra-Short Reach SFP+ will be used to connect to upstream LAN switches
- 10x 10 Gigabit Ethernet copper SPF+ will be used to connect the blade chassis.

Note again that this is the number of ports per individual fabric interconnect because of the requirement that the failure of one should not result in bandwidth halved.

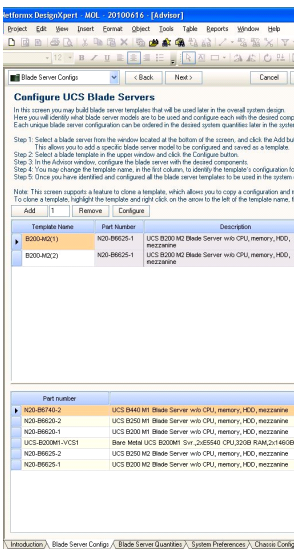
Creating the B-Series BOM

This topic identifies and describes how to create a B-Series solution BOM.

Create BOM

- NetformX DesignXpert
 - Available to UCS ATP partners
 - Acquire credentials
 - Not a sizing tool
- BOM creation process
 - Configuration
 - Service selection
 - Upload

www.ciscoprc.com



Template Name	Part Number	Description
B200A0C1	N20-88625-1	UCS B200 M2 Blade Server w/o CPU, memory, HDD, mezzanine
B200A0C2	N20-88625-1	UCS B200 M2 Blade Server w/o CPU, memory, HDD, mezzanine

Part number	Part number	Description
N20-88745-2	UCS B440 M1 Blade Server w/o CPU, memory, HDD, mezzanine	
N20-88625-2	UCS B200 M1 Blade Server w/o CPU, memory, HDD, mezzanine	
N20-88625-1	UCS B200 M1 Blade Server w/o CPU, memory, HDD, mezzanine	
UCS-B200M1-VCS1	Base Mini UCS B200M1 Srv. 2x5540 CRU, 2x08 RAM, 2x1460B	
N20-88625-2	UCS B200 M2 Blade Server w/o CPU, memory, HDD, mezzanine	
N20-88625-1	UCS B200 M2 Blade Server w/o CPU, memory, HDD, mezzanine	

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-25

The BOM for the Cisco UCS is created using the NetformX DesignXpert tool. The tool is acquired at <http://design.netformx.com/cisco-navigate-to-accelerate/>. In order to be able to use the tool to compile the BOM, proper credentials must be acquired. The credentials are available to Cisco UCS ATP partners.

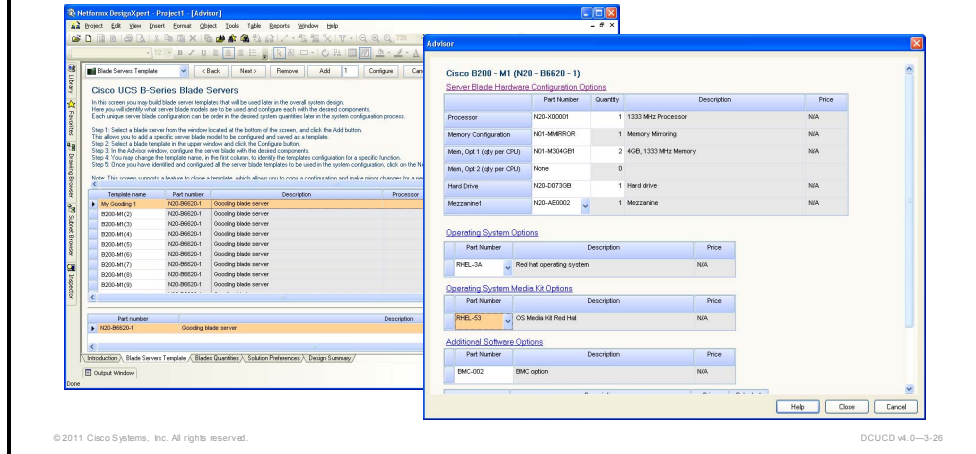
Once the tool is installed with proper credentials, the BOM can be created. This process has three steps:

- Configuration
- Service selection
- Upload

More information about the NetformX DesignXpert tool is available at Cisco Partner Resources Center at <http://www.ciscoprc.com>.

Configuration Process—Blade Library

- Pick blades
- Configure blade flavors



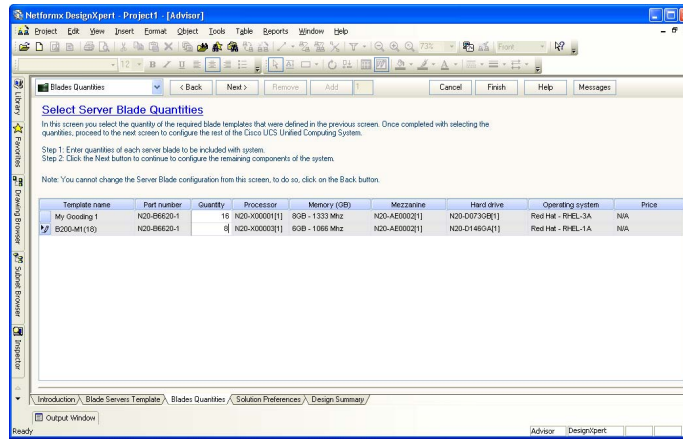
The configuration process uses a bottom-up approach—first the blades are configured, then the chassis is configured, and finally the fabric interconnects. The process consists of the following steps:

- Blade library
- Blade quantity
- Solution preferences
- Design summary

In the first step, the blade library, you pick and configure blade flavors (classes). Based on the requirements, blade types have to be added to the library along with the quantity. An individual blade type can be configured per the requirements, which includes processors, memory, adapters, and so on. After the blade library configuration is finished and the library is populated with the desired blade types, the blade quantities can be entered.

Configuration Process—Blade Quantity

- Pick blade flavor
- Specify number of blade configurations



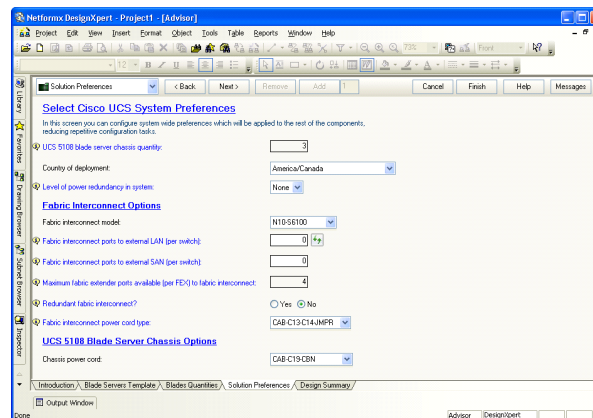
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-27

When the blade classes are defined, the quantities per class have to be entered.

Configuration Process—Solution Preferences

- Allowed solution preferences
- Switch type and cable preferences



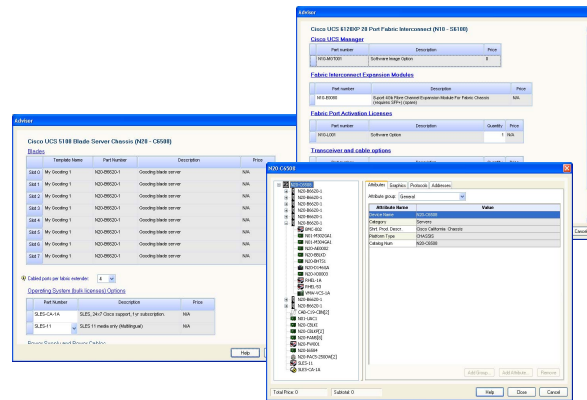
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-28

In the third step, the solution preferences are defined. These preferences include the selection of the fabric interconnect model, ports, cabling, and chassis quantity.

Configuration Process—Design Summary

- Blade distribution into chassis
- Chassis configuration
- Switch configuration



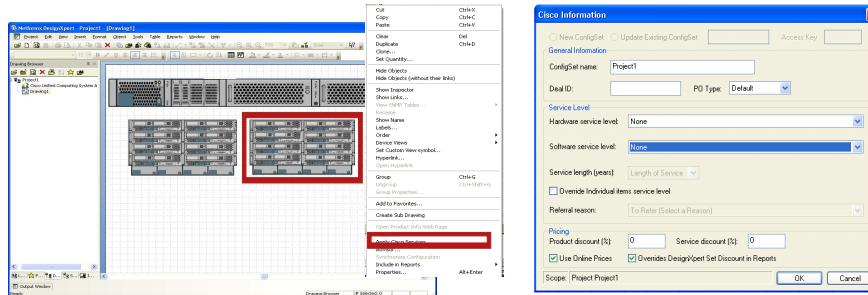
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-29

In the final step, you distribute blades into the chassis, apply the chassis and fabric interconnect configuration, and ensure that the BOM is complete.

Service Selection Process

- Apply services individually or to the entire system
- Review BOM by different categories



© 2011 Cisco Systems, Inc. All rights reserved.

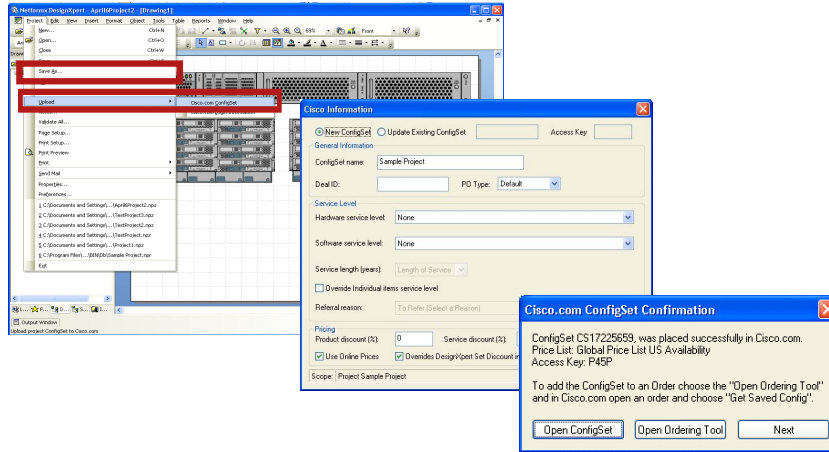
DCUCD v4.0-3-30

The service selection process is used to define the type of services that are to be applied to the selected Cisco UCS components. The services can be selected for the whole BOM or part of the BOM. The selection includes definition of the hardware and software service level, length of the service, and so on.

Before proceeding to the final configuration step—uploading—the BOM should be reviewed to verify that it is correct.

Upload Process

- Upload the ConfigSet to Cisco.com



Once the hardware and services have been selected, the BOM ConfigSet can be uploaded to Cisco.com.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The Cisco UCS B-Series sizing process encompasses selection of server blades and components, chassis and connectivity, and fabric interconnects and upstream connectivity.
- The chassis sizing process should carefully implement required throughput and allowed oversubscription.
- The sizing process defines the BOM—the Cisco UCS hardware components, which includes server blades, chassis, and fabric interconnects.

Planning Physical Deployment

Overview

This lesson discusses how to create a physical deployment plan for the Cisco UCS solution.

Objectives

Upon completing this lesson, you will be able to deploy the Cisco UCS solution. This includes the ability to meet these objectives:

- Identify and describe the Cisco Power Calculator tool
- Discuss the physical deployment plan

Cisco UCS Power Calculator Tool

This topic identifies and describes the Cisco Power Calculator tool.

Cisco UCS Power and Cooling Requirements

Cisco UCS Power Calculator tool

- <http://www.ciscoprc.com/resource/lib.asp?id=937>
- http://www.cisco.com/assets/cdc_content_elements/flash/dataCenter/cisco_ucs_power_calculator/

Enter Cisco UCS configuration

- B-Series and/or C-Series blades, chassis, fabric interconnect

Power and cooling per

- Idle
- 50% load
- Maximum load

Weight

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-4

When designing physical deployment for Cisco UCS, the power consumption has to be calculated for the solution. This calculation is necessary because using one or two processors, disk drives, or DIMMs of different speed and size will result in different power consumption.

The Cisco UCS Power Calculator tool enables the designer to calculate exact values for the environmental parameters of the designed Cisco UCS solution—individual chassis, fabric interconnect switches, and B-Series and C-Series servers:

- Power consumption required and cooling capacity for idle, 50 percent, and maximum load
- Weight

The tool is available to Cisco partners:

- <http://www.ciscoprc.com/resource/lib.asp?id=937>
- http://www.cisco.com/assets/cdc_content_elements/flash/dataCenter/cisco_ucs_power_calculator/

Example—UCS Power Calculation

Fabric Interconnect model: Cisco UCS 6120XP 20-Port
 Redundancy: Yes
 Power supply: Redundant
 Expansion module 1: Fibre-channel

Chassis configuration 1

Idle power per chassis (watts): 1072
 50% load power per chassis (watts): 2062
 Max power per chassis (watts): 3078
 Chassis weight (lbs): 146

B200

Blade configuration 1

Cisco UCS B200 M1

Processors: 2xIntel Xeon E5540 (2.53GHz)
 Memory: 12x8GB
 Disk drives: 2x146GB
 Mezzanine: Cisco UCS M8 1KR VIC cards

Configure 8 blades this way
 Delete Add New Customize

Chassis PSU redundancy: N+1
 Cable connections per fabric extender: 4

Number of racks: 1
 Available ports per Fabric Interconnect: 16

	Fabric Interconnect	Chassis	Rack mount servers	Total	
Number configured	2	1	0		
Idle power	watts	494	1072	0	1566
	btu	1685	3656	0	5341
50% load power	watts	640	2062	0	2702
	btu	2182	7031	0	9213
Max power	watts	786	3078	0	3864
	btu	2680	10496	0	13176
Weight	lbs	68	146	0	214
	US/Metric				
Tons cooling				0.8	
Total number blades configured				8	

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-5

The Power Calculator tool is used to calculate the characteristics for this equipment:

- Cisco UCS C-Series rack-mount servers
- Cisco UCS B-Series blade servers and chassis
- Cisco UCS 6100XP Fabric Interconnect switches

To get the report on required power, cooling capacity, and weight, the design has to select the hardware, which comprises the solution with the proper characteristics (that is, redundancy, quantities, and so on).

Creating the Physical Deployment Plan

This topic discusses the physical deployment plan.

Physical Deployment Considerations

- Rack capacity—chassis density per rack cabinet
 - Usable space
 - Loading
 - Power and cooling—affected by chassis requirement
- Rack density per array (row)
- Fabric interconnect physical placement
 - Depends on racks and chassis density
 - Governs selection of cabling

Cabling Type	Distance (m)	Placement
Copper (Twinax)	1, 3, and 5	MoR
Optical (Short Reach)	82	EoR
Optical (Short Reach)	300	Multirow

© 2011 Cisco Systems, Inc. All rights reserved.

An important part of any data center solution is physical deployment design. This design includes the following:

- Sizing the racks:
 - Determining how much space is available per rack
 - Defining the amount of equipment per rack
 - Calculating the power requirements per rack
 - Determining the number of power supplies per Cisco UCS 5108 chassis to respect the required redundancy level
 - Calculating the heat dissipation per rack
- Sizing the array, which includes determining how many rack cabinets are required per Cisco UCS cluster.
- Fabric interconnect placement, which includes determining where to put the Cisco UCS 6100XP Fabric Interconnect switches, and which cables to use for the I/O Module (IOM) to fabric interconnect physical connectivity.

The table in the figure summarizes the cable lengths, which depend on the media type.

When calculating the heat dissipation, you can use the standard energy-to-BTU rating (1 W = 3.41 BTU/hr).

Rack Capacity Design Example

Space per Rack

Rack type	Usable RU	Chassis per Rack
6-foot rack	42	Up to 7
7-foot rack	44	Up to 7

Power Requirements per Rack

Chassis per Rack	4-Blade Conf. (W)	8-Blade Conf. (W)
2	2474	4114
3	3711	6171
4	4948	8228
5	6184	10,285
6	7422	12,342
7	8659	14,399

Power Requirements per Chassis*

	Power Consumption Under Load (W)
Chassis	417
Per Blade	205
4-Blade Configuration	1237
8-Blade Configuration	2057

N+1 Power Supply Redundancy

Blade Quantity	PS Quantity
4	42
8	44

* Using two Intel Xeon E5540, 16 GB RAM, two 72-GB drives

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-3-8

To calculate the available space per rack, the designer needs information about the size of the rack cabinets that are going to be used. Typically, the rack unit (RU) is used for this purpose. Common rack sizes used are as follows:

- Six feet, where 42 RU of space is available
- Seven feet, where 44 RU of space is available

With these two options, up to seven Cisco UCS 5108 chassis (6 RU) can be installed per rack.


Second, the power requirement per rack is also important in this design, since the power is limited. To calculate how much power is needed per rack, the designer has to determine how much power is used by the equipment. This is achieved with the Cisco UCS Power Calculator tool, which gives information about the power consumption in watts. Note that these values differ under different conditions, such as when the equipment is idle and when it is under 100 percent load.

Based on per-chassis power requirements, the designer can calculate per-rack power requirements for a different number of chassis in the rack as well as the heat dissipation (using the standard energy-to-BTU rating).

Note The power requirements per chassis table lists the values measured for the blade using two Intel Xeon E5540 processors, 16 GB RAM, and two 72-GB disks. The actual numbers might differ in your case.

Array Capacity Design

- High-density rack
 - More power per rack—fewer racks per array
- Low-density rack
 - Less power per rack—more racks per array
- Cabling consideration
 - Chassis density per rack
 - Number of uplinks per chassis
- Rack density per Cisco UCS 6100XP cluster



3 Chassis per Rack

IOM Uplinks	Uplinks per Rack	Racks per UCS 6120XP Cluster	Racks per UCS 6140XP Cluster
One 10 GE	6	Up to 6	Up to 13
Two 10 GE	12	Up to 3	Up to 6
Four 10 GE	24	Up to 1	Up to 3

GE = Gigabit Ethernet

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-9

Next, the rack density per Cisco UCS cluster can be determined. The actual number varies case by case, but in general, it depends on the number of chassis that are supported by the Cisco UCS cluster.

The table in the figure lists an example where up to three chassis per rack are installed. The calculation is done for the maximum number of racks per Cisco UCS 6120XP or 6140XP Fabric Interconnect switches. Note that the number of racks that are calculated in the example also depends on the number of IOM uplinks.

Array Capacity Design Example

4 Chassis per Rack

IOM Uplinks	Uplinks per Rack	Racks per UCS 6120XP	Racks per UCS 6140XP
One 10 GE	8	Up to 5	Up to 10
Two 10 GE	16	Up to 2	Up to 5
Four 10 GE	32	Up to 1	Up to 2

5 Chassis per Rack

IOM Uplinks	Uplinks per Rack	Racks per UCS 6120XP	Racks per UCS 6140XP
One 10 GE	8	Up to 5	Up to 10
Two 10 GE	16	Up to 2	Up to 5
Four 10 GE	32	Up to 1	Up to 2

6 Chassis per Rack

IOM Uplinks	Uplinks per Rack	Racks per UCS 6120XP	Racks per UCS 6140XP
One 10 GE	8	Up to 5	Up to 10
Two 10 GE	16	Up to 2	Up to 5
Four 10 GE	32	Up to 1	Up to 2

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-10

When designing physical deployment, a different number of chassis can be put into a single rack. The three tables seen in the figure list the calculation for four, five, and six chassis per rack design where one, two, or four IOM to fabric interconnect uplinks are used.

Physical Deployment Example

Cisco UCS 6120XP Cluster

- Total power = 19,213 W
- Total heat = 66,516.33 BTU/hr


	Value
Blades* per chassis	8
Chassis per rack	3
Uplinks per IOM	2
Racks per cluster	3
Ports per fabric	18

Cabling Type	Distance (m)	Placement
Copper (Twinax)	1, 3, and 5	MoR

	Value
Power per rack	6171 W
Heat dissipation	21,043.11 BTU/hr
Power per UCS 6120XP	350 W
Heat dissipation	1193.5 BTU/hr

* Using Intel Xeon E5540, 16 GB RAM, two 72-GB drives

© 2011 Cisco Systems, Inc. All rights reserved.



DCUCD v4.0-3-11

This example describes the physical deployment design using the following input information:

- Chassis with eight blades are used.
- Consumption of a chassis was determined by preliminary measurement and is 205 W per blade server and 417 W per chassis.
- You need to install nine chassis and can use three 42-RU rack cabinets.
- From an individual chassis, two IOM to fabric interconnect uplinks per IOM are used.
- The available power per rack is 10 kW.
- The measured power consumption per Cisco UCS 6120XP Fabric Interconnect switch is 350 W.

Based on the input information, the following design has been proposed:

- Install up to three chassis per rack.
- A single rack requires 6171 W of power (counting only chassis) and produces 21,002.19 BTU/hr of heat.
- Two fabric interconnect switches will be placed in the middle rack cabinet.
- That rack will need an extra 700 W of power and produce an extra 1193.5 BTU/hr of heat.

Note Installing three chassis plus two fabric interconnect switches in a rack respects the 10 kW power limit per rack.

- The 5-m Twinax cables will be used between the IOM and fabric interconnects.

In total, the whole cluster at maximum load would need 19,213 W of power and produce 66,516.33 BTU/hr of heat.

Note The power requirements per chassis might differ in your case.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The Cisco UCS Power Calculator tool is used to calculate the power consumption, required cooling capacity, and weight of the designed UCS solution.
- Physical deployment design is necessary to determine rack and array densities with regard to available power and space.

Examining the Cisco UCS Network and Storage

Overview

This lesson discusses the Cisco Unified Computing System (UCS) LAN, SAN, application service, and storage products and technologies that are relevant for a given set of requirements.

Objectives

Upon completing this lesson, you will be able to identify and describe Cisco UCS LAN and SAN products and technologies. This includes the ability to meet these objectives:

- Identify and describe Cisco UCS LAN products and concepts
- Identify and describe the Cisco Nexus 1000V
- Identify and describe Cisco Nexus 1000V integration with VMware vCenter
- Identify and describe Cisco UCS SAN products and concepts
- Identify and describe Cisco UCS storage products and concepts

Cisco UCS LAN

This topic identifies and describes Cisco UCS LAN products and concepts.

LAN Aspects

- Upstream LAN switch(es)
- VLANs
 - No VTP on UCS
 - VLAN overlapping verification
- Interfaces
 - PortChannel configuration
 - Trunk configuration—allow proper VLANs toward UCS

```
interface Vlan11
no shutdown
ip address 10.1.10.252/24
hsrp 1
preempt
priority 150
ip 10.1.10.254

interface Vlan12
no shutdown
ip address 10.1.20.252/24
hsrp 1
preempt
priority 150
ip 10.1.20.254

interface Vlan13
no shutdown
ip address 10.1.30.252/24
hsrp 1
preempt
priority 150
ip 10.1.30.254

interface Vlan21
no shutdown
ip address 10.2.10.252/24
hsrp 1
preempt
priority 150
ip 10.2.10.254
```

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-4

The Cisco UCS clusters and UCS C-series servers are connected via the network. The clients communicate with the server over the network, which must observe certain aspects.

The aspects of the Cisco Data Center UCS solution should be addressed at all layers of the network. The access layer is the first of these, since this is where the servers are attached. Next is the aggregation layer, where traffic manipulation services are typically deployed. Finally, yet importantly, there is the core layer, which must provide maximum throughput and connectivity.

The Cisco UCS deployment LAN aspects govern upstream LAN switch configuration and connectivity characteristics.

Interfaces are the first aspect that needs to be considered – the interfaces where Cisco UCS is connected can be put in the PortChannel configuration to scale the bandwidth and improve convergence. Likewise, the interfaces have to be properly configured as trunk ports in order to allow the VLANs, with which the Cisco UCS is deployed. Since Cisco UCS does not support Virtual Terminal Protocol (VTP), any new VLANs must be manually added on the upstream switches. It is necessary to verify that the purpose and security policy of Cisco UCS VLANs and upstream switch VLANs are the same.

Connectivity is the first aspect to consider. Connectivity for the Cisco Data Center Unified Computing System solution encompasses the Open Systems Interconnection (OSI) layer and Layer 1 physical connectivity to the network connectivity at Layer 3.

LAN Prerequisites

Connectivity

- 10 Gigabit Ethernet for Cisco UCS uplink connectivity
- Copper vs. fiber optics
- Interfaces should be available on different linecards to improve high availability
- Oversubscription level

Redundancy levels

- Dual-fabric design
- Dual hardware elements



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-3-5

These are the key aspects that must be properly addressed by the network design and deployment:

- **Connectivity:** If there is no connectivity or poor connectivity, service will suffer and consequently productivity will be low.
- **High availability:** The Cisco Data Center network must be designed and deployed in a way to easily overcome failure. The goal is zero downtime.

Layer 1 Aspects


Layer 1 connectivity encompasses links between devices. No matter whether the interconnected devices are switches or servers to switches, they need to be properly connected.

The following issues must be considered:

- The type of media that will be used—i.e., copper or fiber optics. This decision depends on the physical cabling availability, server and switch side interface options, and cost.
- If copper-based connections will be used, the speed of the cabling support—100 Mb/s, 1 Gb/s, or 10 Gb/s—must be considered. This determines the copper cable category.
- If the fiber-optic media is chosen, then the fiber type must be determined. This depends on the cabling availability, server and switch side interface options, and cost and connection length. Single-mode fibers can be used over higher distances.
- Depending on the server and switch side interfaces, the connector type of the cable should also be selected. With copper-based connections, this is typically RJ45. For 10 Gigabit Ethernet deployments, Twinax might be used, and for fiber optics, currently the LC connector type is most commonly used.

Layer 2 Network Aspects

- Layer 2 topology:
 - Fault domain size and convergence
 - Server traffic engineering with path selection
 - Network virtualization with VLANs
- Layer 2 adjacency requirements:
 - Clustering
 - VM mobility
- NIC teaming for multipathing
- Capacity planning:
 - Link utilization and load sharing
 - Oversubscription points
 - Uplink/downlink ratio



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-6

Layer 2 is very important for the Cisco Data Center Unified Computing System solution, especially due to Layer 2 adjacency requirements.

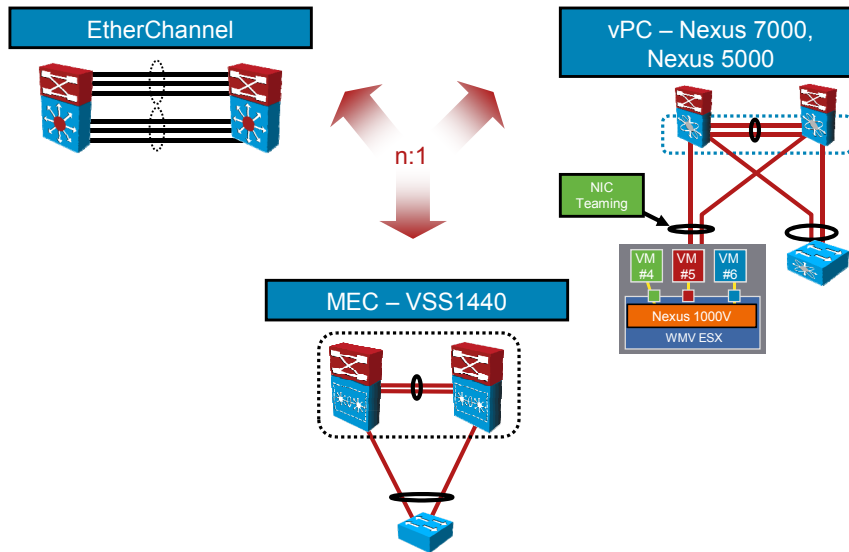
The Cisco Data Center Unified Computing System solution encompasses server virtualization also—for example, with VMware. The benefit of server virtualization is the hardware abstraction, which can be exploited for virtual machine (VM) mobility. VMs can be migrated—live or suspended—to a different ESX server, either to load balance the computing load or to be able to complete maintenance tasks. The VM mobility requires Layer 2 adjacency, thus in the Cisco Data Center Unified Computing System solution access layer, Layer 3 is not an option when VM mobility is required (and typically you would use that) for an ESX cluster.

These are aspects that need to be considered for Layer 2 connectivity:

- Size of the fault or broadcast domain
- Convergence upon failure—does the spanning tree cope with the requirement, or do you need to deploy a faster mechanism?
- Use of traffic engineering to control how traffic is spread over uplinks
- Packing and sending traffic for multiple segments over single wire—a long-known virtualization using VLANs and trunks (although typically this is not seen as link virtualization)
- Servers need redundancy, thus they need to be connected to two or more access layer switches. You must consider how this connectivity will be handled. Server-side multipathing is required, as well as the possible use of port channeling, which also scales the bandwidth.
- Do you want loops? Typically yes, since this is how you achieve redundancy in topology. However, to provide fast convergence, you may not want to use loops with Cisco Data Center.

Layer 2 connectivity must be planned with required capacities and uplink oversubscription in mind. The oversubscription does not affect only server-to-client throughput (vertical) but also server-to-server traffic exchange (horizontal), thus proper oversubscription design must be in place.

Scaling Network Connectivity



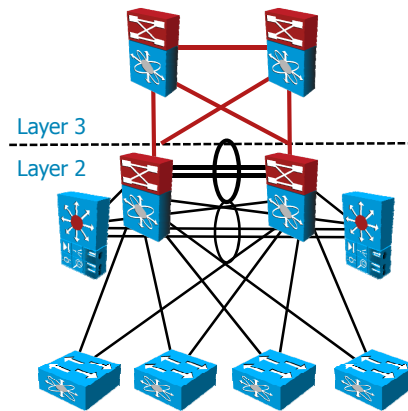
In data centers, you can have network devices such as Cisco Catalyst 6500 series switches employing Virtual Switching System (VSS) technology, thus converting two physical switches to appear as one switch to the rest of the network.

Another example is the bundling of multiple physical links to appear as one logical link. This can be a simple set of parallel links between a pair of devices (that is, EtherChannel or network interface card [NIC] teaming), or an even more robust set of links to multiple devices (Multichassis EtherChannel or Virtual PortChannel).

In a VMware vSphere virtualized environment, the Cisco Nexus 1000V can be used with Network Distributed Switch technology to pool networking resources from individual ESX servers into a single, easily manageable distributed virtual switch.

Layer 3 Network Aspects

- Layer 3 topology:
 - Domain boundary
 - Routing convergence
 - Default gateway placement and availability
- Layer 3 traffic engineering:
 - Path selection
 - ECMP load sharing
- Oversubscription points



© 2011 Cisco Systems, Inc. All rights reserved.


DCUCD v4.0-3-8

As mentioned, Layer 3 would not be typically used in the Cisco Data Center access layer, especially when VM mobility and service clustering is used. Thus, Layer 3 is used in aggregation and core network layers.

The network core-aggregation topology typically uses redundant links that terminate at multiple devices. Since topology is redundantly designed, there are loops. In addition, you want them. Of course, some mechanism must ensure that the traffic does not loop in the network. Typically, you would use routing for the purpose. To utilize all possible paths, the Equal-Cost Multipath (ECMP) can be employed to send the traffic over multiple links. When strict segregation between segments must be achieved, device-based virtualization can be used.

Network High Availability

- Link failure resolution:
 - L2 convergence—STP
 - Port channeling
 - Failover time
- Device failure:
 - Convergence upon failure
 - Redundancy inside a device
- Default gateway failure:
 - First-hop resolution mechanism
 - FHRP and Layer 2 topology synchronization
 - Load sharing between multiple gateways



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-9

Network high availability considers different points in the solution architecture. At base, the link high availability must be achieved—either by employing Spanning Tree Protocol (STP) or preferably with PortChannel, which not only achieves better failover time but also scales the bandwidth.

High availability is implemented using the following techniques:

- Redundant devices—active/standby or active/active setup
- Redundant parts inside a single device—redundant supervisors, redundant power supplies, redundant network modules, switch fabric, fans, and so on
- High-availability mechanisms—Stateful Switchover (SSO)/Nonstop Forwarding (NSF), Extensible Service Transfer Protocol (xSTP), Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), Gateway Load Balancing Protocol (GLBP)

Device High Availability

The device high availability must also be considered, since failure of a network device results in application or service degradation.

When trying to prevent device failure, you must consider the following:

- Speed of switchover from failed to standby device
- Where and how to deploy redundancy—whether inside the device, employing SSO to achieve subsecond switchover, or between the devices

Cisco Catalyst 6500 series switches support fault resistance by allowing a redundant supervisor engine to take over if the primary supervisor engine fails. Cisco NSF works with SSO to minimize the amount of time a network is unavailable to its users following a switchover, while continuing to forward IP packets. Cisco Catalyst 6500 series switches also support Route Processor Redundancy (RPR) and Single Router Mode with Stateful Switchover (SRM with SSO) for redundancy.

The Cisco Nexus 7000 platform also provides 1+1 redundant supervisor modules that can perform a supervisor switchover in critical failure situations. The switchover of active to standby supervisor is nondisruptive to the forwarding plane and can therefore provide nonstop forwarding of data during a supervisor switchover.

Network High Availability

Network high availability must be also considered from the server perspective. First Hop Redundancy Protocols (FHRPs) must be used to overcome broken connectivity upon default gateway failure. Of course, FHRP deployment must take into account the underlying topology—that is, they have to map and match in order to avoid suboptimal traffic flows.

First-Hop Failure Problems

The primary purpose of a network is to provide end users with access to their data and applications. End users typically do not care about routers, links, or LAN switches, or the fact that they are down. Their perception of an enterprise network is that it is a total system.

Redundancy in a network is fundamental for its high availability of delivering traffic. This is achieved with redundant paths and redundant equipment in key locations, which deal with the overall network availability. However, this does not guarantee the availability of vital resources to end users in case of a first-hop router failure.

The goal of IP redundancy is to protect hosts against such first-hop router failures, even when a source host cannot dynamically find an alternate first-hop router. Most end devices—workstations, servers, printers—are not capable of exchanging dynamic routing information and have only a single default gateway configured.

When a default gateway is configured on most devices, there is no means by which to configure a secondary gateway, even if a second route exists to carry packets off the local segment.

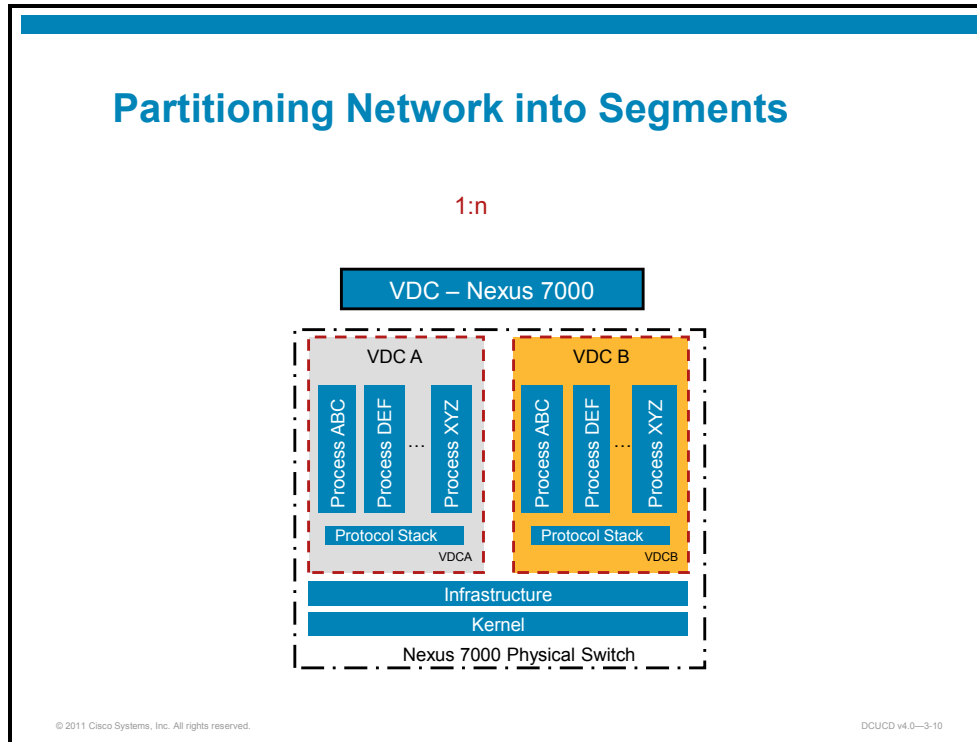
End devices are typically configured with a single default gateway IP address that does not change when network topology changes occur. If the router whose IP address is configured as the default gateway fails, the local device will be unable to send packets off the local network segment, effectively disconnecting it from the rest of the network. Even if a redundant router exists that could serve as a default gateway for that segment, there is no dynamic method by which these devices can determine the address of a new default gateway.

First-Hop Resolution Protocols

With this type of router redundancy, a set of routers works in concert to present the illusion of a single virtual router to the hosts on the LAN. By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a single "virtual" router. The IP address of the virtual router will be configured as the default gateway for the workstations on a specific IP segment. When frames are to be sent from the workstation to the default gateway, the workstation will use Address Resolution Protocol (ARP) to resolve the MAC address associated with the IP address of the default gateway. The ARP resolution will return the MAC address of the virtual router. Frames sent to the MAC address of the virtual router can then be physically processed by any active or standby router that is part of that virtual router group.

A protocol is used to identify two or more routers as the devices responsible for processing frames that are sent to the MAC or IP address of a single virtual router. Host devices send traffic to the address of the virtual router. The physical router that forwards this traffic is transparent to the end stations. The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router. The transition from one forwarding router to another is transparent to the end devices.

Partitioning Network into Segments



The Cisco Nexus Operating System (NX-OS) software supports virtual device contexts (VDCs), which partition a single physical device into multiple logical devices that provide fault isolation, management isolation, address allocation isolation, service differentiation domains, and adaptive resource management. You can manage a VDC instance within a physical device independently. Each VDC appears as a unique device to the connected users. A VDC runs as a separate logical entity within the physical device, maintains its own unique set of running software processes, has its own configuration, and can be managed by a separate administrator.

VDCs also virtualize the control plane, which includes all those software functions that are processed by the CPU on the active supervisor module. The control plane supports the software processes for the services on the physical device, such as the routing information base (RIB) and the routing protocols.

The benefits include VDC-level fault isolation, VDC-level administration, separation of data traffic, and enhanced security.

Fault Isolation

The VDC architecture can prevent failures within one VDC from impacting other VDCs on the same physical device. For instance, an Open Shortest Path First (OSPF) process that fails in one VDC does not affect the OSPF processes in other VDCs in the same physical device.

LAN Configuration Prerequisites

VLAN design

- Free VLAN IDs?
- 802.1Q trunk and allowed VLAN list
- VLAN database

Spanning tree

- UCS uses EHV by default
- PortFast on Cisco UCS attached interfaces

EtherChannel configuration

MAC address overlapping

```
N7k-core1#  
feature lACP  
interface Ethernet1/31  
switchport mode trunk  
switchport trunk allowed vlan 1,11-18  
channel-group 10 mode active
```

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-11

The Cisco Unified Computing System LAN configuration is also governed by the Cisco UCS configuration and design.

When adding Cisco UCS to the existing LAN, the upstream LAN configuration has to be examined. It must be verified that the VLAN IDs uses are free or that they do not overlap with the existing ones.

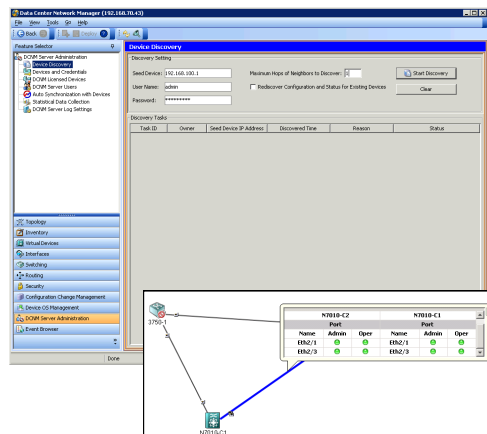
The interfaces connecting the Cisco UCS clusters have to be configured with 802.1Q trunk configuration allowing the VLANs that need to be carried from Cisco UCS to the core LAN network.

Second, Cisco UCS is normally operating in the EHV mode, meaning that the STP is not running on the system. To speed up the convergence of the LAN links where Cisco UCS is connected, the upstream LAN switches can be configured with the STP PortFast option on such links.

Third, Cisco UCS is extremely flexible in terms of configuration. One of the configurable parameters is the MAC addresses that are used by the servers. Thus, the existing LAN has to be verified for overlapping MAC addresses.

LAN Management

- Cisco Data Center Network Manager (DCNM) to manage:
 - Nexus 7000 Series Switches
 - VLANs, IP routing, VDCs, security, vPC, HSRP, and so on
- Management interfaces:
 - Graphical user interface (GUI)
 - Command-line interface (CLI)



In the Cisco Unified Computing solution, the LAN part of the solution will be built using Nexus 7000 as powerful core or distribution layer switches. Nexus 7000 switches can be configured with wide variety of features and protocols from VLANs, IP routing, security, Hot Standby Router Protocol (HSRP), to VDCs, and virtual PCs (vPCs). The application can automatically discover the LAN fabric topology.

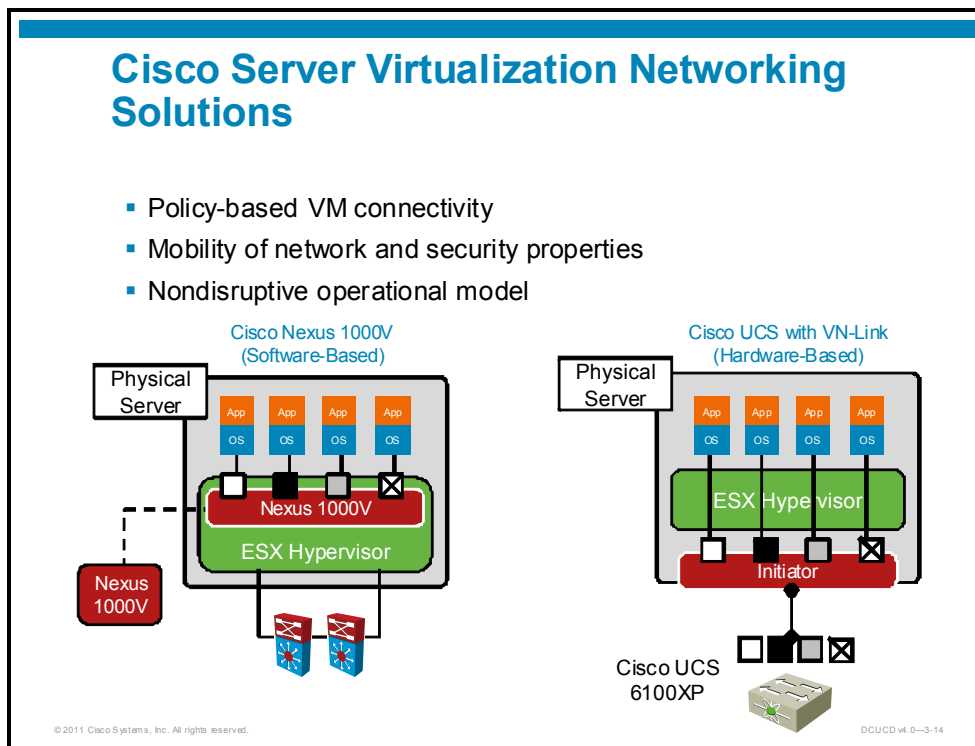
The amount of configuration depends on individual deployment, but with Cisco Data Center Network Manager (DCNM), all these aspects can be configured and managed from a single management application. Apart from the GUI-oriented DCNM application, one can also use command-line interface (CLI).

The DCNM application requires a server where LAN fabric information is stored.

More information about the DCNM application can be found in the course *Implementing Cisco Data Center Networking Infrastructure 2 (DCNI-2) v3.0*.

Cisco Nexus 1000V

This topic identifies and describes Cisco Nexus 1000V.



Cisco server virtualization solution uses technology jointly developed by Cisco and VMware. The network access layer is moved into the virtual environment to provide enhanced network functionality at the VM level.

This can be deployed as a hardware- or software-based solution, depending on the data center design and demands. Both deployment scenarios offer VM visibility, policy-based VM connectivity, policy mobility, and a nondisruptive operational model.

Cisco Nexus 1000V

The Cisco Nexus 1000V is a software-based solution providing VM-level network configurability and management. The Cisco Nexus 1000V works with any upstream switching system to provide standard networking controls to the virtual environment.

VN-Link

VN-Link technology was jointly developed by Cisco and VMware and has been proposed to the IEEE for standardization. The technology is designed to move the network access layer into the virtual environment in order to provide enhanced network functionality at the VM level.

Cisco UCS 6100XP

With the Cisco UCS 6100, VN-Link can be deployed as a hardware-based solution offering VM visibility, policy-based VM connectivity, policy mobility, and a nondisruptive operational model.


Cisco Nexus 1000V

Replaces VMware vDS

- Preserves existing VM management
- NX-OS and IOS look and feel management
- Compatibility with VMware features
- Additional features



Latest release 4.0(4)SV1(3)



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-15

Cisco Nexus 1000V Series Switches are virtual machine access switches that are an intelligent software switch implementation for VMware vSphere environments running the Cisco NX-OS operating system. Operating inside the VMware ESX hypervisor, the Cisco Nexus 1000V Series supports Cisco VN-Link server virtualization technology to provide the following:

- Policy-based virtual machine (VM) connectivity
- Mobile VM security and network policy
- Nondisruptive operational model for your server virtualization, and networking teams

The Cisco Nexus 1000V bypasses the VMware vSwitch with a Cisco software switch. This model provides a single point of configuration for the networking environment of multiple ESX hosts.

When server virtualization is deployed in the data center, virtual servers typically are not managed the same way as physical servers. Server virtualization is treated as a special deployment, leading to longer deployment time, with a greater degree of coordination among server, network, storage, and security administrators.

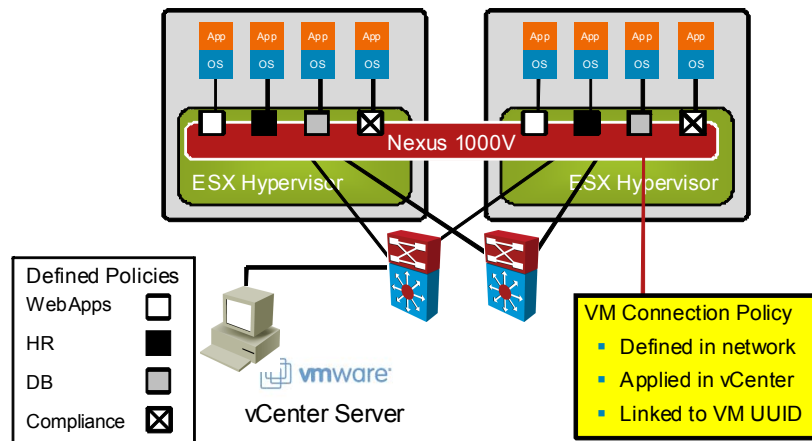
With the Cisco Nexus 1000V Series, you can have a consistent networking feature set and provisioning process all the way from the VM access layer to the core of the data center network infrastructure. Virtual servers can now leverage the same network configuration, security policy, diagnostic tools, and operational models as their physical server counterparts attached to dedicated physical network ports.

Virtualization administrators can access predefined network policy that follows mobile virtual machines to ensure proper connectivity saving valuable time to focus on virtual machine administration.

This comprehensive set of capabilities helps deploy server virtualization faster and realize its benefits sooner.

Cisco Nexus 1000V DVS

Policy-based VM connectivity

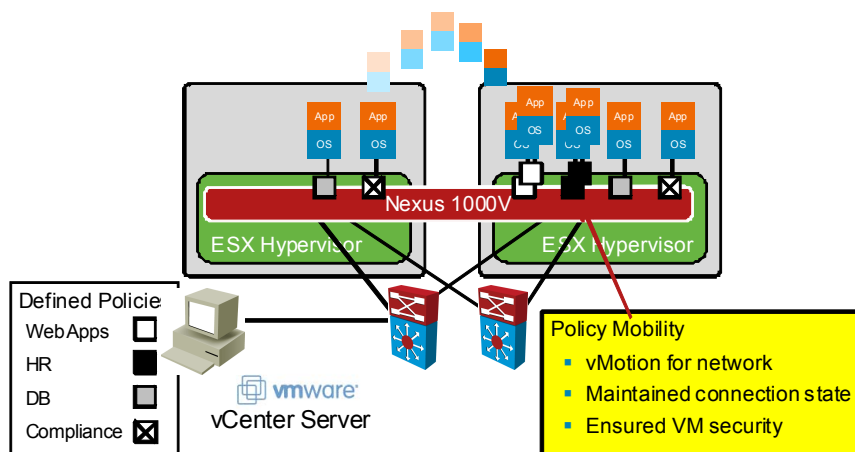


VM connection policies are defined in the network and applied to individual VMs from within VMware vCenter. These policies are linked to the Universally Unique ID (UUID) of the VM and are not based on physical or virtual ports.

To complement the ease of creating and provisioning VMs, the Cisco Nexus 1000V includes the Port Profile feature to address the dynamic nature of server virtualization from the network's perspective. Port Profiles enable you to define VM network policies for different types or classes of VMs from the Cisco Nexus 1000V VSM, then apply the profiles to individual VM virtual NICs through VMware's vCenter GUI for transparent provisioning of network resources. Port Profiles are a scalable mechanism to configure networks with large numbers of VMs.

Cisco Nexus 1000V DVS (Cont.)

Mobility of network and security properties



Through the VMware vCenter APIs, the Cisco Nexus 1000V monitors VM migration and ensures policy enforcement as machines transition between physical ports. Security policies are applied and enforced as VMs migrate through automatic or manual processes.


Network and security policies defined in the Port Profile follow the VM throughout its lifecycle whether it is being migrated from one server to another, suspended, hibernated, or restarted.

In addition to migrating the policy, the Cisco Nexus 1000V Virtual Supervisor Module also moves the VM network state, such as the port counters and flow statistics. VMs participating in traffic monitoring activities, such as Cisco NetFlow or ERSPAN, can continue these activities uninterrupted by vMotion operations.

When a specific Port Profile is updated, the Cisco Nexus 1000V automatically provides live updates to all of the virtual ports using that same Port Profile. With the ability to migrate network and security policies through vMotion, regulatory compliance is much easier to enforce with the Cisco Nexus 1000V, because the security policy is defined in the same way as physical servers and constantly enforced by the switch.

Cisco Nexus 1000V DVS (Cont.)

- Nondisruptive operational model
- Server benefits
 - Existing VM management preserved
 - Reduced deployment time and operational workload
 - Improved scalability
 - VM-level visibility
- Network benefits
 - Unified network management and operations
 - Improved operational security
 - Enhanced VM network features
 - Policy persistence
 - VM-level visibility



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-18


When using the Cisco Nexus 1000V, the management model for VMs stays the same and is handled by the VM administrator. Network administrators create security profiles and these policies are applied to individual VMs by the VM admin. Deployment time is reduced through pre-configured repeatable processes, reducing the operational workload. The benefits include unified network management and operations, enhanced network features at the VM level, and VM-level visibility.

Because of its close integration with VMware vCenter Server, the Cisco Nexus 1000V allows virtualization administrators to continue using VMware tools to provision VMs. At the same time, network administrators can provision and operate the VM network the same way they do the physical network using Cisco CLI and SNMP along with tools such as ERSPAN and NetFlow.

While both teams work independently, using familiar tools, the Cisco Nexus 1000V enforces consistent configuration and policy throughout the server virtualization environment. This level of integration lowers the cost of ownership while supporting various organizational boundaries among server, network, security, and storage teams.

Cisco Nexus 1000V Features

- Layer 2
 - VLAN, PVLAN, 802.1Q
 - LACP
 - vPC host mode
- QoS classification and marking
- Security
 - Layer 2, 3, 4 access lists
 - Port security
- SPAN and ERSPAN
- Compatibility with VMware
 - vMotion, Storage vMotion
 - DRS, HA, FT



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-19

Cisco Nexus 1000V supports the same features as physical Cisco Catalyst or Nexus switches while maintaining compatibility with VMware advanced services like vMotion, DRS, FT, HA, Storage vMotion, Update Manager, and vShield Zones.

vPC Host Mode

Virtual PortChannel host mode (vPC-HM) allows member ports in a port channel to connect to two different upstream switches. With vPC-HM, ports are grouped into two subgroups for traffic separation. If Cisco Discovery Protocol is enabled on the upstream switch, then the subgroups are automatically created using Cisco Discovery Protocol information. If Cisco Discovery Protocol is not enabled on the upstream switch, then the subgroup on the interface must be manually configured.

Layer 2 Features

The following Layer 2 features are supported by Nexus 1000V:

- Layer 2 switch ports and VLAN trunks
- IEEE 802.1Q VLAN encapsulation
- Link Aggregation Control Protocol (LACP): IEEE 802.3ad
- Advanced port channel hashing based on Layer 2, 3, and 4 information
- Virtual PortChannel host mode
- Private VLANs with promiscuous, isolated, and community ports
- Private VLAN on trunks
- Internet Group Management Protocol (IGMP) snooping versions 1, 2, and 3
- Jumbo frame support; up to 9216 bytes
- Integrated loop prevention with Bridge Protocol Data Unit (BDPU) filter without running Spanning Tree Protocol (STP)

QoS Features

The following QoS features are supported by Nexus 1000V:

- Classification per access group (ACL), IEEE 802.1p CoS, IP Type of Service: IP precedence or DSCP (RFC 2474), User Datagram Protocol (UDP) ports, Packet length
- Marking per two rate three color marker (RFC 2698), IEEE 802.1p CoS marking, IP Type of Service: IP precedence or DSCP (RFC 2474)
- Traffic policing (transmit- and receive-rate limiting)
- Modular QoS CLI (MQC) compliance

Security Features


The following security features are supported by Nexus 1000V:

- Ingress and egress ACLs on Ethernet and virtual Ethernet ports
- Standard and extended Layer 2 ACLs
- Standard and extended Layer 3 and 4 ACLs
- Port-based ACLs (PACLs)
- Named ACLs
- ACL statistics
- Cisco integrated security features
- Virtual service domain for Layer 4 through 7 services virtual machine

Cisco Nexus 1000V Architecture

VSM—Virtual Supervisor Module

- Management, monitoring, and configuration
- Integrates with VMware vCenter
- Uses NX-OS
- Configurable via CLI



Cisco Nexus 1000V is licensed per each server CPU regardless of the number of cores. It comprises the following:

- **Virtual Supervisor Module (VSM):** Performs management, monitoring, and configuration tasks for the Cisco Nexus 1000V and is tightly integrated with the VMware vCenter—the connectivity definitions are pushed from Cisco Nexus 1000V to the vCenter.
- **Virtual Ethernet Module (VEM):** Enables advanced networking capability on the ESX hypervisor and provides each VM with a virtual dedicated switch port.

A Cisco Nexus 1000V deployment consists of VSM (one or two for redundancy) and multiple VEMs installed in the ESX hosts—a vNetwork Distributed Switch.

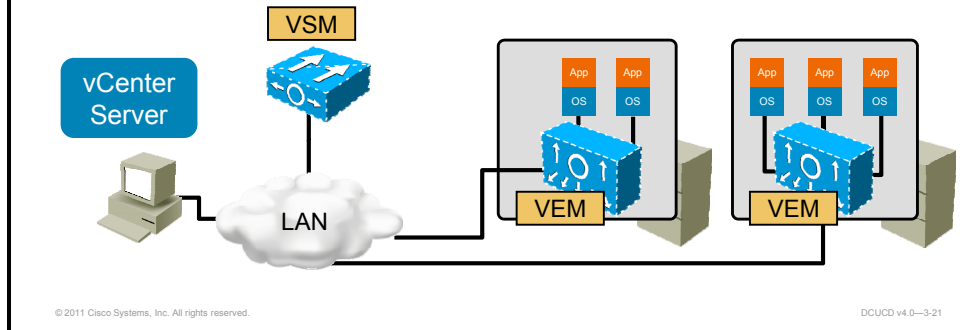
A VSM is a supervisor module much like in regular physical modular switches, whereas VEMs are remote Ethernet line cards to VSM.

In Cisco Nexus 1000V deployments, VMware provides the vNIC and drivers while the Cisco Nexus 1000V provides the switching and management of switching.

Cisco Nexus 1000V Architecture (Cont.)

VEM—Virtual Ethernet Module

- Replaces ESX virtual switch
- Enables advanced networking on ESX hypervisor
- Provides each VM with dedicated port



The VEM is a software replacement for the VMware vSwitch on a VMware ESX Host. All traffic-forwarding decisions are made by the VEM.

The VEM leverages the VMware vNetwork Distributed Switch (vDS) API, which was developed jointly by Cisco and VMware, to provide advanced networking capability to virtual machines. This level of integration ensures that the Cisco Nexus 1000V is fully aware of all server virtualization events, such as VMware vMotion and Distributed Resource Scheduler (DRS). The VEM takes configuration information from the Virtual Supervisor Module and performs layer 2 switching and advanced networking functions:

- PortChannel
- Quality of service (QoS)
- Security: Private VLAN, access control lists, port security
- Monitoring: NetFlow, Switched Port Analyzer (SPAN), Encapsulated Remote SPAN (ERSPAN)

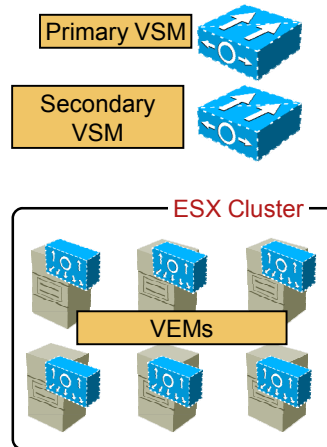
Cisco Nexus 1000V Architecture (Cont.)

ESX cluster = single Cisco Nexus 1000V switch

- Two VSM for high availability
- One VEM per ESX host

License per CPU socket for each VEM

- Installed on the VSM
- One or more license files can be installed
- Installation is nondisruptive to operation
- Dual supervisor environment runs licensed software on both



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-22

The Cisco Nexus 1000V uses distributed architecture. This architecture separates control and data plane functionality. The control plane functionality is represented by VSM, which manages multiple distributed data planes (VEM in each ESX host). Thus, a VSM acts as a supervisor module for remote VEMs.

All configuration and supervisor functions are handled by the VSM. Using Console, Telnet, or SSH, an administrator makes all configuration changes on the VSM. When a change is made on the VSM, the configuration is passed to vCenter and the changes are made on the distributed virtual switch (DVS). DVS changes are passed down to the corresponding VEM.

Each VEM will act as a module on the VSM and the VSM/VEM(s) appear as a single switch to Cisco Discovery Protocol neighbors. The VSM does not reside in the data path and therefore cannot directly receive or respond to Cisco Discovery Protocol messages. Cisco Discovery Protocol and other network management packets are transferred between the VEM to the VSM on one of three required VLANs, known as the Packet VLAN.

The Cisco Nexus 1000V Series is licensed based on the number of physical CPUs on the server on which the VEM is running.


VSM—Control Plane

Centralized control plane—VSM

- Manages multiple data planes
- Software appliance on physical server or VM
- Two VSMs for high availability

Maximum configuration:

- 64 VMware hosts per VSM
- 512 active VLANs
- 2048 vEth ports per vDS
- 216 vEth per host
- 32 physical NICs per host
- 256 port channels per vDS
- 8 port channels per host



VSM VM	Minimum Req.
OS	Other Linux (64-bit)
Memory	2 GB
vCPU	1
Network	3x e1000
Disk	3 GB

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-23

VSM can be deployed as a software appliance on a physical server (Control Plane Physical Appliance—CPPA) or on a VM (Control Plane Virtual Appliance—CPVA). A redundant Cisco Nexus 1000V deployment would incorporate two VSM appliances.

An individual Cisco Nexus 1000V supports the following:

- One data plane per ESX host
- Up to 64 ESX hosts
- 512 active VLANs
- 32 physical NICs per ESX host
- 2048 virtual Ethernet ports
- 216 virtual Ethernet ports per ESX host

The VSM has the following minimum system requirements when run as a VM:

- 2-GB memory
- Single vCPU
- Three e1000 network adapters named Management, Control, and Packet in the right sequence
- 4-GB disk

VEM—Data Plane


Distributed data plane—VEMs

- Each operates independently
- No address learning across VEM
- No backplane between VEM
- No forwarding from ingress line card to egress line card
- No EtherChannel across VEMs

Internal switching—VSM not required

- VSM failure does not fail data path
- Traffic continues to flow by VEM

No STP



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-24

Each VEM acts as a separate switching line card with no concept of a fabric between VEMs. This means there are no port channels or connections between VEMs and they do not rely on one another for operation. Switching decisions and frame forwarding all happen on the VEM and are not reliant on the VSM.

Only the uplinks in a host can be bundled in a port channel for load balancing and high availability. The Cisco Nexus 1000V does not support EtherChannels across different VEMs.

The Cisco Nexus 1000V does not run STP because it will deactivate all but one uplink to an upstream switch, preventing full utilization of uplink bandwidth. Instead, each VEM is designed to prevent loops in the network topology.

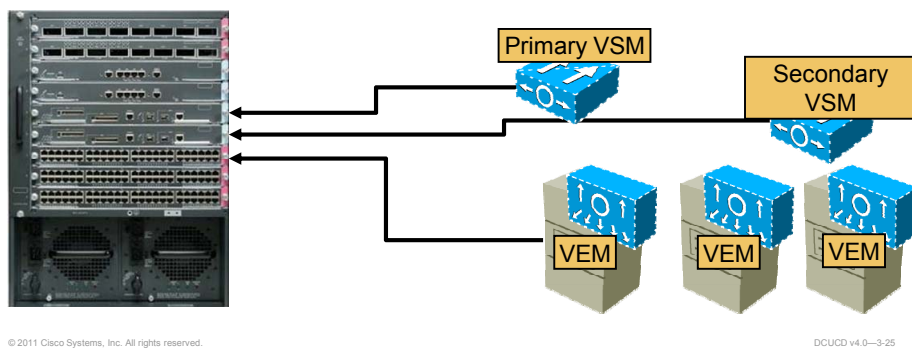
In the event of loss of communication with the Virtual Supervisor Module, the VEM has Nonstop Forwarding capability to continue to switch traffic based on last known configuration. In short, the VEM provides advanced switching with data-center reliability for the server virtualization environment.

Unlike with the VSM, VEM resources are unmanaged and dynamic. Although the storage footprint of the VEM is fixed (approximately 6.4 MB of disk space), RAM utilization on the VMware ESX host is variable, based on the configuration and scale of the Cisco Nexus 1000V Series deployment. In a typical configuration, each VEM can be expected to require 10 to 50 MB of RAM, with an upper hard limit of 150 MB for a fully scaled solution with all features turned on and utilized to their design limits.

Cisco Nexus 1000V Virtual Chassis

```
n1kv# show module
```

Mod	Ports	Module-Type	Model	Status
1	0	Virtual Supervisor Module	Nexus1000V	active *
2	0	Virtual Supervisor Module	Nexus1000V	ha-standby
3	248	Virtual Ethernet Module	NA	ok
4	248	Virtual Ethernet Module	NA	ok
5	248	Virtual Ethernet Module	NA	ok
... output omitted ...				



The Cisco Nexus 1000V Series uses a virtual chassis model to represent a pair of VSMs and their associated VEMs.

Like any Cisco chassis base platform, the Cisco Nexus 1000V Series virtual chassis has slots and modules, or line cards, associated with it. The VSMs are always associated with slot numbers 1 and 2 in the virtual chassis.

The VEMs are sequentially assigned to slots 3 through 66 based on the order in which their respective hosts were added to the Cisco Nexus 1000V Series Switch. With the 64 VEMs and the redundant supervisors, the Cisco Nexus 1000V can be viewed as a 66-slot modular switch.

When a VEM comes online for the first time, the VSM assigns the module number and tracks that module using the Unique User ID (UUID) of the VMware ESX server, helping ensure that if the VMware ESX host loses connectivity or is powered down for any reason, the VEM will retain its module number when the host comes back online.

The VSM maintains a heartbeat with its associated VEMs. This heartbeat is transmitted at 2-second intervals. If the VSM does not receive a response within 8 seconds, the VSM considers the VEM removed from the virtual chassis. If the VEM is not responding because of a connectivity problem, the VEM will continue to switch packets in its last known good state. When communication is restored between a running VEM and the VSM, the VEM is reprogrammed, causing a slight (1 to 15 seconds) pause in network traffic.

Cisco Nexus 1000V VMware Integration

This topic identifies and describes Cisco Nexus 1000V integration with VMware vCenter.

Cisco Nexus 1000V Operation

Cisco Nexus 1000V is distributed switch

- VSM programs VEM over the network
- Uses control messaging based on Nexus 7000 and MDS

Two options for VSM-to-VEM communication

- Using Layer 2 control and packet VLANs
- Using Layer 3 control capability (minimum 4.0(4)SV1(2) release required)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-27

In many ways, the Cisco Nexus 1000V Series Switches are similar to physical Ethernet switches. For packet forwarding, the Cisco Nexus 1000V Series uses the same techniques that other Ethernet switches apply, with a MAC address-to-port mapping table used to determine where packets should be forwarded.

The Cisco Nexus 1000V Series maintains forwarding tables in a slightly different manner than other modular switches. Unlike physical switches with a centralized forwarding engine, each VEM maintains a separate forwarding table. There is no synchronization between forwarding tables on different VEMs. In addition, there is no concept of forwarding from a port on one VEM to a port on another VEM. Packets destined for a device not local to a VEM are forwarded to the external network, which in turn may forward the packets to a different VEM.

MAC Address Learning

This distributed forwarding model within a centrally managed switch is demonstrated by the way the Cisco Nexus 1000V Series handles MAC address learning. A MAC address can be learned multiple times within a single Cisco Nexus 1000V Series Switch in either of two ways: statically or dynamically. Static entries are automatically generated for virtual machines running on the VEM—these entries do not time out. For devices not running on the VEM, the VEM can learn a MAC address dynamically, through the physical NICs in the server.

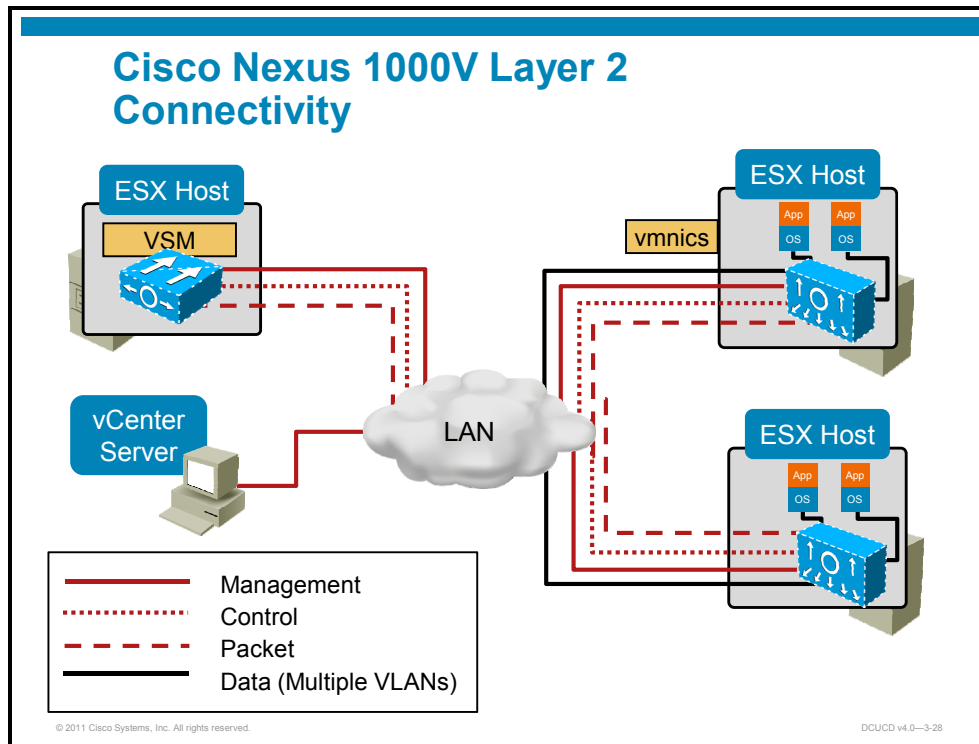
Each VEM maintains a separate MAC address table. Thus, a single Cisco Nexus 1000V Series Switch may learn a given MAC address multiple times, as often as once per VEM. For example, one VEM may be hosting a virtual machine, and the MAC address of the virtual machine will be statically learned on the VEM. A second VEM, in the same Cisco Nexus 1000V Series Switch, may learn the MAC address of the virtual machine dynamically. Thus, within the Cisco NX-OS CLI, you may see the MAC address of the virtual machine twice: a dynamic entry and a static entry.

Loop Prevention

Another differentiating characteristic of the Cisco Nexus 1000V Series is that it does not run STP. Although this might seem to be a significant departure from other Ethernet switches, potentially causing catastrophic network loops, in reality the Cisco Nexus 1000V Series implements a simple and effective loop-prevention strategy that does not require STP.

Because the Cisco Nexus 1000V Series does not participate in STP, it does not respond to Bridge Protocol Data Unit (BPDU) packets, nor does it generate them. BPDU packets that are received by Cisco Nexus 1000V Series Switches are dropped.

The Cisco Nexus 1000V Series uses a simple technique to prevent loops. Like a physical Ethernet switch, the Cisco Nexus 1000V Series performs source and destination MAC address lookups to make forwarding decisions. The VEM applies loop-prevention logic to every incoming packet on Ethernet interfaces. This logic is used to identify potential loops. Every ingress packet on a physical Ethernet interface is inspected to help ensure that the destination MAC address is internal to the VEM. If the destination MAC address is external, the Cisco Nexus 1000V Series will drop the packet, preventing a loop back to the physical network.



The VSM, VEM, vCenter, and VM connectivity uses dedicated VLANs—management, control, packet, and one or more data networks.

Like the VSM, each VEM has a control and a packet interface. These interfaces are unmanaged and not directly configurable by the end user. The VEM uses the opaque data provided by VMware vCenter Server to configure the control and packet interfaces with the correct VLANs. The VEM then applies the correct uplink port profile to the control and packet interfaces to establish communication with the VSM. There are two ways of connecting the VSM and the VEM.

If the VSM and the VEM are in the same Layer 2 domain, the best way to connect them is to use the Layer 2 connectivity mode.

Cisco Nexus 1000V Layer 2 Connectivity

Layer 2 connectivity required between VSM and VEMs

Management VLAN—OOB for VSM (mgmt0 port)

- Should be the same as vCenter and ESX management VLAN

Domain ID

- Single Nexus 1000V instance—dual VSM and VEMs

Control VLAN

- Exchange control messages between VSM and VEM

Packet VLAN

- Used for protocols like Cisco Discovery Protocol, LACP, IGMP

Data VLANs

- One or more VLANs for VM connectivity

Recommendation is to maintain separate VLANs

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-29

The Cisco Nexus 1000V VLANs consist of a management, control, packet, and one or more data networks.

Management VLAN

Each VMware ESX Host, the VSM, and the vCenter Server must all reside in the same management network and be part of the same Layer 2 domain.

Domain ID

A single Cisco Nexus 1000V instance, including dual redundant VSMs and managed VEMs, forms a switch domain. Each Cisco Nexus 1000V domain within a VMware vCenter Server needs to be distinguished by a unique integer called the domain identifier.

Control and Packet VLANs

Control and packet VLANs from the VSM must be accessible by uplink profiles on each VEM. The control VLAN and the packet VLAN are used for communication between the VSM and the VEMs within a switch domain.

The packet VLAN is used by protocols such as Cisco Discovery Protocol, LACP, and IGMP.

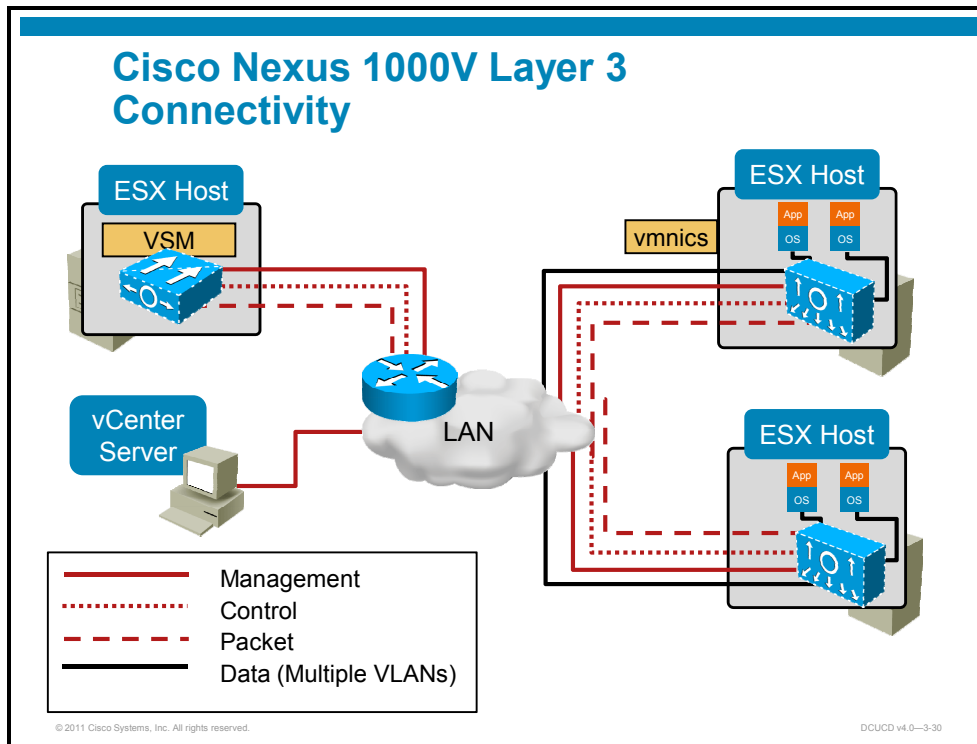
The control VLAN is used for the following:

- VSM configuration commands to each VEM, and their responses
- VEM notifications to the VSM; for example, a VEM notifies the VSM of the attachment or detachment of ports to the DVS
- VEM NetFlow exports are sent to the VSM, where they are then forwarded to a NetFlow Collector.

Data VLANs

The data networks carry VM packet traffic (server data). One or more data VLANs are defined for this purpose. Data traffic from the VM is not sent to the VSM and the VSM does not require access to the data VLANs. All VSM management is out-of-band and switching decisions do not rely on the VSM.

Note It is recommended that the control VLAN and packet VLAN should be separate VLANs, and that they should be on separate VLANs from those that carry data.



If the VSM and the VEM are in different Layer 2 domains, the Layer 3 connectivity mode should be used.

The Layer 3 mode encapsulates the packet of the Layer 2 mode using Generic Routing Encapsulation (GRE).

This process requires configuration of a VMware VMkernel interface on each VMware ESX host. This VMware VMkernel interface will need to be configured and attached to a port profile with the **l3control** option.

The **l3control** configuration tells the VEM that it can use this interface to send Layer 3 packets, so even if the Cisco Nexus 1000V Series is a Layer 2 switch, it can send IP packets.

Cisco Nexus 1000V Layer 3 Connectivity

Minimum release 4.0(4)SV1(2) required

Layer 3 connectivity between VSM and VEMs

- VSM—SVS Layer 3 mode with control or management interface
- VEM—VM kernel interface and GRE to tunnel control traffic to VSM
- Requires per VEM Layer 3 control port profile

Option 1—management interface

- OOB management for VSM—mgmt0 port
- Should be the same as vCenter and ESX management VLAN
- VSM-to-VEM traffic mixed with vCenter management traffic

Option 2—special control interface with own IP address

- Dedicated control0 interface for VSM-to-VEM communication

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-31

For Layer 3 connectivity between VSM and VEM, the following options are available:

- OOB management interface (mgmt0 port), which must be in the same VLAN as vCenter Server and ESX management interfaces. With this option the VSM-to-VEM traffic is mixed with vCenter management traffic.
- Dedicated control interface (Control0) with its own IP address for VSM-to-VEM communication.

Cisco Nexus 1000V Switch Interfaces

Ethernet port (**eth**)

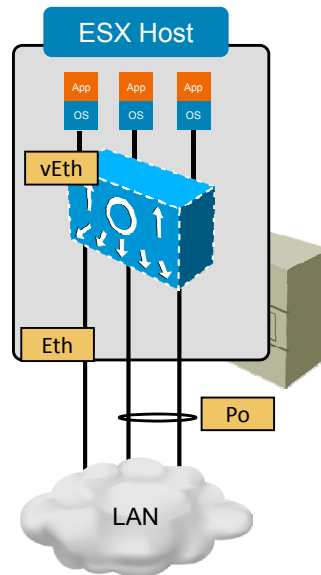
- ESX/ESXi **vmnic**

PortChannel (**po**)

- Supported on eth interfaces only
- Supports all Cisco and LACP port channels
- Created automatically if port profile used

Virtual Ethernet (**veth**)

- Access port of VM (1—2048)
- Autocreated when port profile attached to VM vNIC via vCenter
- No port number change on vMotion, DRS



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-32

The Cisco Nexus 1000V Series supports multiple switch-port types for internal and external connectivity: virtual Ethernet (vEth), Ethernet (Eth), and PortChannel (Po). The most common port type within a Cisco Nexus 1000V Series environment is a new concept called a virtual Ethernet interface. This interface type represents the switch port connected to a virtual machine's vNIC or connectivity to specialized interface types such as the vswif or vmknic interface.

vEth Interface

A vEth interface has several characteristics that differentiate it from other interface types. Besides the obvious fact that vEth interfaces are virtual and therefore have no associated physical component, the interface naming convention is unique.

Unlike a traditional Cisco interface, a vEth interface name does not indicate the module with which the port is associated. Whereas a traditional physical switch port may be notated as GigX/Y, where X is the module number and Y is the port number on the module, a vEth interface is notated like this: vEthY. This unique notation is designed to work transparently with VMware vMotion, keeping the interface name the same regardless of the location of the associated virtual machine.

The second characteristic that makes a vEth interface unique is its transient nature. A given vEth interface appears or disappears based on the status of the virtual machine connected to it. The mapping of a virtual machine vNIC to a vEth interface is static. When a new virtual machine is created, a vEth interface is also created for each vNIC of the virtual machine. The vEth interfaces will persist as long as the virtual machine exists. If the virtual machine is temporarily down (the guest operating system is shut down), the vEth interfaces will remain inactive but still bound to that specific virtual machine. If the virtual machine is deleted, the vEth interfaces will become available for connection to newly provisioned virtual machines.

Eth Interface

The Cisco Nexus 1000V Series contains two interface types related to the VMNICs (physical NICs) within a VMware ESX host.

An Ethernet, or Eth, interface is the Cisco Nexus 1000V Series representation of a VMNIC. An Eth interface is represented in standard Cisco interface notation (EthX/Y) using the Cisco NX-OS naming convention “Eth” rather than a speed such as “Gig” or “Fast,” as is the custom in Cisco IOS Software. These Eth interfaces are module-specific and are designed to be fairly static within the environment.

PortChannel Interface

PortChannel interfaces are the third interface type supported by the Cisco Nexus 1000V Series. A PortChannel interface is an aggregation of multiple Eth interfaces on the same VEM.

Cisco Nexus 1000V Port Profiles

Port profiles


- To configure multiple similar ports
- Define VLAN, ACL, QoS, port security, etc.
- VMware port-group equivalent

Uplink port profiles (System, VM uplink)

- Outbound connectivity from VEM
- Used for VEM to VSM and VM data traffic
- Assigned to vmnic (physical NIC)

VM port profiles

- Provide configuration for VM ports
- Assigned to VM vNIC



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-33

In the Cisco Nexus 1000V, port profiles are used to configure interfaces. A port profile can be assigned to multiple interfaces, giving them all the same configuration. Changes to the port profile can be propagated automatically to the configuration of any interface assigned to it.

In the VMware vCenter Server, a port profile is represented as a port group. The vEthernet or Ethernet interfaces are assigned in vCenter Server to a port profile for the following:

- Defining port configuration by policy
- Applying a single policy across a large number of ports
- Supporting both vEthernet and Ethernet ports

Port Profile Configuration

A port profile is a set of interface configuration commands that can be dynamically applied to either the physical (uplink) or virtual interfaces. A port profile can define a set of attributes including the following:

- VLAN
- Port channels
- Private VLAN (PVLAN)
- Access control list (ACL)
- Port security
- NetFlow
- Rate limiting
- Quality of service (QoS) marking

The network administrator defines port profiles in the VSM. When the VSM connects to vCenter Server, it creates a DVS, and each port profile is published as a port group on the DVS. The server administrator can then apply those port groups to specific uplinks, VM vNICs, or management ports, such as virtual switch interfaces or VM kernel NICs.

A change to a VSM port profile is propagated to all ports associated with the port profile. The network administrator uses the Cisco NX-OS CLI to change a specific interface configuration from the port profile configuration applied to it. For example, a specific uplink can be shut down or a specific virtual port can have ERSPAN applied to it, without affecting other interfaces using the same port profile.

Note Although the configuration can be applied to the individual virtual port on Cisco Nexus 1000V, it is recommended to apply the entire configuration via port profiles.

By using port profiles, hundreds or thousands of VMs can be provisioned rapidly with detailed network configurations such as port state, QoS tagging, and ACL controls. Additionally, port profiles reduce the risk of misconfiguration on groups of similar ports by maintaining the configuration in a central location. While individual port configuration is still possible using the Cisco Nexus 1000V, it is recommended to use port profiles rather than configuring ports directly. This method reduces administration time and simplifies network troubleshooting.

Uplink Port Profiles

The server administrator can assign port profiles that are configured as uplinks to physical NICs.

System Uplink Port Profiles

When a server administrator adds a host to the DVS, the VEM in that host needs to be able to configure the VSM. Since the ports and VLANs for this communication are not yet in place, system port profiles and system VLANs are configured to meet this need. VSM sends minimal early configuration to the vCenter Server, which then propagates it to the VEM when the host is added to the DVS.

A system port profile is designed to establish and protect vCenter Server connectivity. It can carry the following VLANs:

- System VLANs or vNICs used when bringing up the ports before communication is established between the VSM and VEM
- The uplink that carries the control VLAN
- Management uplinks used for VMware vCenter Server connectivity or SSH or Telnet connections. There can be more than one management port or VLAN—for example, one dedicated for vCenter Server connectivity, one for SSH, one for SNMP, a switch interface, and so on.
- VMware kernel NIC for accessing VMFS storage over Internet Small Computer Systems Interface (iSCSI) or NFS.

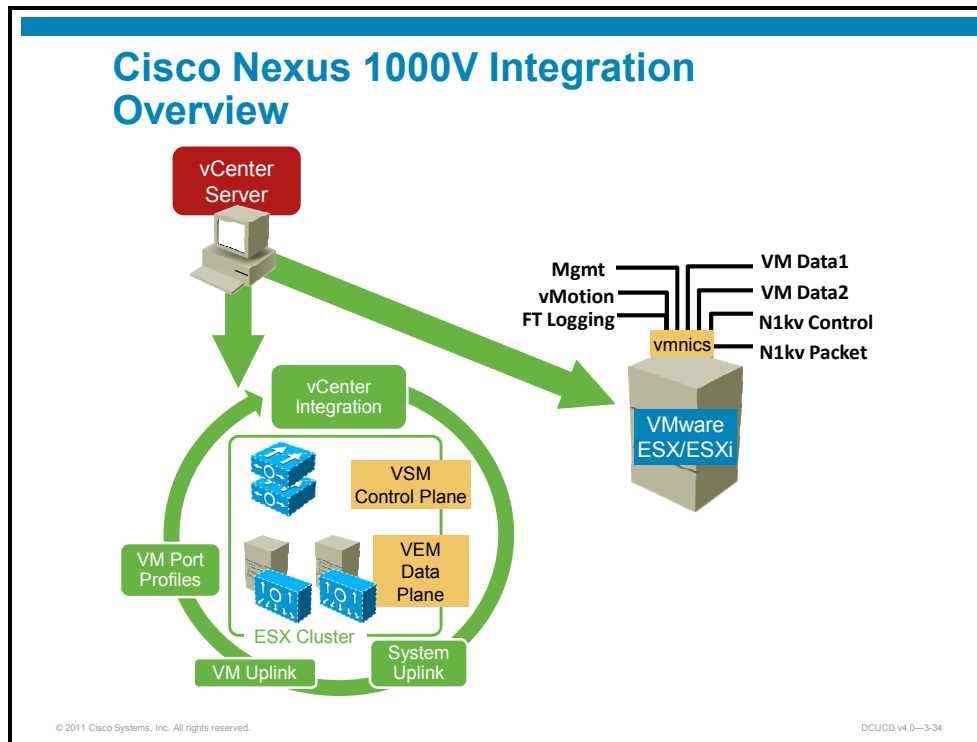
VM Uplink Port Profiles

VM uplink port profiles are used to define VM uplink characteristics that are separate from the control and packet VLANs.

VM Port Profiles

VM profiles are used to provide configuration for VMs and typically require an uplink profile to access the physical network.

When a VM profile is configured without a corresponding uplink profile, it creates internal VM networks. If a VM profile is created accessing a VLAN that is not trunked to the physical network, then the assigned VMs will be able to communicate only with other VMs assigned to the profile in the same host. This configuration is similar to creating internal-only vSwitches or port groups within a standard VMware networking environment.



The Cisco Nexus 1000V Series is tightly integrated with VMware vCenter. This integration enables the network administrator and the server administrator to collaborate efficiently without each having to learn a different management tool. The network administrator uses the Cisco NX-OS CLI on the VSM, and the server administrator continues to use VMware vCenter. Because of the tight relationship between the VSM and VMware vCenter, the communication between the two needs to be reliable and secure.

Nexus 1000V and vCenter Server Communication

The VSM maintains a link to VMware vCenter Server that is used to maintain the definition of the Cisco Nexus 1000V Series within VMware vCenter Server as well as propagate port profiles. The server and network administrators both have roles in establishing the link between the Cisco Nexus 1000V Series and VMware vCenter Server.

After installation is finished, the communication between the VSM and VMware vCenter Server is enabled under the `svs` configuration. The installer application will register the VSM plug-in with VMware vCenter Server, which will establish the link and create the instance of the Cisco Nexus 1000V Series within VMware vCenter Server.

Each VSM contains a unique extension key used to bind that specific VSM to VMware vCenter Server.

In creating the Cisco Nexus 1000V Series within VMware vCenter Server, the VSM propagates any port profiles that are already defined, as well as important information required for VEM installation called opaque data. The opaque data provides limited configuration details to the VEM so that it can communicate with the VSM after installation. The VSM is considered the authoritative container for all configuration information. If the connection between the VSM and VMware vCenter Server is disrupted, the VSM helps ensure that any configuration changes that have been made during this period of disrupted communication are propagated to VMware vCenter Server when the link is restored.

After the connection between the VSM and VMware vCenter Server is established, the link is primarily used to propagate new port profiles and any changes to existing port profiles.

Cisco Nexus 1000V Series VMware vCenter Server Extension

VMware vCenter Server is an extensible application that allows third-party management plug-ins, thus enabling external applications to extend the capabilities of VMware vCenter Server and its companion GUI, VMware vSphere Client.

The Cisco Nexus 1000V Series uses a VMware vCenter Server extension to properly display a representation of the Cisco Nexus 1000V Series and its main components within VMware vSphere Client.

The Cisco Nexus 1000V Series extension is a small XML file (`cisco_nexus_1000v_extension.xml`) that is directly installed and registered using the installation GUI. The plug-in can also be downloaded from the management IP address of the VSM using a web browser. This plug-in must be installed before the VSM can establish a link to VMware vCenter Server.

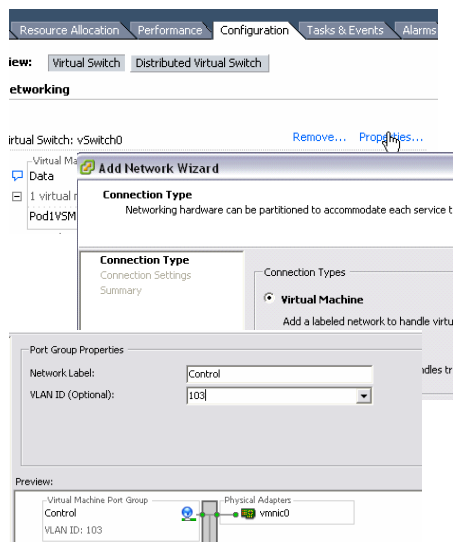
Prepare VMware Hosts for Nexus 1000V

Define management, control, packet VLANs

- Required when VSM is running as VM

Create VM port groups for each VLAN on VSM hosts

VLAN	ID	Port Group
Mgmt	1	Management
Control	999	Control
Packet	998	Packet



To deploy and integrate Cisco Nexus 1000V with VMware vCenter and hosts, the following prerequisites have to be met:

- vCenter Server installed and prepared.
- The VMware Enterprise Plus license must already be installed on the hosts.
- All hosts for the VEM must be running ESX or ESXi 4.0.
- There is at least one ESX or ESXi 4.0 host. If you plan to use vMotion, you need two ESX or ESXi 4.0 hosts.
- An ESX host is available to run the VSM VM.
- The VSM virtual machine can be hosted on the VEM in an ESX host that it is managing, or it can be hosted on a separate ESX or ESXi host (3.5 or 4.0) running the regular VMware vSwitch.
- The ESX host requires a minimum of 4 GB of physical RAM to host a VSM VM as the ESX server alone requires a minimum of 2 GB of physical RAM. Additional memory may be required to run the vCenter Server VM on the same host.
- Each host has a minimum of the following physical NICs (PNICs)—one PNIC for a Service Console or management, one PNIC for the traffic between VSM and VEM and for VM data traffic.
- All ESX hosts within a Cisco Nexus 1000V must have Layer 2 connectivity to each other.
- If you are using a set of switches, make sure that the inter-switch trunk links carry all relevant VLANs, including control and packet VLANs. The uplink should be a trunk port carrying all VLANs configured on the ESX host.
- On the host running the VSM VM, the control and packet VLANs are configured through the VMware switch and the VMNIC.
- The optional VMware Update Manager (VUM) manages the Cisco Nexus 1000V software installation for ESX hosts in a data center.

- The VSM VM control interface must be on the same Layer 2 VLAN as the ESX 4.0 host that it manages. If you configure Layer 3, then you do not have this restriction. In each case however, the two VSMS must run in the same IP subnet.
- You have completed configuring valid network mapping for control, management, and packet VLAN access.
- The host for the VSM VM runs on 64-bit server hardware and can run ESX 4.0 or ESX 3.5.

Cisco Nexus 1000V Requirements

Dedicated VLANs

- Management, control, packet

High availability solution

- Primary and secondary VSM on different hosts

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-36

Planning the Cisco Nexus 1000V deployment must take into consideration Cisco Nexus 1000V, VMware, and uplink switch aspects.

From the Cisco Nexus 1000V perspective, the following must be addressed:

- **Licensing:** Cisco Nexus 1000V is licensed per server CPU. The designer must thus know how many ESX hosts will be initially used and how many will be used in the future to plan and select the proper licensing pack.
- **VLAN scheme:** For the Cisco Nexus 1000V deployment, a minimum of three VLANs are required—management, control, and packet. The designer must reserve and assign these VLAN IDs from the free VLAN ID pool. Although the VLAN IDs could be any of those already used, it is recommended to separate VLAN IDs.
- **Design VSM deployment.**

VSM Deployment Design

VSM deployment design must address the following Cisco Nexus 1000V aspects:

- Select the ESX host or hosts where the VSM appliance will be running—the host should have sufficient resources for the VSM VM.
- Define the management parameters like management IP address and login credentials.
- Define the domain ID parameters for multiple Cisco Nexus 1000V deployments. A unique domain ID per Cisco Nexus 1000V domain is recommended.
- If you must use the same control and packet VLAN pair for multiple domains, you must ensure that their domain identifiers are different.
- Define VSM restrictions: prohibited vMotion, disabled Distributed Resource Scheduler (DRS), and Fault Tolerance (FT) for the VSM appliance.
- Define system uplink port profile for control and packet VLANs.

- Define the redundancy scheme for the VSM—standalone or active-standby (primary-secondary). Select the ESX host for the secondary VSM appliance. The selected host should be different from the one selected for the primary VSM.

Cisco Nexus 1000V Integration Overview

Cisco Nexus 1000V installed as VM

Integration steps:

1. Install Cisco Nexus 1000V VSM.
2. Generate control plane certificate and authenticate to vCenter.
3. Create appropriate VLANs (management, control, packet).
4. Install VEM on VMware hosts (manually or with VUM).
5. Assign vmnic to VEM.
6. Add VMware hosts to Nexus1000V vDS.
7. Create Nexus 1000V configuration.
8. Create port profile for VM connectivity.
9. Add standby VSM (optional).

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-37

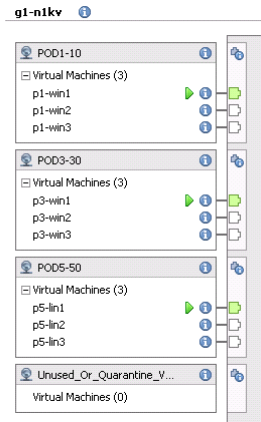
The process of deploying and integrating the Cisco Nexus 1000V can be divided into nine steps, as listed in the figure.

On the VMware side, the physical NICs have to be selected for the system uplink port profiles to enable VSM to VEM communication.

On the uplink switch side, the interface where ESX is attached has to be configured for 802.1Q trunking and must allow the expected VLANs—control, packet, and management.

VM Deployment Considerations

- Cisco Nexus 1000V considerations
 - Define data VLANs
 - Define VM uplink port profile—trunk data VLANs
 - Define VM port profile—for VM connectivity
 - Per VM policies—security, QoS, etc.
- VMware vCenter considerations
 - Select vmnic for VM uplink port profile for VM connectivity
- Upstream switch considerations
 - ESX attached interface for trunk with allowed VLANs



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-38

Inside VMware vCenter Server, VMs are configured as before. Instead of defining network configuration in vCenter Server, port profiles defined on the Cisco Nexus 1000V Virtual Supervisor Module and are displayed by vCenter as port groups.

Virtualization administrators can take advantage of preconfigured port groups and focus on VM management, while network administrators can use port profiles to apply policy for a large number of ports at the same time. Together, both teams can deploy server virtualization more efficiently and with lower operational cost.

Once the Cisco Nexus 1000V deployment is designed, the VM deployment can be planned. This includes Cisco Nexus 1000V, VMware, and upstream switch settings.

Cisco Nexus 1000V Design for VM Deployment

The design should define the following:

- A VLAN scheme for data traffic from different VMs. The scheme would typically define multiple VLANs for VM connectivity.
- VM port profiles for VM connectivity. A common VM profile for applications with the same connectivity requirements should be defined.
- A per-VM port profile policy, which includes QoS, security, and other settings.

VMware Design for VM Deployment

The design should specify the port groups derived from port profiles for different VMs. The port group description should specify the connectivity policy, which is derived from the port profiles.

Upstream Switch Design

The design should specify the interface settings of the upstream switch where the ESX server is connected for the VMNIC carrying data VLANs.

Cisco UCS SAN

This topic identifies and describes Cisco UCS LAN products and concepts.

SAN Aspects

- Upstream SAN switches
 - Must enable NPIV
- VSANs
 - Must be created
 - VSAN overlapping verification
- Interfaces
 - Disable trunking
 - Proper access VSAN per port or trunk configuration (1.4 software version)
- Zoning
 - Add Cisco UCS servers to zones configuration
- 1.4 software version
 - Direct attached storage support (default zoning equals permit)

```
vsan 1 wwn 50:06:01:60:4b:a
| [emc2-spa]
vsan 1 wwn 50:06:01:60:4b:a
| [emc3-spa]
vsan 11 wwn 50:06:01:60:4b:
| [emc1-spa]
vsan 11 wwn 50:06:01:60:4b:
| [emc2-spa]
vsan 11 wwn 50:06:01:60:4b:
| [emc3-spa]
vsan database
vsan 11 interface fc1/1
vsan 11 interface fc1/2
vsan 11 interface fc1/3
vsan 11 interface fc1/4
vsan 11 interface fc1/5
vsan 11 interface fc1/6
vsan 11 interface fc1/7
vsan 11 interface fc1/8
vsan 11 interface fc1/9
vsan 11 interface fc1/10
vsan 11 interface fc1/11
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
```

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-40

The Cisco Data Center Unified Computing System solution storage component provides space for servers to store data. It consists of the servers, the SAN, and storage devices (such as disk arrays and tape libraries).

These are key aspects addressed by storage design and deployment:

- **Connectivity:** If there is no connectivity or if connectivity is limited, service will suffer and consequently productivity will be low.
- **Security:** SAN security is typically neglected. Although the SAN is inside the network, it is still susceptible to intentional or unintentional attacks.
- **High availability:** Storage must be designed and in such a way as to easily overcome failure. As with the network, the goal is zero downtime.
- **Capacity:** This defines the amount of space available to the servers and the oversubscription ratio between the storage and server connectivity (fan-in ratio).

SAN Configuration Prerequisites

Firmware

Supported version

NPIV feature

VSAN design

Free VSAN IDs?

Zoning

Zones and zonesets

Include Cisco UCS information

WWN address overlapping

```
zoneset name zset11 vsan 11

zone name zone11
member pwn 20:00:20:b0:ca:fe:00:01
member pwn 21:31:00:02:ac:00:07:0f
member pwn 20:32:00:02:ac:00:07:0f

zoneset activate name zset11 vsan 11
```

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-41

To enable SAN communication, the prerequisites listed in the figure need to be fulfilled.

SAN Prerequisites

Connectivity

- 1/2/4/8 Fibre Channel for Cisco UCS uplink connectivity
- Short vs. long range interfaces
- Interfaces should be available on different linecards to improve high availability

Oversubscription levels

- Oversubscribed vs. full-rate port mode

Redundancy levels

- Dual-fabric design
- Dual hardware elements



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-42

SAN Oversubscription

Oversubscription is typically done at the storage connectivity level. The number of servers along with the throughput of the server connection surpasses the storage connectivity. Typically, the servers were not capable of producing the same amount of traffic that storage devices can receive until recently—or, more accurately, until virtualization is used.

With virtualization, a single physical server produces a larger amount of traffic, thus storage also receives more traffic. So in order to deploy a virtualization solution without storage access problems, storage connectivity must be properly sized also.

SAN High Availability


At the SAN link level, the high availability is handled by deploying multiple links. There are two options for how the redundant connections are utilized:

- Without any special configuration, the FSPF routing protocol used in Fibre Channel networks can handle the failover to the backup link when the primary fails. Although this happens automatically, it can result in fabric reconfiguration, which is disruptive for storage connectivity.
- A better approach, which includes also bandwidth scaling, is to deploy PortChannel. With PortChannel, the traffic is balanced over the links in a bundle while the individual physical ports are not presented to FSPF. FSPF sees the port channel as a single connection, thus no FSPF-based convergence happens upon link failure. The failover time is better than in the first case.

From the server and storage device perspective, the devices are connected to two fabrics to achieve end-to-end storage connectivity redundancy. Typically, the multipathing software used on servers for the host bus adapters (HBAs) handles this in an active/standby fashion; i.e., one fabric is actively used, while the second one is used only when the path through the primary fails. Active/standby fabric design must be properly done in order to utilize resources of both—that is, since multiple servers are connected to the SAN, they belong to different fabrics (VSAN-based), and with proper configuration the traffic utilizes both.

Virtualization with VSAN

- SAN island consolidation
- Fabric ports utilization increase
- Support for large fabrics deployment:
 - Dividing a large fabric into smaller
 - Isolation of disruptive events per VSAN
 - Advanced traffic management



© 2011 Cisco Systems, Inc. All rights reserved.

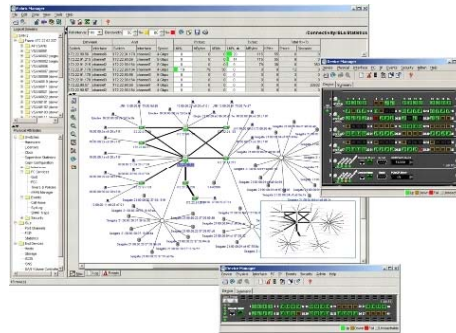
DCUCD v4.0-3-43

SAN networking is used to connect servers to the central storage. Similarly to VLANs in a data network, VSANs can be used to isolate storage networking, thus giving the ability to use a single physical infrastructure while maintaining logical separation end to end (from servers to storage arrays).

VSANs are an integral part of the Cisco UCS configuration. They are used to segregate server SAN traffic and to allow Cisco UCS to connect to different logical fabrics.

SAN Management

- Cisco Fabric Manager and Cisco Device Manager to manage:
 - MDS 9000 Series SAN and Nexus 5000 Series Switches
 - VSANs, zoning database, traffic engineering, security, etc.
- Management interfaces:
 - GUI
 - CLI



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-44

If a data center is built using Cisco MDS and/or Nexus series switches, they can be managed using the Cisco Fabric Manager and Cisco Device Manager.

The management applications enable an administrator to address all the aspects of the SAN environment—VSAN configuration and management, Fibre Channel port provisioning, zoning configuration and management, traffic engineering if required, security, and physical devices inventory. The application offers automatic fabric discovery with all the attached devices.

Apart from the GUI-oriented Cisco Fabric Manager and Cisco Device Manager, the MDS 9000 switches can also be managed using CLI.

The Cisco Fabric Manager requires a server where the database that stores the fabric information resides.

More information about Cisco Fabric Manager and Cisco Device Manager applications can be found in the course *Implementing Cisco Storage Networking Solutions (ICSNS) v3.0*.

Cisco UCS Storage

This topic identifies and describes Cisco UCS storage products and concepts.

Storage Prerequisites

Resources

- Can additional LUNs be created?
- Free space


LUN masking

- Expose proper LUNs to Cisco UCS
- Include Cisco UCS WWN information

Firmware

- Supported version

Multipathing support



© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-3-46

The servers can connect to network-attached storage (NAS) or SAN remote storage, which governs the protocols used—the Common Internet File System (CIFS) or Network File System (NFS) protocols can be used to access the NAS remote storage, whereas for the SAN, typically Fibre Channel or iSCSI protocols are utilized.

In enterprise environments where storage is a vital resource, the remote storage is usually available via SAN, based on the Fibre Channel protocol.

Apart from the protocol used, the server storage connectivity also defines other connectivity characteristics.

The high availability is also affected by the storage device. It is of no use to deploy multiple links from the server, if the failure of a disk in storage device results in storage being unavailable. For that purpose, different disk redundancy schemes are deployed; for example, Redundant Array of Independent Disks (RAID) level.

The link utilization and oversubscription affect the storage connective service quality and need to be properly designed. The throughput or bandwidth that is required on the link from the server depends on the amount of traffic that a server is capable of producing. For instance, in a server virtualization environment where a single physical server is hosting multiple virtual machines, the amount of storage traffic from the server would be higher than in the case of single-application server deployment. In cases of high throughput requirement, multiple links connecting the server to the storage might be deployed.

Since remote storage is vital for server operation (a loss of connectivity might result in an application being unavailable), redundant connectivity to the remote storage is typically provided. If a single HBA is sufficient from the throughput perspective, a second HBA is used to implement redundant connectivity. However, providing the redundant HBA only is not enough. To be able to utilize redundant connectivity through the second HBA, multipathing software must be installed on the server.

There are also some server storage connectivity challenges that need to be addressed, depending on the solution used.

In the server virtualization environment, the virtual machines are sharing the physical server HBA. If no special configuration is used, a virtual machine ends up using the same world wide name (WWN) as other virtual machines and thus also the same Fibre Channel ID (FC-ID), which in reality means that there is no distinction between different virtual machines in the storage network. In some cases a virtual machine needs to access a dedicated logical unit number (LUN) and a WWN and an FC-ID shared with other virtual machines cannot be used to restrict LUN access to that particular virtual machine only. In such cases, an N_port identifier virtualization (NPIV) mechanism is used with which the WWN and FC-ID granularity per virtual machine is achieved.

With the advent of blade servers, where a blade chassis hosts various I/O modules (possibly including a Fibre Channel switch), the large computing solutions reach the Fibre Channel SAN topology limitation. This is the maximum number of Fibre Channel switches by the DomainID pool (from 1 to 239). To overcome that limitation, the Fibre Channel switches that are installed in the blade chassis are configured to operate in N_port virtualization (NPV) edge mode. In this mode the Fibre Channel switch operates like an HBA—that is, all Fibre Channel SAN services such as name server and fabric login are operated by the upstream NPV core switch to which the NPV edge switch connects. This way, the blade chassis Fibre Channel switch does not consume the DomainID, which prevents DomainID saturation.

Finally, yet importantly, a server accessing a remote storage via Fibre Channel SAN typically would access a LUN, which is a logical volume that resides on a disk array. Physically the LUN is typically spread across the disk array disk drives, and the disk array addresses the redundancy and high availability for the LUN. With LUNs, the internal disk array architecture is hidden from the server.


SAN is the network which connects servers to the storage devices with which storage space is decoupled from the physical server and better storage utilization is achieved.

As with the LAN, storage also requires proper design and has certain requirements:

- The fault domain size must be properly selected, as too large a domain can result in improper storage connectivity. To achieve segmentation virtual storage area networks (VSANs) can be used—VSANs virtualize connections between Fibre Channel switches, which allow transfer of traffic for multiple VSANs over a single link. VSANs virtualize SAN switches as well by segregating the domains into separate fabrics.
- Storage space is presented in the form of LUNs to the servers. All servers should not be capable of connecting to all LUNs. SAN-based LUN masking can be used to tie the individual LUN exposure to a certain server(s).
- Storage connectivity can be achieved using different protocols—Fibre Channel, iSCSI, or Fibre Channel over Ethernet (FCoE). The benefit of FCoE over Fibre Channel or iSCSI is that it minimizes the deployment and management overhead and costs from the server perspective while still preserving the same level of service.

Storage Virtualization

- Segment physical storage to LUNs.
- Map LUNs to server virtual volumes:
 - Hosts and applications see virtual volumes.



© 2011 Cisco Systems, Inc. All rights reserved.

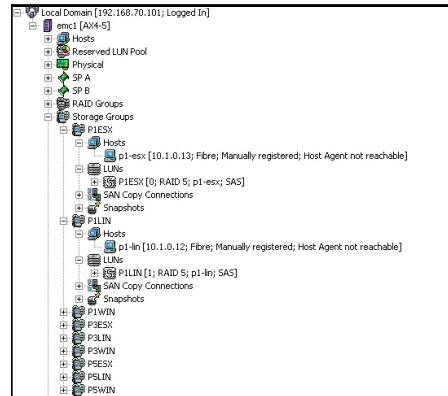
DCUCD v4.0-3-47

One of the most basic virtualization techniques in a centralized physical storage is to create LUNs. Traditionally, one LUN would be mapped to one application server. In VMware environments, however, a single LUN can be mapped to a cluster of physical ESX servers.

From the perspective of Cisco UCS configuration and server deployment, LUNs are an integral part of the configuration when SAN boot functionality is deployed.

Storage Management

- Storage management to application to manage drive arrays:
 - SAN connectivity
 - Managing LUNs and LUN masking
 - Binding hosts and LUNs
 - Performing backup and restore, expanding volumes
- Management through:
 - GUI
 - CLI



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4. 0-3-48

An important part of the Cisco Unified Computing System solution management application portfolio is also storage management. Depending on the storage solution vendor, different applications are used. In the case of EMC drive arrays, the EMC NaviSphere management can be used to provision and manage storage. The application enables the administrator to manage physical space and SAN connectivity, to create LUNs and define LUN masking, and to bind hosts (servers) with LUNs.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Network high availability is important from the Cisco UCS perspective—the solution must be resilient to device failures utilizing redundant hardware and must also be able to respond to topology changes.
- Cisco Nexus 1000V architecture provides a virtual switch chassis with a supervisor module (VSM) and switching line cards (VEM).
- Cisco Nexus 1000V deployment requires management, control, packet, and data VLANs.
- Cisco UCS systems can be combined with Cisco Nexus 1000V technology to provide physical network controls at the virtual network level.
- The NPIV feature is required to connect the Cisco UCS cluster to the SAN network.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-49

Lesson 7

Designing the Cisco UCS Network and Storage

Overview

This lesson discusses the design of the LAN and SAN components of the Cisco Unified Computing System solution.

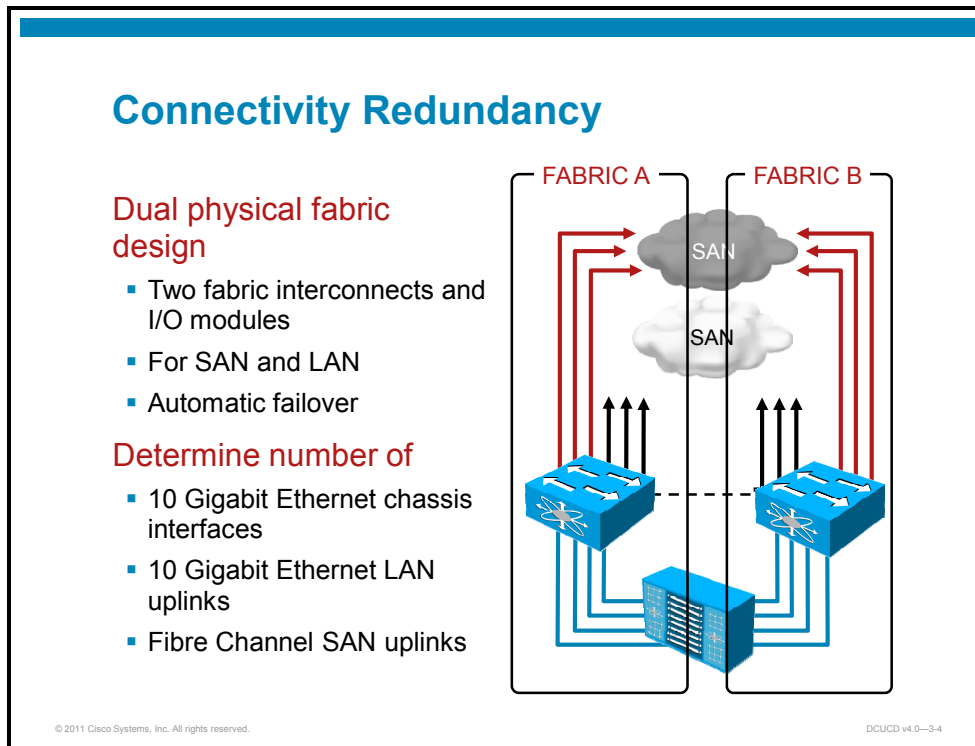
Objectives

Upon completing this lesson, you will be able to design the LAN and SAN of the Cisco UCS solution. This includes the ability to meet these objectives:

- Identify solution network and storage requirements
- Propose LAN hardware per requirements
- Propose SAN hardware per requirements

Analyzing Requirements

This topic identifies and describes solution network and storage requirements.



The Cisco UCS solution is typically designed with redundancy in mind. This means that a Cisco UCS cluster is connected to two physical fabrics—Fabric A and Fabric B. From the Cisco UCS LAN and SAN perspective, the Cisco UCS connectivity requirements define the number and type of 10 Gigabit Ethernet LAN uplinks (or 1 Gigabit Ethernet uplinks when used) and Fibre Channel SAN uplinks.

Depending on the network adapters and configuration, Cisco UCS can also incorporate automatic failover for individual adapters from the primary to secondary fabric.

Cisco UCS LAN Design Process

- Defines the number and type of physical LAN devices

1. Identify Cisco UCS network connectivity requirements.



2. Identify and size Cisco UCS network devices.



3. Define a Cisco UCS network connectivity scheme.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-5

The design of the Cisco Data Center Unified Computing System network can be divided into three steps:

- Reviewing the Cisco UCS sizing output to gather the network requirements
- Identifying and sizing the network device(s)
- Defining how Cisco UCS clusters(s) will be connected to the network

Designing Cisco UCS LAN

Identify network device(s) for LAN connectivity based on Cisco UCS sizing and design

- Type and number of **10 Gigabit Ethernet interfaces**
- Device(s) and connectivity **redundancy** levels
- Throughput

Define connectivity scheme

- Physical connectivity
- **VLANs**
- High-availability mechanisms (HSRP, vPC, and so on)



© 2011 Cisco Systems, Inc. All rights reserved.

The first step in the design is to review the Cisco UCS Ethernet uplink connectivity requirements. From these requirements, you can gather the input information to be able to size the network.

The following questions are important for the design process:

- What type of network interfaces are needed? Since you are connecting to Cisco UCS, the interfaces will be 10 Gigabit Ethernet. However, apart from the speed, you also need to know the connector and media type—for example, Twinax versus USR SFP+.
- What is the number of interfaces required? This will be the number of Cisco UCS cluster Ethernet uplink ports.
- What kind of throughput is needed on these interfaces? Is any oversubscription allowed, or do you need to ensure that no oversubscription is present on the Cisco UCS-facing interfaces?
- Do you need to leave a place for future expansion? Typically, some room should be left to be able to scale the solution in size in the future.

Once you select the network devices, you need to plan how the Cisco UCS clusters will be connected. This includes physical port connectivity—in other words, to which module on network device certain Cisco UCS interfaces will be connected. Second, can you scale the bandwidth by using PortChannel on the Cisco UCS Fabric Interconnect side? What about the Spanning Tree Protocol (STP) operation? Do you leave the Cisco UCS to operate in an EHV mode?

You also need to define the VLAN topology—this requires precise analysis of the services and applications to be deployed and the connectivity requirements.

Finally, yet importantly, high availability also needs to be addressed. This involves network device redundancy, default gateway redundancy protocol selection, and PortChannel deployment, among other high-availability features.

Cisco UCS SAN Design Process

- Defines the number and type of physical SAN devices

1. Identify Cisco UCS SAN connectivity requirements.



2. Identify and size Cisco UCS SAN devices.



3. Define a Cisco UCS SAN connectivity scheme.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-7

The design of the Cisco Unified Computing System SAN can be divided into three steps:

- Reviewing the Cisco UCS sizing output to gather the SAN requirements
- Identifying and sizing SAN device(s)
- Defining how Cisco UCS cluster(s) will be connected to the SAN

Designing Cisco UCS SAN

Identify MDS switch required for SAN connectivity based on Cisco UCS sizing and design

- Type and number of **Fibre Channel interfaces**
- Device and connectivity **redundancy** levels
- Required throughput

Define connectivity scheme

- Physical connectivity
- **VSANs**
- Define high-availability mechanisms
- Host multipathing software



© 2011 Cisco Systems, Inc. All rights reserved.

The first step in design is to review the Cisco UCS Fibre Channel uplink connectivity requirements. From these requirements, you can gather the input information to be able to size the SAN.

The following questions should be asked:

- What type of Fibre Channel interfaces do you need? Since you are connecting Cisco UCS, the interfaces will be 1/2/4G Fibre Channel. However, apart from the speed you also need to know the connector and media type—for example, small form-factor pluggable (SFP) multimode (MM).
- What is the number of interfaces required? This is the number of Cisco UCS cluster Fibre Channel uplink ports.
- What kind of throughput do you need on these interfaces? Is it allowable to have some oversubscription, or do you need to ensure that no oversubscription is present on the Cisco UCS-facing interfaces?
- Do you need to leave room for future expansion? Typically, some room should be left to be able to scale the solution size in the future.

Once you select the network devices, you need to plan how the Cisco UCS clusters will be connected. This includes the physical port connectivity—that is, to which module on the SAN device will the Cisco UCS Fibre Channel interfaces be connected?

You also need to define the virtual storage area network (VSAN) topology—this requires precise analysis of the server and SAN connectivity requirements. Depending on the SAN setup, you may need to define SAN pin groups to influence path selection for storage traffic exiting your Cisco UCS cluster(s).

Finally, yet importantly, high availability also needs to be addressed. If the VSAN topology allows, you can provide a first level of redundancy with multiple Fibre Channel uplinks, which are members of the same VSAN on a single fabric interconnect. However, typically host multipathing software will be needed to properly address the redundancy requirements.

Designing the Cisco UCS LAN

This topic identifies how to propose LAN hardware per requirements.

Greenfield Implementation Summary

- 2x Cisco UCS 6140 fabric interconnects:
 - 2x BC-class1 => 6 PS-class1 blades
 - 1x BC-class1 => 6 PS-class1 blades + PS-class2 blade
 - 2x BC-class1 => 8 PS-class2 blades
 - 3x BC-class2 => 5 VS-class1 blades
- Redundancy required:
- Installed in main data center in San Jose

	PS-class1 Server Blade	PS-class2 Server Blade	VS-class1 Server Blade
Adapter	UCS M71KR-Q CNA	UCS M71KR-Q CNA	UCS M81KR VIC
Processor	Intel Xeon E5540	Intel Xeon E5520	2x Intel Xeon E5570
Memory	48 GB	12 GB	96 GB
Quantity	18	17	15
Connectivity	LAN, SAN	LAN, SAN (boot)	2xLAN, SAN


© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-10

The design for the Greenfield implementation has already established the Cisco UCS requirements of the solution. Based on that information, network equipment can be selected. First, the summary of the Cisco UCS design for the solution should be observed—in other words, the part of the design that influences the network equipment selection. The information above summarizes the use for the implementation of two Cisco UCS 6140 Fabric Interconnects.

Greenfield Implementation Summary (Cont.)

- Single Cisco UCS cluster
- Connectivity requirements:
 - LAN connectivity per fabric interconnect—8x 10 Gigabit Ethernet ultra-short reach SFP+
 - 2:1 uplink oversubscription is allowed.
 - Upstream switch failure should not result in no connectivity.



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-11

The next thing you need to establish is how many Ethernet uplink ports you need for the fabric interconnect connectivity. For this implementation, you have established in the Cisco UCS part of the design that from each Cisco UCS 6100 in a cluster you need to connect eight 10 Gigabit Ethernet ports. Thus, altogether you need 16 10 Gigabit Ethernet ports to connect the Cisco UCS cluster.

Selecting Network Equipment

- Number of required 10 Gigabit Ethernet interfaces
 - 2x 8 x SFP+ ultra-short reach (USR) for uplinks from both Cisco UCS 6140XP
- **Two Catalyst 6509 in VSS setup**
- Individual Catalyst 6509 hardware configuration
 - Supervisor 720-10G for VSS mode of operation
 - 8-port 10 Gigabit Ethernet line card
 - Two 16-port 10 Gigabit Ethernet line cards for connectivity to the rest of the network
 - Two 6748 10/100/1000 GE line cards—other devices



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-12

Now you can select the network equipment. Since you need 16 10 Gigabit Ethernet interfaces for the connectivity, you can select the Catalyst 6500 as the solution.

In the data centers, high availability is very important. You can achieve better redundancy by deploying two Catalyst 6500 switches for the Cisco UCS cluster connectivity. In case of complete failure by one of them, the second should take over.

When Ethernet uplinks from the Cisco UCS are connected, you can scale the bandwidth by deploying the EtherChannel. In the Cisco UCS, this can currently be done from a single fabric interconnect—that is, you cannot put ports from the Fabric A and Fabric B switches into the channel. However, you can terminate the uplink ports at a different Catalyst 6500 chassis by deploying the Virtual Switching System (VSS). For that functionality, you need Supervisor 720-10G, which is the one that supports deployment of VSS.

To connect the 16 10 Gigabit Ethernet uplink ports from the Cisco UCS, you will use the 8-port 10 Gigabit Ethernet line card, one in each chassis.

The Catalyst 6500 switches also need to connect to the rest of the network. For that purpose you should additionally install two 16-port 10 Gigabit Ethernet line cards and two 48-port 10/100/1000 line cards.

Selecting Network Equipment (Cont.)

- WS-X6708-10GE-3C/CXL:
 - Cisco UCS connectivity
 - 2:1 oversubscription
- WS-X6716-10GE-3C/CXL:
 - VSL for VSS—use eight 10 Gigabit Ethernet in PortChannel
 - Rest of the network
 - 4:1 oversubscription
- WS-X6748-GE-TX (CEF720):
 - Other devices
 - 1:1.2 oversubscription




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-13

You have selected the following line cards for the Catalyst 6500. The line cards have appropriate oversubscription rates for your deployment.

Connecting Cisco UCS to the Network

- Physical connectivity
 - Connect Cisco UCS uplinks to Catalyst 6500 8-port 10 Gigabit Ethernet line card (smaller oversubscription rate)
 - Cisco UCS 6100-A Fabric Interconnect connectivity:
 - EtherChannel over 8 10 Gigabit Ethernet uplink ports on Cisco UCS
 - Connect first 4 10 Gigabit Ethernet uplinks to the first Catalyst 6500
 - Connect second 4 10 Gigabit Ethernet uplinks to the second Catalyst 6500
 - Deploy MEC on the Catalyst 6500 VSS side
 - Repeat the same connectivity for the Cisco UCS 6100-B Fabric Interconnect.



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-14

Once you know the quantity and type of the line cards and switches, you need to decide how you are going to connect the Cisco UCS.

You know you must connect eight 10 Gigabit Ethernet ports from each fabric interconnect. You have decided to deploy Catalyst 6500 in VSS mode to be able to create the multichassis EtherChannel.

Therefore, for this Greenfield implementation, you can use the following connectivity scheme:

- Split the eight 10 Gigabit Ethernet ports from a single fabric interconnect between the two Catalyst 6500 chassis—i.e., four 10 Gigabit Ethernet ports to the first Catalyst 6500 chassis and four 10 Gigabit Ethernet ports to the second Catalyst 6500 chassis.
- Create the EtherChannel on the fabric interconnect and add all eight ports to the channel.
- Create a multichassis EtherChannel (MEC) on the Catalyst 6500 VSS and add all eight ports from a single fabric interconnect to the channel.
- Repeat the same setup for the second fabric interconnect.

Existing Solution Migration Summary

- Installed in main data center in Berlin
- Redundancy required

Cisco UCS System Quantity and per UCS Blade Chassis Population

UCS Class	Quantity	Max. Blade Quantity	BC-class1	BC-class2	BC-class3	BC-class4
UC-class1	6	96	12	-	-	-
UC-class2	2	64	-	6	2	-
UC-class3	1	160	10	1	-	9

Per Chassis Maximum Blade Population

Chassis Class	PS-class1	PS-class2	PS-class3	PS-class4	VS-class1
BC-class1	6	2	-	-	-
BC-class2	-	8	-	-	-
BC-class3	-	-	-	-	8
BC-class4	-	-	4	4	-

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-15

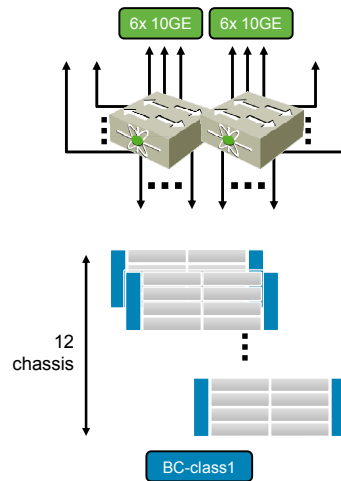
For the design of the existing solution migration, you have established the Cisco UCS requirements by transforming the existing servers into the Cisco UCS. The design process has established that three different Cisco UCS classes will be used, with their own characteristics. The information above summarizes the information about the Cisco UCS classes designed.

You know from this information that you need to connect the following:

- Six UCS-class1 systems
- Two UCS-class2 systems
- One UCS-class3 system

Existing Solution Migration Summary (Cont.)

- **UCS-class1:**
 - 12x BC-class1
 - 2x Cisco UCS 6140XP
- **Per fabric interconnect**
 - LAN uplinks – 6x 10GE USR SFP+



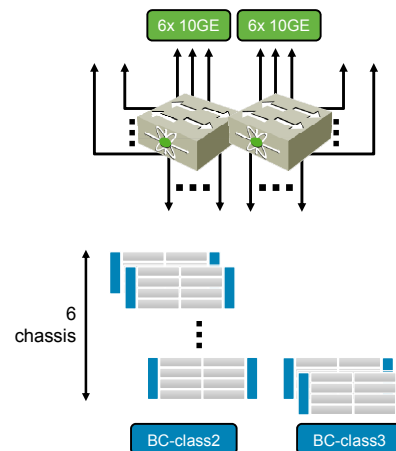
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-16

Next, you need to establish the quantity of uplink parts per individual Cisco UCS class. First, look at the UCS-class1 system. From the Cisco UCS design, you know you need to connect six 10 Gigabit Ethernet uplink ports from each fabric interconnect.

Existing Solution Migration Summary (Cont.)

- **UCS-class2:**
 - 6x BC-class2
 - 2x BC-class3
 - 2x Cisco UCS 6140XP
- **Per fabric interconnect:**
 - LAN uplinks—7x 10 Gigabit Ethernet USR SPF+



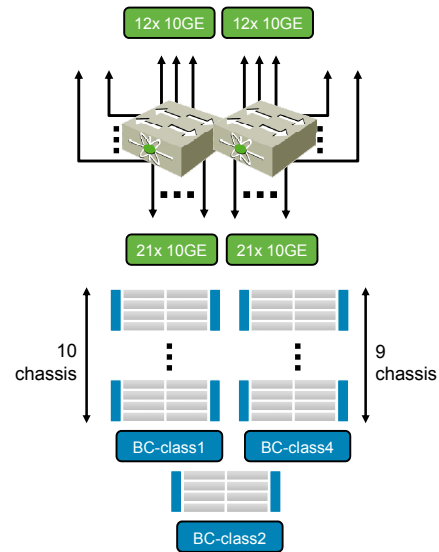
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-17

Next, you determine the requirements from the UCS-class2 system. Here you need to connect seven 10 Gigabit Ethernet uplink ports from each fabric interconnect.

Existing Solution Migration Summary (Cont.)

- **UCS-class3:**
 - 10x BC-class1
 - 1x BC-class2
 - 9x BC-class4
 - 2x Cisco UCS 6140XP
- **Per fabric interconnect:**
 - LAN uplinks—12x 10 Gigabit Ethernet USR SFP+



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-18

For the third Cisco UCS class type—UCS-class3—you need to connect 12 10 Gigabit Ethernet uplink ports from a single fabric interconnect.

Assessing Network Requirements

- Physical Ethernet fabrics A and B
- Per Ethernet fabric => 62 10 Gigabit Ethernet non-oversubscribed interfaces:
 - $6 * 6 = 36$ 10 Gigabit Ethernet non-oversubscribed (6 * System1 => 6*6 10 Gigabit Ethernet)
 - $2 * 7 = 14$ 10 Gigabit Ethernet non-oversubscribed (2 * System2 => 2*7 10 Gigabit Ethernet)
 - $12 = 12$ 10 Gigabit Ethernet non-oversubscribed (1 * System3 => 12 10 Gigabit Ethernet)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-19

Knowing the port count information per Cisco UCS class, you can calculate the total port count for fabric A and fabric B:

- You need to deploy 6 UCS-class1 systems requiring 6 10 Gigabit Ethernet uplink ports per fabric interconnect, thus you need 36 10 Gigabit Ethernet interfaces from the network side.
- Next, you need to deploy 2 UCS-class2 systems requiring 7 10 Gigabit Ethernet uplink ports per fabric interconnect, thus you need an additional 14 10 Gigabit Ethernet interfaces from the network side.
- For the UCS-class1 system—since only one needs to be deployed—you need to add 12 10 Gigabit Ethernet uplink ports per fabric interconnect to the total port count.

Altogether, you need 62 10 Gigabit Ethernet interfaces for fabric A connectivity and 62 10 Gigabit Ethernet interfaces for fabric B connectivity.

Selecting Network Equipment

- Two Nexus 7018 chassis with vPC deployment
- Nexus 7000 hardware configuration:
 - 8 * 32-port 10G line cards:
 - 8 ports per line card used in dedicated mode for Cisco UCS 6140 connectivity
 - 3 * fabric modules for fabric redundancy
 - 1 * 48-port line card to connect to other network devices
 - Additional 32-port 10G module(s) used for connectivity to the rest of the network and between the Nexus 7000



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-20

The number of 10 Gigabit Ethernet interfaces required indicates that you need to use the Nexus 7000 switches. To observe the importance of high availability in the solution, you will deploy two Nexus 7010 switches for the Cisco UCS connectivity. In the case of complete failure of one of the switches, the second should take over.

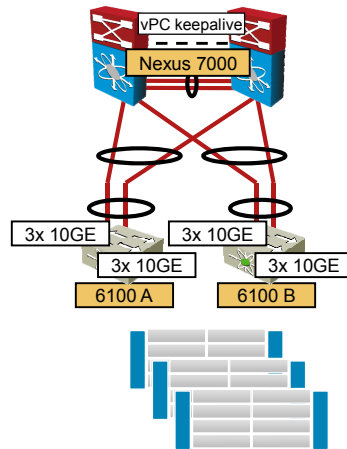
Nexus 7000 supports a virtual PortChannel (vPC) functionality, which allows the interfaces in a PortChannel to terminate at different chassis. To scale the bandwidth, you decide to use it.

Taking into account the interface requirements, this is the Nexus 7010 hardware list:

- Eight 32-port 10 Gigabit Ethernet line cards. You will use eight ports per line card in the dedicated mode of operation to get the full 10 Gigabit Ethernet throughput on each of the Cisco UCS uplink ports.
- For the fabric redundancy, you need to install three fabric modules into a single chassis.
- Additionally, you will add a 48-port 10/100/100 line card and a 32-port 10 Gigabit Ethernet line card for the connectivity to the other devices and the rest of the network.

Connecting Cisco UCS to the Network

- Nexus 7000 with vPC to scale the bandwidth
- UCS-class1 physical connectivity
 - Cisco UCS 6100 A Fabric Interconnect
 - Connect 3 10 Gigabit Ethernet interfaces to the first Nexus 7000
 - Connect 3 10 Gigabit Ethernet interfaces to the second Nexus 7000
 - Create a PortChannel on the Cisco UCS side
 - Create a vPC on the Nexus side
 - Repeat the same connectivity configuration for the Cisco UCS 6100 B Fabric Interconnect




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-21

Now you need to decide how the Cisco UCS systems will be connected to the Nexus 7000 switches. First, look at the UCS-class1 system. You will split the six uplink ports from a single fabric interconnect between the two chassis and put them into the channel on the Cisco UCS side. On the Nexus 7000, you will use the vPC and put the six interfaces into the channel. You will repeat the same connectivity configuration for the second fabric interconnect.

Connecting Cisco UCS to the Network (Cont.)

- UCS-class2 physical connectivity
 - Cisco UCS 6100 A Fabric Interconnect
 - Connect 4 10 Gigabit Ethernet interfaces to the first Nexus 7000
 - Connect 3 10 Gigabit Ethernet interfaces to the second Nexus 7000
 - Create a PortChannel on the Cisco UCS side
 - Create a vPC on the Nexus side
 - Cisco UCS 6100 B Fabric Interconnect
 - Connect 3 10 Gigabit Ethernet interfaces to the first Nexus 7000
 - Connect 4 10 Gigabit Ethernet interfaces to the second Nexus 7000
 - Create a PortChannel on the UCS side
 - Create a vPC on the Nexus side



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-22


For the UCS-class2 system, you will split the seven uplink ports from a single fabric interconnect between the two chassis in the following manner:

- For fabric interconnect A, four ports will terminate at the first Nexus 7000 chassis and three ports will terminate at the second Nexus 7000 chassis.
- For fabric interconnect B, three ports will terminate at the first Nexus 7000 chassis and four ports will terminate at the second Nexus 7000 chassis.

For both fabric interconnects, you will create a channel on the Cisco UCS side and a vPC on the Nexus 7000.

Connecting Cisco UCS to the Network (Cont.)

- UCS-class3 physical connectivity
- Cisco UCS 6100 A Fabric Interconnect
 - Connect 6 10 Gigabit Ethernet interfaces to the first Nexus 7000
 - Connect 6 10 Gigabit Ethernet interfaces to the second Nexus 7000
 - Create two PortChannels on the Cisco UCS side:
 - Put 6 interfaces into a channel
 - Create a regular PortChannel on the Nexus side
- Repeat the same connectivity configuration for the second Cisco UCS 6100



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-23

For the UCS-class3 system, you will split the 12 uplink ports from a single fabric interconnect into two 6-port groups. Then each 6-port group will be split between the two Nexus 7000 chassis (three ports per chassis).

You will use two channels on the Cisco UCS side with six uplink ports in a channel and two vPCs on the Nexus 7000 side. The same connectivity configuration will be repeated for the second fabric interconnect.

Designing the Cisco UCS SAN

This topic identifies how to propose SAN hardware per requirements.

Greenfield Implementation Summary

- 2x Cisco UCS 6140 Fabric Interconnect:
 - 2x BC-class1 => 6 PS-class1 blades
 - 1x BC-class1 => 6 PS-class1 blades + PS-class2 blade
 - 2x BC-class1 => 8 PS-class2 blades
 - 3x BC-class2 => 5 VS-class1 blades
- Redundancy required
- Installed in main data center in San Jose

	PS-class1 Server Blade	PS-class2 Server Blade	VS-class1 Server Blade
Adapter	UCS M71KR-Q CNA	UCS M71KR-Q CNA	UCS M81KR VIC
Processor	Intel Xeon E5540	Intel Xeon E5520	2x Intel Xeon E5570
Memory	48 GB	12 GB	96 GB
Quantity	18	17	15
Connectivity	LAN, SAN	LAN, SAN (boot)	2xLAN, SAN


© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-25

As with network design, for the SAN design you also need to examine what you have established during the Cisco UCS design phase—that is, the Cisco UCS requirements of the solution. Based on that information, SAN equipment can be selected. The information above summarizes that for implementation you will use two Cisco UCS 6140 Fabric Interconnects.

Greenfield Implementation Summary (Cont.)

- SAN connectivity per fabric interconnect:
 - 2x N10-E0080 expansion module
 - 16x 1/2/4G Fibre Channel MM SFP



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-26

Next, you need to establish how many Fibre Channel uplink ports are needed for the fabric interconnect connectivity. For this implementation, you have established in the Cisco UCS part of the design that from each Cisco UCS 6100 in a cluster you need to connect 16 4G Fibre Channel ports. Thus, altogether you need 32 4G Fibre Channel ports to connect the Cisco UCS cluster.

Selecting SAN Equipment

- Number of required Fibre Channel interfaces:
 - 2x 16 4G Fibre Channel interfaces
- Cisco MDS 9506 hardware configuration:
 - Two supervisors
 - 3x 24-port 1/2/4/8G Fibre Channel line card
 - All ports on line cards capable of wire rate 4G Fibre Channel



© 2011 Cisco Systems, Inc. All rights reserved.


DCUCD v4.0-3-27

Now you are ready to select the SAN equipment. To connect 32 4G interfaces in a redundant fashion, you will use the Cisco MDS 9506 director class switches in the following setup:

- Two supervisors for the purpose of redundancy
- Three 24-port 1/2/4/8G Fibre Channel line cards—all ports on such line cards can operate at wire speed at 4G Fibre Channel, which is required for the Cisco UCS uplink ports.

Connecting Cisco UCS to SAN

- Cisco UCS 6100-A connectivity:
 - First 4 FC ports => MDS line card 1
 - Second 4 FC ports => MDS line card 2
 - Third 4 FC ports => MDS line card 3
 - Fourth 4 FC ports => MDS line card 1
- Cisco UCS 6100-B connectivity:
 - First 4 FC ports => MDS line card 2
 - Second 4 FC ports => MDS line card 3
 - Third 4 FC ports => MDS line card 1
 - Fourth 4 FC ports => MDS line card 2
- Remaining Fibre Channel ports for drive array connectivity:
 - Use round robin approach
 - 4-port Cisco MDS line card 1, 2, 3



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-28

To achieve better redundancy from a connectivity perspective, you will spread the Cisco UCS Fibre Channel uplink ports between the line cards as indicated in the figure.

Existing Deployment Migration Summary

- Installed in main data center in Berlin
- Redundancy required

Cisco UCS System Quantity and per UCS Blade Chassis Population

UCS class	Quantity	Max. Blade Quantity	BC-class1	BC-class2	BC-class3	BC-class4
UC-class1	6	96	12	-	-	-
UC-class2	2	64	-	6	2	-
UC-class3	1	160	10	1	-	9

Per Chassis Maximum Blade Population

Chassis Class	PS-class1	PS-class2	PS-class3	PS-class4	VS-class1
BC-class1	6	2	-	-	-
BC-class2	-	8	-	-	-
BC-class3	-	-	-	-	8
BC-class4	-	-	4	4	-

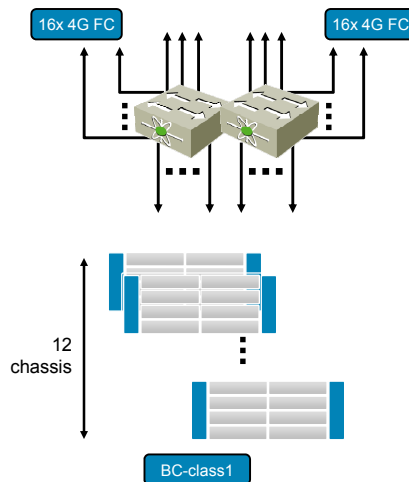
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-29

Just as with the new implementation use case example, for the existing solution migration you need to first establish the port count requirements. Based on that information, SAN equipment can be selected. You know that you have three different Cisco UCS system classes and you know the quantity of each as stated in the tables in the figure.

Sizing System—Fabric Interconnect Classes

- **UCS-class1:**
 - 12x BC-class1
 - 2x Cisco UCS 6140XP IOMs
- Per fabric interconnect:
 - 2x N10-E0080 expansion module
 - SAN uplinks—16x 1/2/4G Fibre Channel MM SFP



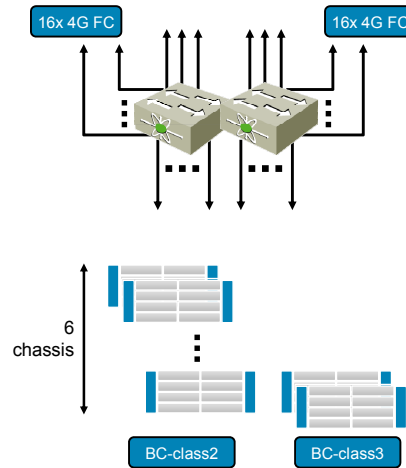
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-30

For UCS-class1 systems, you need to connect 16 1/2/4G Fibre Channel uplink ports.

Sizing System—Fabric Interconnect Classes (Cont.)

- **UCS-class2:**
 - 6x BC-class2
 - 2x BC-class3
 - 2x Cisco UCS 6140XP IOMs
- Per fabric interconnect:
 - 2x N10-E0080 expansion module
 - SAN uplinks—16x 1/2/4G Fibre Channel MM SFP



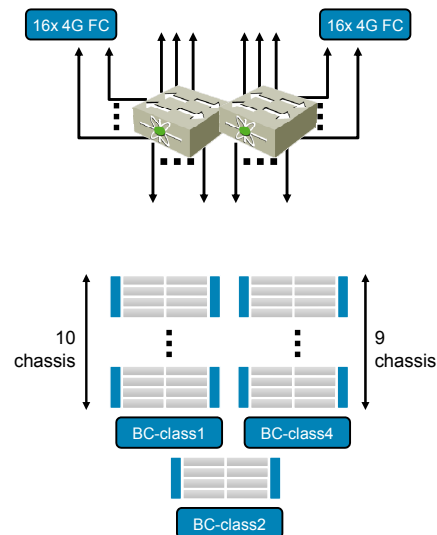
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-31

For UCS-class2 systems, you also need to connect 16 1/2/4G Fibre Channel uplink ports.

Sizing System—Fabric Interconnect Classes (Cont.)

- **UCS-class3:**
 - 10x BC-class1
 - 1x BC-class2
 - 9x BC-class4
 - 2x Cisco UCS 6140XP IOMs
- Per fabric interconnect
 - 2x N10-E0080 expansion module
 - SAN uplinks—16x 1/2/4G Fibre Channel MM SFP



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—3-32

The same applies to UCS-class3 systems—you need to connect 16 1/2/4G Fibre Channel uplink ports.

SAN Requirements

- Physical Fibre Channel fabric A, B
- Per Fibre Channel fabric => 143 4G Fibre Channel non-oversubscribed interfaces
- UCS-class1: 6*16 4G Fibre Channel => 96 4G Fibre Channel
- UCS-class2: 2*16 4G Fibre Channel => 32 4G Fibre Channel
- UCS-class3: 16* 4G Fibre Channel => 16 4G Fibre Channel




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-33

The requirements for the SAN fabric for the Cisco UCS connectivity state that you need two physical fabrics—A and B. You need a total of 143 4G Fibre Channel interfaces per fabric.

Selecting SAN Equipment

- Use two existing Cisco MDS 9509s:
 - Use 3 * 48 1/2/4/8G Fibre Channel modules per Cisco MDS for 6140 connectivity:
 - Connect ports interchangeably.
 - Half of the ports represent fabric A and half represent fabric B.
 - Use two existing 1/2/4G 48 Fibre Channel port modules for 36 disk array ports.
 - Use two existing 1/2/4G 12 Fibre Channel port modules for Inter-Switch Link (ISL) between Cisco MDS switches.



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-34

Since you are migrating the solution, the customer already should have some existing equipment that can be used. Thus you will use two Cisco MDS9509 switches for the Cisco UCS deployment. You do need to add additional line cards since you need 143 4G wire rate Fibre Channel interfaces. For this purpose, you will use three 48 1/2/4/8G Fibre Channel line cards in each Cisco MDS switch.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- To size the network for the Cisco Data Center Unified Computing System solution, the Ethernet uplink ports from the Cisco fabric interconnects must be counted.
- Deploy VSS with MEC when connecting the Cisco UCS to the Catalyst 6500 to scale the bandwidth, if appropriate.
- Deploy a vPC when connecting the Cisco UCS to the Nexus 7000 to scale the bandwidth, if appropriate.
- Cisco UCS Fabric Interconnect Fibre Channel ports govern the number of Fibre Channel ports on the SAN switch.
- Interfaces from a single Cisco UCS fabric interconnect should be connected to two different uplink SAN switches whenever possible.

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- Cisco UCS uses 10 Gigabit Ethernet with FCoE support to consolidate Ethernet and Fibre Channel traffic on a common media.
- The Cisco UCS B-Series is a system that consists of fabric interconnect switches, chassis with IOMs, and B-Series Server blades and that is managed from the Cisco UCS Manager application.
- Cisco UCS fabric interconnect switches are typically deployed in a cluster to provide connectivity and management high availability.
- EHV is the default switching mode for Cisco UCS fabric interconnect, which eliminates the need for STP.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-1

Module Summary (Cont.)

- The Cisco UCS B-Series sizing process defines server blade classes, chassis classes, and, if necessary, cluster classes.
- Cisco UCS LAN addresses Layer 1, Layer 2, and Layer 3 aspects.
- Cisco UCS LAN design consists of three steps—identifying Cisco UCS connectivity requirements, selecting network devices, and defining the Cisco UCS network connectivity configuration.
- Cisco UCS SAN design selects the appropriate SAN devices among Cisco MDS Series switches.
- NPIV functionality is required on the core SAN switch to connect the Cisco UCS.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-3-2

Module Summary (Cont.)

- The Cisco Nexus 1000V distributed virtual switch enables mobility for network configuration.
- VSM is the Cisco Nexus 1000V supervisor module that is integrated with VMware vCenter.
- Cisco Nexus 1000V deployment requires management, control, and packet VLANs for VSM-to-VEM communication.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which C-Series server adapter supports FCoE? (Source: Evaluating Cisco UCS C-Series Architecture)
- A) Broadcom Ethernet adapters
 - B) Emulex 4-port HBA
 - C) QLogic 4-port HBA
 - D) Cisco P81E VIC
 - E) Intel Ethernet adapters
- Q2) Which C-Series server can be equipped with the Intel Xeon 5600 and memory of 256 GB? (Source: Evaluating Cisco UCS C-Series Architecture)
- A) C200 M2
 - B) C210 M2
 - C) C460 M2
 - D) C250 M1
 - E) C250 M2
- Q3) What is required to integrate the C-Series server with Cisco UCS 6120XP? (Source: Evaluating Cisco UCS C-Series Architecture)
- A) special management adapter
 - B) Cisco Nexus 2248
 - C) preconfigured service profiles
 - D) two Fibre Channel HBAs
 - E) C460 M2 server
- Q4) Which two tools can be used to create the C-Series BOM? (Choose two.) (Source: Sizing the Cisco UCS C-Series Solution)
- A) Capacity Planner
 - B) Dynamic Configuration Tool
 - C) MAP toolkit
 - D) NetformX DesignXpert
 - E) Cisco UCS Manager
- Q5) How many server ports are available on the Cisco UCS 6120XP Fabric Interconnect switch? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) 20
 - B) 26
 - C) 24
 - D) 10
 - E) 0

- Q6) What is the maximum amount of memory per Cisco UCS B200-M1 Server Blade?
(Source: Evaluating Cisco UCS B-Series Architecture)
- A) 384 GB
 - B) 96 GB
 - C) depends on DIMM size
 - D) 64 GB
 - E) 32 GB
 - F) 128 GB
- Q7) Which protocol is native to Cisco UCS? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) SNMP
 - B) CIM-XML
 - C) IPMI
 - D) XML API
 - E) SoL
- Q8) Where does Cisco UCS Manager run? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) on dedicated server
 - B) on virtual machine
 - C) on Cisco UCS 6100XP Fabric Interconnect
 - D) on Cisco UCS B200-M1 Blade Server
- Q9) Which interfaces are used for Cisco UCS cluster connectivity? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) first two ports
 - B) last two ports
 - C) first two ports on expansion module
 - D) ports L1 and L2
- Q10) Which fabric port on the left IOM would be used by the server blade in slot 3 if all four ports on IOM are connected to a fabric interconnect? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) port 1 on IOM
 - B) depends on configuration
 - C) port 1 or port 2
 - D) port 3 on IOM
- Q11) Which interface in EHV mode is used to connect Cisco UCS to the external LAN?
(Source: Evaluating Cisco UCS B-Series Architecture)
- A) bridge link
 - B) border link
 - C) server link
 - D) active link
- Q12) Which configuration option allows administrative control over uplink selection?
(Source: Evaluating Cisco UCS B-Series Architecture)
- A) vNIC redundancy
 - B) LAN pin group
 - C) traffic shaping
 - D) cost

- Q13) What is the Fibre Channel uplink port type on the Cisco UCS Fabric Interconnect? (Source: Evaluating Cisco UCS B-Series Architecture)
- A) F_Port
 - B) NPV_Port
 - C) NP_Port
 - D) TE_Port
- Q14) What do you need to determine when designing physical deployment for Cisco UCS? (Source: Sizing the Cisco UCS B-Series Solution)
- A) number of server uplink ports
 - B) available power per rack cabinet
 - C) server blade form factor
 - D) server adapter type
- Q15) Which tool can be used to create the C-Series BoM? (Source: Sizing the Cisco UCS B-Series Solution)
- A) Capacity Planner
 - B) Dynamic Configuration tool
 - C) MAP toolkit
 - D) NetformX DesignXpert
 - E) Cisco UCS Manager
- Q16) Data center design also involves a physical deployment plan. Which tool can be used to help with planning for physical deployment? (Source: Planning Physical Deployment)
- A) Cisco UCS Manager
 - B) NetformX DesignXpert
 - C) Microsoft Vision
 - D) Cisco UCS Power Calculator tool
 - E) Cisco UCS ROI tool
- Q17) Which two mechanisms can be used to scale high availability and bandwidth when connecting the Cisco UCS to two Cisco Catalyst 6500 Series Switches? (Choose two.) (Source: Examining the Cisco UCS Network and Storage)
- A) VSS
 - B) VDC
 - C) vPC
 - D) FCoE
 - E) MEC
- Q18) Which protocol can be used to address default gateway high availability? (Source: Designing the Cisco UCS Network and Storage)
- A) HSRP
 - B) BFD
 - C) ISSU
 - D) EHV
- Q19) Which is an important aspect of designing a Cisco UCS network? (Source: Designing the Cisco UCS Network and Storage)
- A) routing protocol selection
 - B) physical connectivity configuration
 - C) context placement
 - D) FCoE VLAN ID

- Q20) Which mechanism must be enabled on the core SAN switch in order to connect Cisco UCS? (Source: Designing the Cisco UCS Network and Storage)
- A) interop mode 1
 - B) NPIV
 - C) NPV edge mode
 - D) FSPF
- Q21) Which solution is used to achieve storage high availability from the server perspective? (Source: Designing the Cisco UCS Network and Storage)
- A) VRRP
 - B) host multipathing software
 - C) PortChannel
 - D) HBA automatic failover
- Q22) How many VEMs can be managed from a single VSM? (Source: Designing the Cisco UCS Network and Storage)
- A) 256
 - B) 128
 - C) 64
 - D) 16
- Q23) What happens upon VSM failure? (Source: Designing the Cisco UCS Network and Storage)
- A) All VEMs seize to forward the traffic.
 - B) All VEMs revert to standard vSwitch operational mode.
 - C) VM traffic continues to be switched by VEMs.
 - D) The ESX host reverts to isolated mode.
- Q24) Which is the proper way of ensuring VSM high availability? (Source: Designing the Cisco UCS Network and Storage)
- A) Enable VMware HA for VSM VM.
 - B) Enable VMware FT for VSM VM.
 - C) Run VSM on the physical server.
 - D) Deploy a secondary VSM on a different ESX host.
- Q25) Which practice ensures proper operation of multiple Cisco Nexus 1000V domains on the same ESX servers that are sharing control and packet VLANs? (Source: Designing the Cisco UCS Network and Storage)
- A) Use different domain IDs for individual domains.
 - B) Configure port groups with different names.
 - C) Place VSMs on different management VLANs.
 - D) Put ESX servers in the same cluster.

Module Self-Check Answer Key

Q1)	D
Q2)	E
Q3)	B
Q4)	B, D
Q5)	A
Q6)	B
Q7)	D
Q8)	C
Q9)	D
Q10)	D
Q11)	B
Q12)	B
Q13)	C
Q14)	B
Q15)	D
Q16)	D
Q17)	A, E
Q18)	A
Q19)	B
Q20)	B
Q21)	B
Q22)	D
Q23)	C
Q24)	D
Q25)	A

Design Server Deployment

Overview

This module evaluates how to propose a basic server deployment plan for Cisco UCS.

Objectives

Upon completing this module, you will be able to understand the Cisco UCS basic server deployment model. This ability includes being able to meet this objective:

- Apply the Cisco UCS server implementation model to deploy a new server

Designing the Cisco UCS Server Deployment Model

Overview

This lesson identifies and describes how to apply the Cisco UCS server implementation model to deploy a new server.

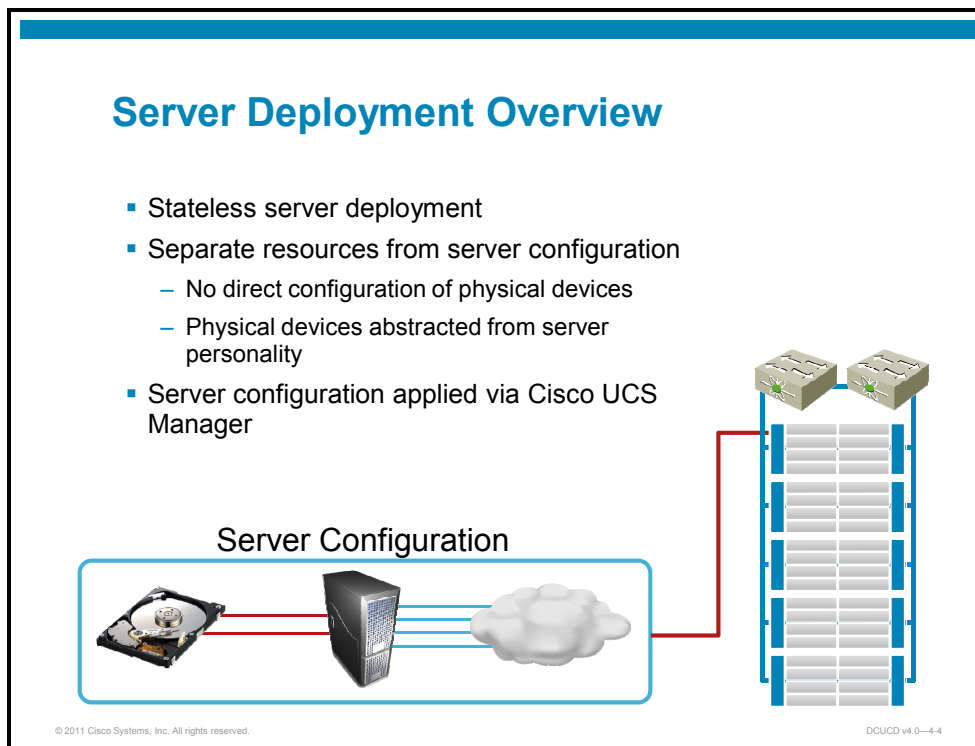
Objectives

Upon completing this lesson, you will be able to deploy a new server using the Cisco UCS server implementation model. This includes the ability to meet these objectives:

- Design Cisco UCS B-Series basic server deployment
- Evaluate the UCS advanced server deployment model
- Evaluate UCS identity pools
- Evaluate UCS resource pools
- Evaluate UCS policies
- Evaluate Cisco UCS release 1.4 enhancements

Designing Basic B-Series Server Deployment

This topic discusses the design of basic server deployment for the Cisco UCS B-series.



The Cisco UCS server deployment model separates the physical hardware of the server from the server configuration and personality by decoupling the logical domain from the physical resources. With this deployment model, no configuration is performed on the physical components. The configuration is applied implicitly upon the server personality that is assigned to a physical server.

With this type of separation, stateless computing can be deployed where the applied attributes are not tied to the physical hardware. The server or computing hardware is simply a generic computing capacity that is flexible and can be dynamically controlled. The flexibility of the computing capacity enables the dynamic provisioning of servers, where the entire server-related configuration is applied via the Cisco UCS Manager, with server elements also being controlled from that single management point.

Using Service Profiles

Every server is stateless and must be associated with a service profile in order to gain its identities and personality. A service profile is associated with a stateless server via manual association, or automatically via a blade pool.

After you associate a service profile with a server, the server is ready to have an operating system and applications installed, and you can use the service profile to review the configuration of the server. When a service profile is disassociated from a server, the identity and connectivity information for the server is reset to factory defaults, and the server is returned to the pool (if it originated from a pool).

Servers and service profiles have a one-to-one relationship. Each server can be associated with only one service profile. Each service profile can be associated with only one server. A service profile can be modified, cloned, or used to instantiate a template.

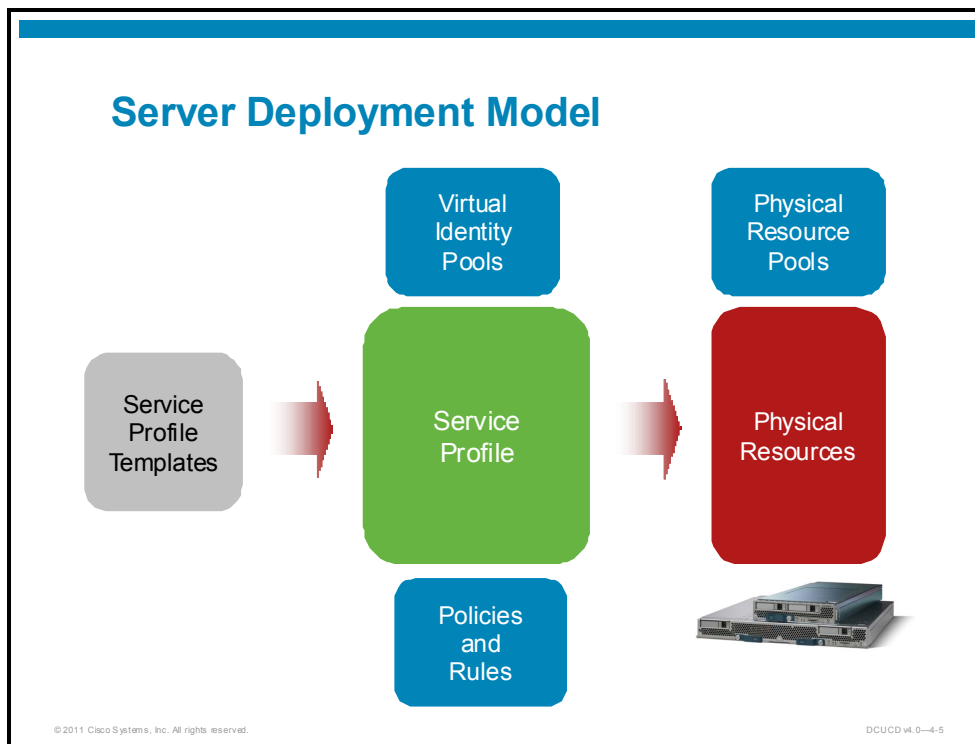
Note At any given time, each server can be associated with only one service profile. Similarly, each service profile can be associated with only one server at a time.

Note that within the Cisco Unified Computing System, the service profiles must be used for any server deployment, even if the blades will be managed as traditional individual rack servers. In such cases, the service profile provides LAN and SAN connectivity configuration.

A server that is associated with a service profile is similar to a traditional bare metal server in the way that it is ready to be used for production and run business services or application software.

Server Mobility

Server mobility enables the transfer of server identity seamlessly between the server blades and service profiles. When a server must be moved from one blade to another, a simple operation of associating a service profile to another available server in a Cisco UCS system automatically includes complete migration of identities, firmware, and connectivity to the LAN and the SAN.



The Cisco UCS server deployment model is built around the following elements, which will be described in detail later.

Server Personality

A server personality is composed of a service profile, which is a description of an individual server and all its related characteristics. A service profile can use various identity pools (for example, a MAC pool, a world wide name [WWN] pool, and so on) from which identity parameters are taken. Apart from the identity, the service profile is also assigned various policies that define rules, which govern how the server is initialized and managed.

Service profile templates offer an approach to server deployment that eases large-scale server deployments.

Server Physical Resources

Server physical resources (for example, server blades) can be used individually or can be pooled into common resource pools upon which service profiles are applied. A blade must be associated with a service profile in order to be used for the operating system and application deployment.

Cisco UCS Resources

Storage

Optionally specify disk policies and SAN configuration information.

Local Storage: If nothing is selected, the default Local Storage policy will be assigned to the service profile.

Create Local Disk Configuration Policy

Scrub Policy: Create Scrub Policy

How would you like to configure SAN connectivity? Simple Expert No WRE

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN profile.

World Wide Node Name

WWNN Assignment: 20:00:00:25:85:01:10:00

Create WWNN Pool

World Wide Node Name:

Click [here](#) to verify if this WWNN is available.

Name	WWNN
vHBA HBA-SANfabricA	20:00:00:25:85:01:10:01
vHBA LP VSAN11	

Parameter

UUID

LAN—vNICs

- MAC address
- Fabric selection
- VLANs

SAN—vHBAs

- nWWN address
- pWWN per HBA
- Fabric selection
- VSAN

Policies—boot, firmware, etc.

Server association

- Manual vs. from pool

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-4-6

As mentioned earlier, in order to provide a flexible server deployment model, Cisco UCS separates resources and identities (personality) from the physical server (resources). The abstraction provides a flexible environment in which resources can be provisioned and migrated between physical devices rapidly.

Cisco UCS categorizes resources separately from the configurations that may reside on them. This is to allow rapid, stateless provisioning of the computing cloud.

LAN Resources

The Cisco UCS management system that runs on the Cisco UCS fabric interconnect controls various LAN resources. These resources include MAC addresses, VLANs, and LAN ports.

SAN Resources

The Cisco UCS management system is responsible for VSANs, world wide node name (nWNN), and world wide port name (pWWN) address assignment, as well as SAN ports. The central assignment of WWNs is critical to the portability of service profiles.

Compute Resources

The Cisco UCS management system manages compute nodes—blade servers, firmware, and universally unique identifiers (UUIDs) within the Cisco UCS environment.

Management Resources

All system management is processed through the Cisco UCS Manager that runs on the fabric interconnect. The management VLANs, server keyboard, video, mouse (KVM) attachments, and low-level control are also Cisco UCS resources.

Basic Service Profile Model

Matches typical standalone server deployment:

- No mobility without manual reconfiguration
- Profile still required (no network or SAN connectivity without it)
- Distinct profile for each physical blade server you want to boot concurrently
- Clone profiles or create from templates to simplify deployment

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—4.7

Service Profile

A service profile is a self-contained logical representation (object) of a desired physical server, including connectivity, configuration, and identity.

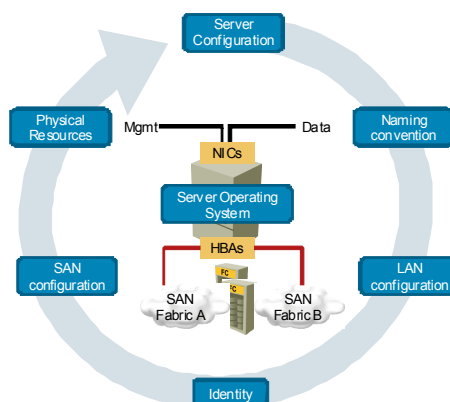
Each service profile serves a specific purpose—to ensure that the associated server hardware has the necessary configuration, identities, and connectivity to LAN and SAN based on the requirements from the applications the server will host.

The service profile maintains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity—the server personality. This includes all of the unique information for the server, including MAC, pWWN, UUID, boot order, and so on.

The personality information is stored in a format that can be managed through the Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnects.

Basic Profile Parameters

- UUID
 - Uses hardware defaults
 - Called “derived” addressing
- vNIC and vHBA
 - Must create vNIC or vHBA for every adapter you want connected in blade server operating system
- vNIC and vHBA identity (MAC or WWN)
 - Specify or use derived (not available with M81KR and P81E)
- Boot order
 - Set through profile
 - Leave empty to control via server BIOS



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-4-8

The service profile defines server hardware (configuration, firmware, identity, and boot information), fabric connectivity, policies, external management, and high availability information.

Cisco UCS uses pools, policies, and service profiles to abstract state and configuration information from all the components that define a server, its service quality, and its connectivity—this makes the Cisco UCS system a stateless computing system. Cisco UCS enables an operational state and services to be separated from the server hardware and physical connectivity.

Network Connectivity

Each service profile specifies the LAN and SAN network connections for the server through the Cisco UCS infrastructure and out to the external network. An administrator is not required to manually configure the network connections for Cisco UCS servers and other components. All network configuration is performed through the service profile. When a service profile is associated with a server, the Cisco UCS internal fabric is configured with the information in the service profile. If the profile was previously associated with a different server, the Cisco UCS network infrastructure is reconfigured to support identical network connectivity to the new server.

Hardware Component Configuration

When a service profile is associated with a server, the following components are configured according to the data in the profile:

- Server, including BIOS and baseboard management controller (BMC)
- Adapters
- Fabric interconnect

The administrator is not required to configure these hardware components directly.

Server Identity Management

The administrator can use the network and device identities that are burned into the server hardware at the time of manufacture, or use the identities that are specified in the associated service profile either directly or through identity pools, such as MAC, WWN, and UUID.

The following are examples of configuration information that the administrator can include in a service profile:

- Profile name and description
- Unique server identity (the UUID)
- LAN connectivity attributes, such as the MAC address
- SAN connectivity attributes, such as the WWN

Operational Aspects Configuration

The administrator can configure some of the operational functions for a server in a service profile, such as these:

- Firmware packages and versions
- Operating system boot order and configuration
- Intelligent Platform Management Interface (IPMI) and KVM access

vNIC Configuration

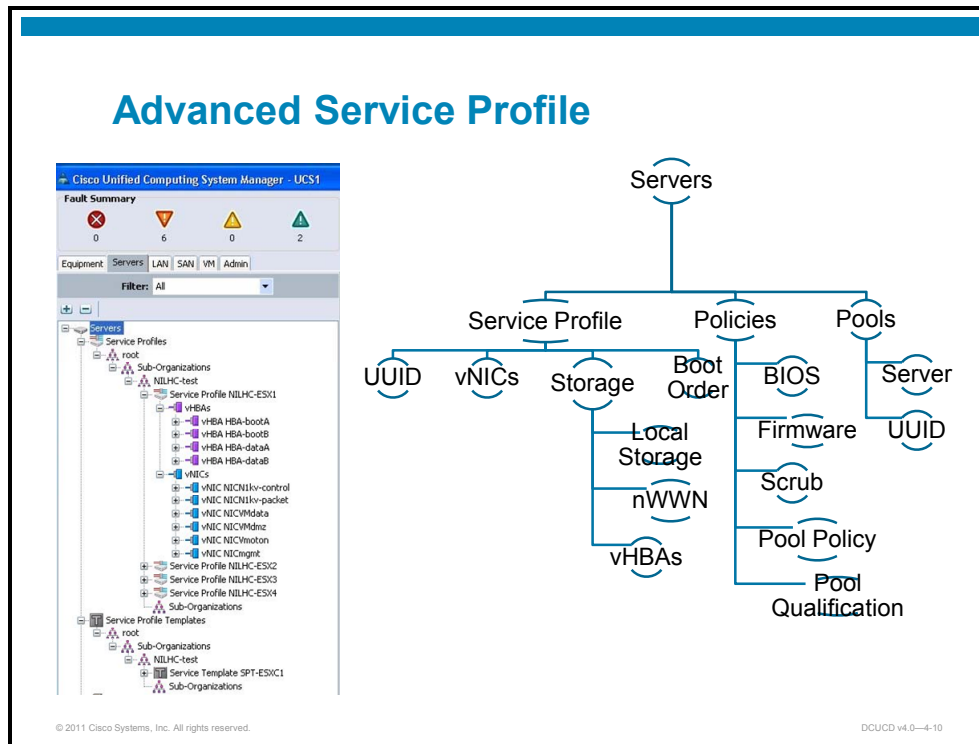
A vNIC is a virtualized network interface card that is configured on a physical network adapter and appears to be a physical NIC to the operating system of the server. The type of adapter in the system determines how many vNICs the administrator can create. For example, a Cisco UCS CNA M71KR adapter has two NICs, which means a maximum of two vNICs for each of those adapters can be created. A vNIC communicates over Ethernet and manages LAN traffic. At a minimum, each vNIC must be configured with a name and with fabric and network connectivity.

vHBA Configuration

A vHBA is a virtualized host bus adapter that is configured on a physical network adapter and appears to be a physical HBA to the operating system of the server. The type of adapter in the system determines how many vHBAs the administrator can create. For example, a Cisco UCS CNA M71KR has two HBAs, which means a maximum of two vHBAs for each of those adapters can be created. In contrast, a Cisco UCS 82598KR-CI does not have any HBAs, thus no vHBAs can be created. A vHBA communicates over Fibre Channel over Ethernet (FCoE) and manages SAN traffic. At a minimum, each vHBA must be configured with a name and with fabric connectivity.

Advanced Server Deployment Model

This topic evaluates the UCS advanced server deployment model.




Cisco UCS resource pools are collections of identities, or physical or logical resources, that are available in the system. Pools are containers for resources and identity definitions. A pool contains only one type of resource; for instance, a MAC pool must contain only MAC addresses, and a nWWN pool must contain only nWWN addresses.

All pools increase the flexibility of server deployments and allow an administrator to centrally manage the Cisco UCS system resources, which reduces the need for administrators to actively manage the use of resources.

Pool Overview

- Pool = collection of identities and resources (physical or logical)
 - Groups resources with similar characteristics
 - Assigned to service profiles or templates
 - Identity and blade pool type
- Enables central resource management
 - Eases system management and maintenance
- Default vs. user-created pools

Pool Name	→	Web Srv MAC
Pool Members	→	00 : 25 : b5 : 00 : 00 : 01
		00 : 25 : b5 : 00 : 00 : 02
		00 : 25 : b5 : 00 : 00 : 03
		00 : 25 : b5 : 00 : 00 : 04



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-11

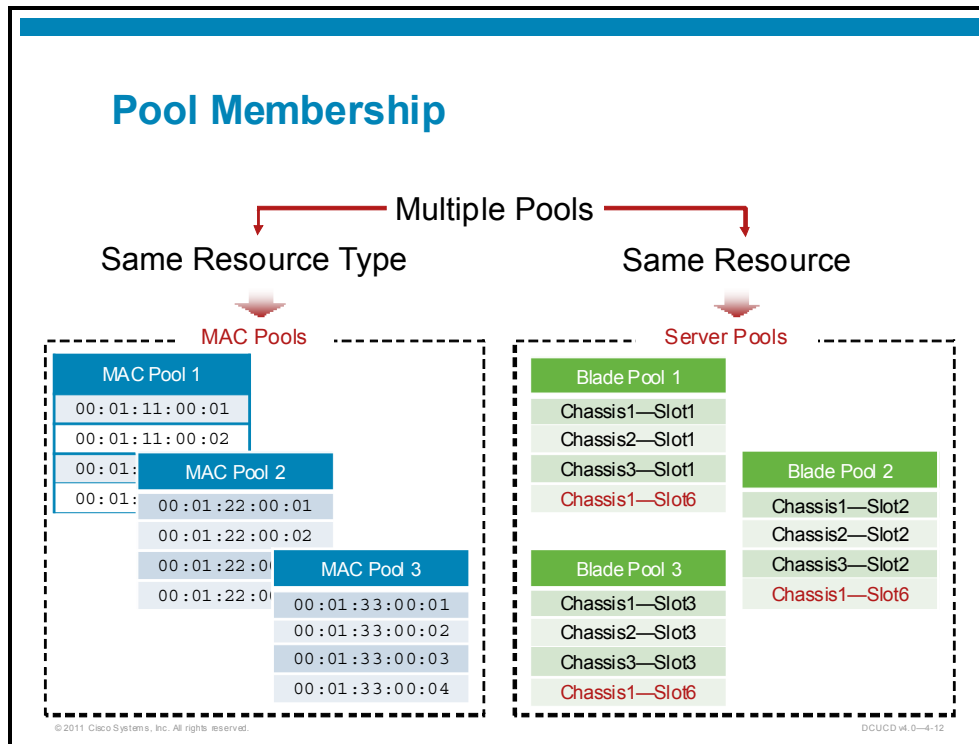
There are two pool types—identity pools and blade pools.

If a pool is used to group identity information such as MAC addresses, it can be preassigned with identity ranges for servers that will host specific information. For example, database servers could be configured to use the same range of MAC addresses, UUIDs, and WWN addresses.

Blade pools can be used to segment unconfigured servers with similar characteristics. Later, they can be associated with service profiles by means of a policy.

Default vs. User-Created Pools

Cisco UCS also has default pools that are predefined, but not prepopulated. Default pools are used as a last resource when a regular pool has been drained of its resources in the local organization, and there are no other resources available in either the parent or grandparent organization.

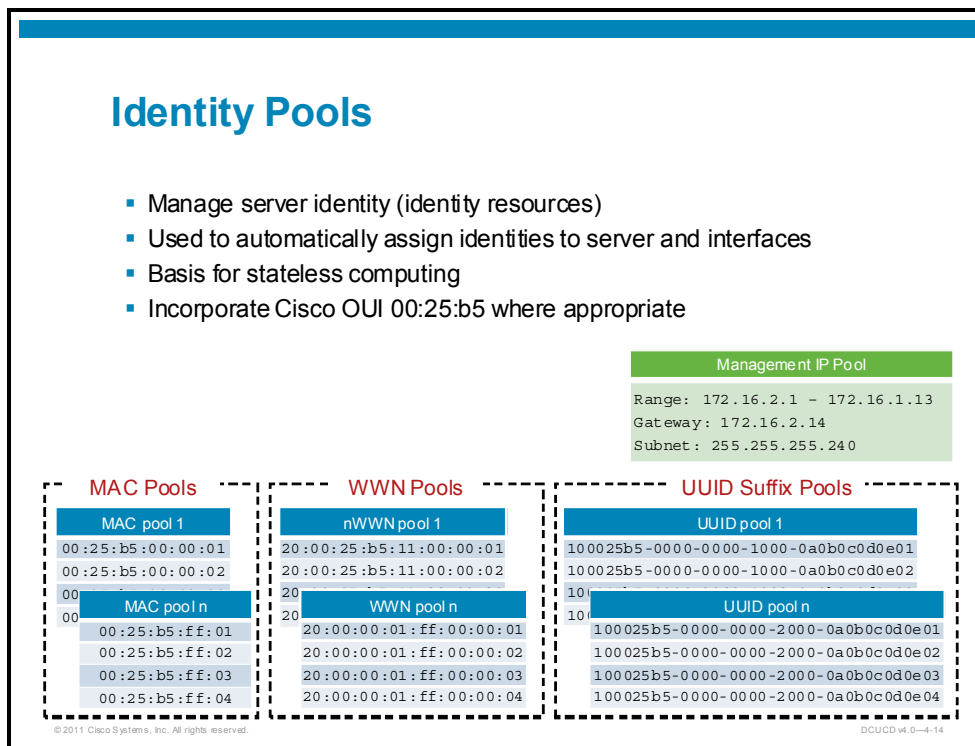


Multiple pools with different names can contain the same types of resources, and the same resources as well.

For example, a blade can be a member of multiple pools. Once a blade is consumed and assigned by the system, that blade is no longer available to any resource pools.

Designing Identity Pools

This topic evaluates UCS identity pools.



Identity resource pools are used for server identity management and assignment. They allow the Cisco UCS solution to automatically assign identities to servers and their interfaces.

Using identity pools minimizes the administrative tasks that are needed for network, storage, and server administration. Identity pools are also key components for the stateless computing model, since the physical servers are unaware of the configuration, and any server can be assigned to any identity and workload at any time.

There are various identity resource pools in existence, such as network MAC pools, storage WWN pools, UUID pools, and management IP pools.


IEEE OUI

The Organizationally Unique Identifier (OUI) is assigned by the IEEE to allow hardware manufacturers to guarantee universal uniqueness of the addressing schemes for MAC and WWN addresses.

The OUI that is assigned to Cisco is 00:25:b5.

MAC Pools

- MAC address
 - 48-bit address
 - EUI-48 standard
 - Uniquely identifies node on LAN
 - Restricted to 00:25:b5:nn:nn:nn
 - System generated or administrator assigned
 - Assigned to the server LAN adapter



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—4-15

A MAC address is a hardware address that uniquely identifies each node on a LAN. A MAC address—a 48-bit structure that is defined by the EUI-48 standard, which is a 48-bit extended universal identifier—has a form of MM:MM:MM:nn:nn:nn, where:

- MM:MM:MM (the first 24 bits) represents the OUI of the device manufacturer
- nn:nn:nn (the last 24 bits) is assigned by the organization

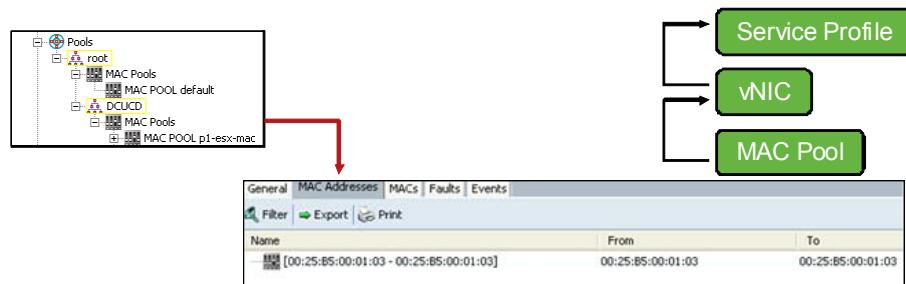
MAC addresses should be globally unique, or at least remain within an Ethernet broadcast domain.

Note A Cisco UCS MAC pool is restricted to MAC addresses that remain within the 00:25:b5:nn:nn:nn range.

When defining MAC addresses that are part of a certain pool, an administrator can encode meaning into the organizationally assigned field (the last 24 bits). Such a meaning could be the Cisco UCS cluster ID and data center, which would help preserve uniqueness within a Cisco UCS management organization.

MAC Pools (Cont.)

- MAC address pool definition
 - One or more blocks of MAC addresses
- MAC address block definition
 - First MAC address
 - Number of addresses
- Address assigned to a vNIC in service profile



MAC Resource Pools

The MAC address resource pools allow dynamic assignment of unique MAC addresses, thus allowing servers to be stateless.

A MAC pool is a collection of network MAC addresses. The pool MAC addresses are unique in the Layer 2 environment and are available to be assigned to the vNICs on the servers. An administrator who creates one or more blocks of MAC addresses in a pool populates the MAC pool.

A MAC address block is a sequence with a first MAC address and a last MAC address. The block is created by specifying a starting MAC address and number of addresses.

An example of a MAC address block is presented here:

- First MAC: 00:25:b5:00:00:01
- Last MAC: 00:25:b5:00:01:00

The MAC address resource pools allow dynamic assignment of unique MAC addresses and thus allow servers to be stateless.

When a MAC address resource pool is defined, the size of the address block should remain within reasonable boundaries, in order to omit unnecessary large blocks that add size and overhead to the Cisco UCS Manager database.

Using MAC Resource Pools

Using MAC pools with service profiles releases an administrator from having to manually configure the MAC addresses that are used by the server that is associated with the service profile.


MAC pools can be implemented in multitenancy environments. In such environments, you should use an organizational hierarchy, in order to ensure that MAC pools can be used only by specific applications or business services.

For Cisco UCS to assign the MAC address from the pool to a server, the MAC pool must be included in a vNIC policy, which must be included in the service profile that represents the server.

The MAC address pools typically should be defined in the initial phase of Cisco UCS deployment to allow all stakeholders to agree upon the addressing protocols that will be used later on in the service profiles.

WWN Pools

- WWN address
 - 64-bit address
 - Uniquely identifies node or port on SAN
 - Restricted to
 - 20:nn:nn:nn:nn:nn:nn:nn range
 - 20:00:00:25:b5:nn:nn:nn range
 - 5n:nn:nn:nn:nn:nn:nn:nn range
 - System generated or administrator assigned
- nWWN address pool
 - nWWN assigned to server
- pWWN address pool
 - pWWN assigned to vHBA



A WWN is an address that uniquely identifies a server in a SAN environment. WWN addresses are used for two purposes:

- An nWWN is used to address the server.
- A pWWN is used to address each port on the host bus adapter (vHBA).

The nWWN and pWWN address resource pools remove the server dependency of SAN resources so that hardware changes remain transparent to the storage network and devices. The zoning configuration, which couples initiators and targets, does not need to change upon hardware changes.

Note A WWN pool can include only nWWNs or pWWNs in the 20:nn:nn:nn:nn:nn:nn:nn, 20:00:00:25:b5:nn:nn:nn, or 5n:nn:nn:nn:nn:nn:nn:nn range. All other WWN ranges are reserved.

Defining WWN Addresses

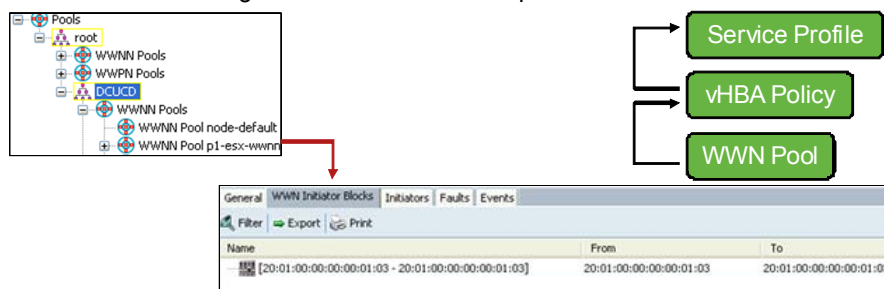
First, WWN addresses should be globally unique. To achieve this within Cisco UCS and maintain standards, it is recommended that the WWN address incorporate an OUI. This way, the WWN address will not conflict with any other WWN addresses used in the SAN of an organization.

Second, the WWN address can also encode a Cisco cluster ID and data center, which would help preserve uniqueness within an organization.

Third, to distinguish between the nWWN and pWWN addresses, the second octet can encode meaning for the nWWN or pWWN—for example, 00 for nWWN and 01 for pWWN.

WWN Pools (Cont.)

- WWN address pool
 - One or more blocks of WWN addresses
 - Optionally one or more initiators
- WWN address block definition
 - First WWN address
 - Number of addresses
- Address assigned to a vHBA in service profile



WWN Resource Pool

A WWN pool is a collection of WWNs for use by the Fibre Channel vHBAs in a Cisco UCS solution. The WWN pool contains blocks of WWN addresses, similar to the MAC pool, and is created and populated by an administrator who defines and configures one or more blocks of WWN addresses in the pool. Each WWN block or individual WWN can be assigned a boot target.

A WWN address block is a sequence that starts with the first WWN address and ends with the last WWN address. The block is created by specifying a starting WWN address and number of addresses.

An example of a WWN address block is presented here:

- First WWN: 20:00:00:01:11:00:00:01
- Last WWN: 20:00:00:01:11:00:01:00

When a WWN address resource pool is defined, the size of the address block should remain within reasonable boundaries, in order to omit unnecessary large blocks that add size and overhead to the Cisco UCS Manager database.

Within the Cisco UCS, separate nWWN and pWWN pools are created.

nWWN Pools

An nWWN pool is a WWN pool that contains only world wide node names. If included in a service profile, the associated server will be assigned an nWWN from that pool.

pWWN Pools

A pWWN pool is a WWN pool that contains only world wide port names. If included in a service profile, the port on each vHBA of the associated server will be assigned a pWWN from that pool.

Using WWN Resource Pools

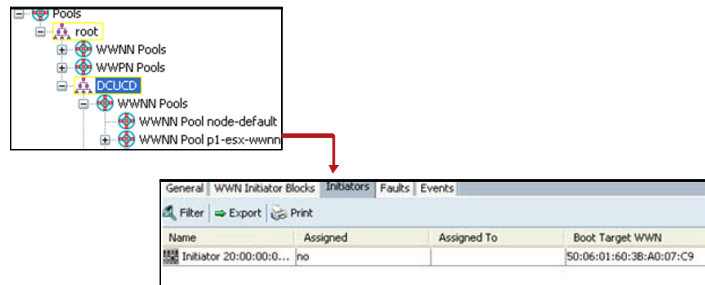
Using WWN pools in service profiles relieves an administrator from manually configuring the WWNs that will be used by the server that is associated with the service profile.

In multitenancy environments, a WWN pool can be used to control the WWNs used by each organization.

Similar to the MAC address pools, the WWN address pools should typically be defined in the initial phase of Cisco UCS deployment to allow all stakeholders to agree upon the addressing that will be used later on in the service profiles.

WWN Initiator

- SAN boot configuration control
 - Associates SAN boot LUN with specific WWN
 - Part of the WWN pool configuration
 - Based on the assigned nWWN or pWWN
 - Typical pWWN is used



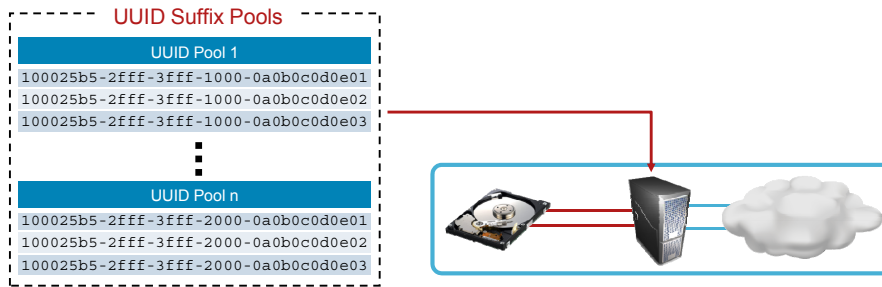
When defining a WWN address block in a WWN address pool, the administrator can also define the initiator-target relationship to associate SAN boot logical unit number (LUN) information with a specific entry in the WWN pool.

The WWN initiator can be defined for either the nWWN or pWWN address. The relationship is typically defined with a pWWN server address.

UUID Suffix Pools

- UUID
 - 128-bit number
 - Uniquely identifies computer node (server)
 - System generated or administrator assigned
 - Consists of prefix (64 bits) and suffix (64 bits)

UUID Prefix (64 bits)	UUID Suffix (64 bits)
FFFFFFFF-FFFF-FFFF	HHHH-HHHHHHHHHHHHH



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-20

A UUID is an identifier that uniquely identifies each computing resource-server.

A UUID consists of a prefix and a suffix. Within Cisco UCS, a suffix UUID is 8 bytes long, consists of 16 hexadecimal characters, and uses the format HHHH-HHHHHHHHHHHHH. The pool is prefixed with eight bytes that are unique to Cisco UCS, in the format FFFFFFFF-FFFF-FFFF.

The UUID assigned to the server is a “prefix-suffix” number in the format FFFFFFFF-FFFF-FFFF-HHHH-HHHHHHHHHHHHH.

Like the MAC and WWN addresses, the UUID can also encode the OUI to maintain global uniqueness. The administrator can also encode the UCD cluster ID and data center into the UUID to further maintain uniqueness in the organization.

UUID Suffix Pools (Cont.)

- UUID suffix collection
 - Configurable UUID prefix per pool
 - One or more blocks of UUID suffixes
- UUID suffix assigned to a service profile

The screenshot displays the 'UUID Suffix Pools' configuration page in Cisco UCS Manager. It features a table with columns for Name, Pool Name, and UUID Prefix. Below the table are 'Filter', 'Export', and 'Print' options. To the right, a diagram shows a 'Service Profile' box connected to a 'UUID Pool' box. Below this, a detailed view of a 'UUID Suffix' is shown with tabs for 'General', 'UUID Suffixes', 'UUID Blocks', 'Faults', and 'Events'. The 'UUID Suffixes' tab is active, showing a table with columns for Name and Assigned. The first row shows the name '1000-000000000003' and the assigned status 'no'.

Name	Pool Name	UUID Prefix
Pool p1-esx-uuid	p1-esx-uuid	10000000-0000-0000
Pool p1-ln-uuid	p1-ln-uuid	10000000-0000-0000
Pool p1-win-uuid	p1-win-uuid	10000000-0000-0000
Pool p2-esx-uuid	p2-esx-uuid	10000000-0000-0000

Name	Assigned
1000-000000000003	no

A Cisco UCS UUID suffix pool is a collection of BIOS UUID suffixes that are available to be assigned to servers. The UUID suffix pool ensures that these variable values are unique for each server that is associated with a service profile, which uses that particular pool to avoid conflicts.

The suffix UUIDs are populated by an administrator, who creates one or more blocks of UUID numbers in the pool. The UUID suffix pool is created by specifying the UUID prefix and the block of UUID suffixes. A UUID suffix block is created by specifying the first UUID suffix and the number of UUID suffixes.

An example of UUID suffix block definition is presented here:

- UUID prefix: 10000000-0000-0000
- First UUID: 1000-000000000001
- Last UUID: 1000-0000000000255

When a UUID suffix resource pool is defined, the size of the address block should remain within reasonable boundaries in order to omit unnecessary large blocks that add size and overhead to the Cisco UCS Manager database.


Using UUID Suffix Pools

Using UUID suffix pools in service profiles relieves the administrator from having to manually configure the UUID of the server that is associated with the service profile.

Management IP Pool

- External IP addresses collection
 - Used for management access to the servers
 - One or many IP address blocks per pool
 - One pool per organization
- Used for external access to the server (blade)
 - Keyboard, video, mouse (KVM)
 - Serial over LAN
 - IPMI

Management IP Pool
Range: 172.16.1.1 - 172.16.1.253
Gateway: 172.16.1.254
Subnet: 255.255.255.0
Range: 172.16.2.1 - 172.16.1.13
Gateway: 172.16.2.14
Subnet: 255.255.255.240



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-22

The management IP pool is a collection of external IP addresses. The Cisco UCS Manager reserves each block of IP addresses in the management IP pool for external access that terminates in the server BMC within a server. The Cisco UCS Manager uses the IP addresses in a management IP pool for external access such as KVM, Serial over LAN (SoL), and IPMI.

The management IP pool address block is created by specifying the following:

- First and last IP address in the address range
- Gateway IP address
- Subnet mask

Remember that the management IP pool addresses should be from the same IP subnet as the Cisco UCS fabric interconnect switches and UCS cluster management IP addresses.

Management IP pools can be created per organization.

Note You cannot create (or delete) a management IP pool. You can only enter (scope to) the existing default pool.

Designing Resource Pools

This topic evaluates Cisco UCS resource pools.

Server Pools

- Set of servers (blades)
 - Typically share same characteristics (memory capacity, CPU type, and so on)
 - Blade from any chassis in the system
- Individual server blade can be member of multiple server pools

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-4-24

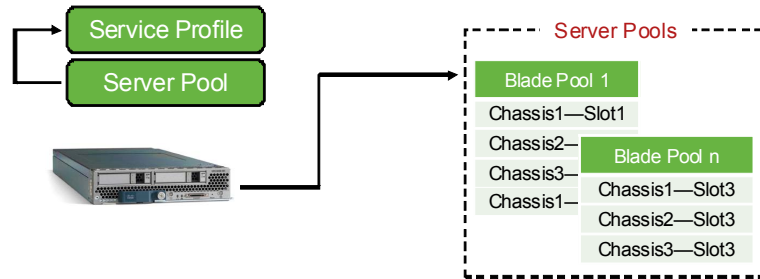
Server blades can also be pooled. A server pool contains a set of available stateless servers. These servers typically share the same characteristics. These shared characteristics may be the location of each server in the chassis or an attribute such as server type, amount of memory, local storage, type of CPU, or local drive configuration.

A server pool can include servers from any chassis in the system. A given server can belong to multiple server pools.

Server pools can be used in multitenancy deployments, in which organizations designate one or more server pools to be used by a specific organization. For example, a pool that includes all servers with two CPUs could be assigned to the Marketing organization, whereas all servers with 64 GB of memory could be assigned to the Finance organization.

Server Pools (Cont.)

- Population of server pools
 - Manual—administratively defined
 - Automatic—using policies upon discovery
- Assigned to a service profile
 - Discovered, nonassociated blades can be assigned



The Cisco UCS administrator has two options to populate the server pool:

- Manually assign a server to a server pool
- Use server pool policies and server pool policy qualifications to automate the assignment

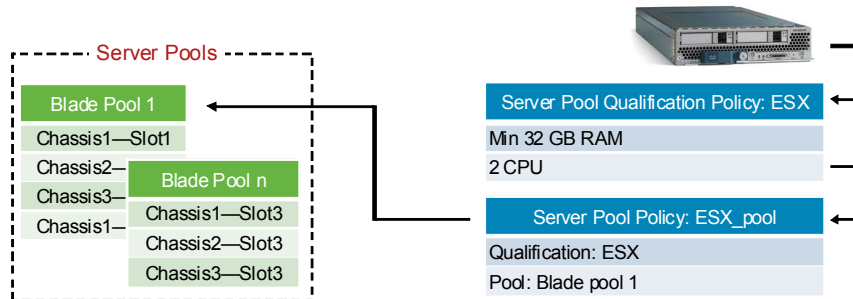
Whether the server pools will be manually or automatically populated, each pool must first be defined with a name.

Manual Population of Pools

Manual population is a simple—but not very efficient—way to assign pool membership for a server. The administrator manually maintains and manages the server membership, and must manually evaluate inventory of each server in order to assign it to one or more server pools.

Automatic Server Pool Population

- Upon server blade discovery only
- Using policies
 - Server pool qualification policy—defines qualification criteria
 - Server pool policy—ties together qualification and pool



A server pool can be automatically populated via policies. Two different policies are used together to determine which pool a server should be a member of:

- Server pool qualification policy
- Server pool policy

Note Server pool automatic population happens when an individual server blade is discovered. If the administrator wants to apply automatic population to existing server blades, they must be reacknowledged.

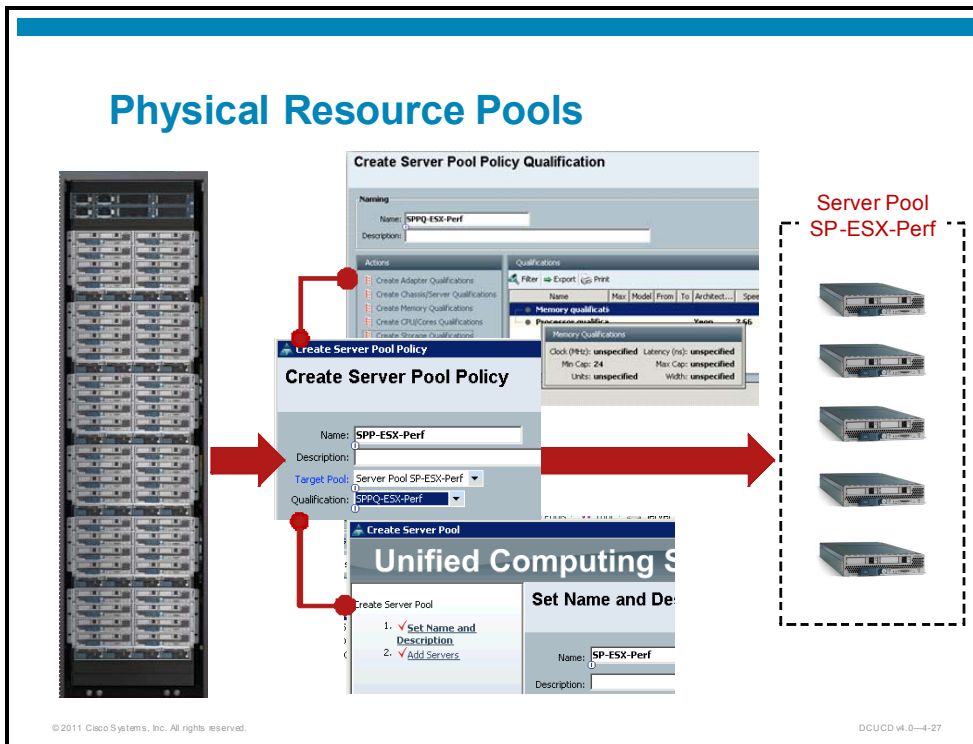
Server Pool Qualification Policy

A server pool qualification policy describes the hardware qualification criteria for servers, like number of CPUs, memory capacity, type of adapters, and so on. A server that fulfills all of the defined criteria in the policy is considered a qualified server.

Server Pool Policy

A server pool policy describes which server pools the server becomes a member of, if it has qualified for a certain server pool qualification. The same server can meet multiple qualifications; thus, multiple pool policies can point to the same pools.

Physical Resource Pools



An administrator can use a server pool instead of a physical server blade to quickly deploy Cisco UCS service profiles that are based on predefined requirements. An individual server blade must be associated with only one service profile at a time.

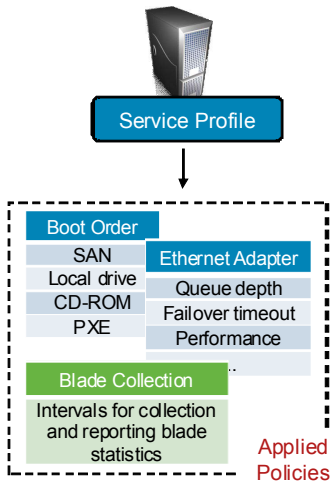
When a server pool is associated with a service profile, Cisco UCS Manager automatically selects an available server blade—one that is currently discovered, is not associated with any other service profile, and is not in the process of being associated or disassociated with a service profile.

Design Policies

This topic evaluates UCS policies.

Policies Overview

- Policy-driven management of Cisco UCS components
 - Ensures consistent computing environment
- Policy
 - Determines how Cisco UCS components act in given circumstances
 - Separates functions within a system
 - Different policies defined for network, storage, server
- Used within service profiles



The diagram illustrates the relationship between a Service Profile and Applied Policies. A Service Profile (represented by a server icon) points to a dashed box containing a list of policies. The policies are categorized into three groups: Boot Order (SAN, Local drive, CD-ROM, PXE), Ethernet Adapter (Queue depth, Failover timeout, Performance), and Blade Collection (Intervals for collection and reporting blade statistics). The label 'Applied Policies' is positioned to the right of the dashed box.

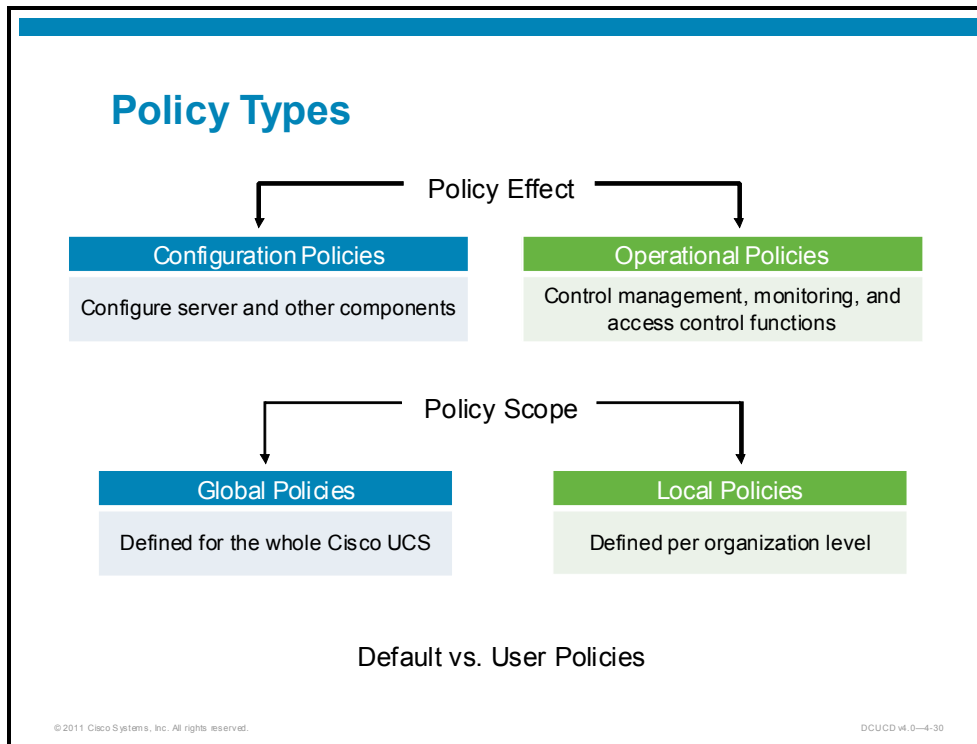
© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-4-29

Cisco UCS Manager is a policy-driven management device manager. Policy-driven management is one of the key features of Cisco UCS Manager, and it helps IT organizations to better define and implement their own best practices. Since data centers are becoming increasingly more dynamic, the definition, consumption, and resolution of policies is a key enabling technology for making infrastructure devices portable and able to be reconfigured dynamically.

Policies help to ensure that consistent, tested, and compliant systems are used, reducing the risk of issues that are caused by repetitive manual tasks. Rules are defined in the form of policies inside the Cisco UCS Manager. Policies are centrally defined and enforced at the endpoint; they can be used to perform policy enforcement on any device within the same Cisco UCS. Such an approach removes state from a device (blade) and thus improves mobility and scalability. An administrator can import and export policies into Cisco UCS.

Policies determine how Cisco UCS components will act in specific circumstances. An administrator can create multiple instances of most policies. For example, an administrator might want different boot policies, so that some servers can Preboot Execution Environment (PXE) boot, some can SAN boot, and others can boot from local storage.

Policies also allow separation of functions within the system. A subject matter expert can define policies that are used in a service profile, which can then be created by someone who does not have that subject matter expertise. For example, a LAN administrator can create adapter policies and quality of service policies for the system. Someone who has limited or no subject matter expertise with LAN administration can then use these policies to create a service profile.



Policy Types

An administrator can create and use two types of policies in the Cisco UCS Manager:

- Configuration policies, which configure the servers and other components
- Operational policies, which control certain management, monitoring, and access control functions

Configuration policies make up the majority of the policies in the Cisco UCS and are used to describe configurations of different components of the system. Operational policies determine how the system behaves under specific circumstances.

Policy Scope

Policies also have a scope—they can be either global or local. Global policies are defined for the entire Cisco UCS solution, whereas local policies are defined per organization level. Examples of global policies include fault collection policies, Call Home policies, and the various statistics collection policies.

Default vs. User Policies

Cisco UCS has default policies that are used in the absence of other user-specific created policies. The predefined default policies typically default to normal system operation behavior.

Configuration Policies

Policy	Policy Description
Autoconfiguration	New server autoconfiguration actions
Boot	Server boot location (SAN, LAN, local disk, virtual media)
Chassis discovery	System action upon new chassis discovery
Dynamic connection	VNTag to VM and dynamic vNICs connectivity configuration
Ethernet adapter	Ethernet adapter settings (queue depth, failover timeout, performance, etc.)
Fibre Channel adapter	Fibre Channel adapter traffic settings (FLOGI and PLOGI timeouts, error handling, and so on)
Host firmware package	Common set of server firmware versions
IPMI access profile	IPMI access definition

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-31

The table lists and describes the available Cisco UCS configuration policies.

Note The policy must typically be included in a service profile, and that service profile must be associated with a server for the policy to take effect; otherwise, the server uses the default settings.

Server Autoconfiguration Policy

This policy determines whether one or more of the following is automatically applied to a new server:

- A server pool policy qualification that qualifies the server for one or more server pools
- An organization
- A service profile template that associates the server with a service profile created from that template

Boot Policy

This policy determines the following:

- Configuration of the boot device
- Location from which the server boots
- Order in which boot devices are invoked

An administrator can choose to have associated servers boot from a local device, such as a local disk or virtual CD-ROM, or can select a SAN boot or a LAN (PXE) boot.

Note Changes to a boot policy may be propagated to all servers that are created with an updating service profile template that includes that boot policy. Reassociation of the service profile with the server to rewrite the boot order information in the BIOS is automatically triggered.

SAN Boot

Boots the server from an operating system image on the SAN. A primary and a secondary SAN boot can be specified so that if the primary boot fails, the secondary will be used. It is recommended to use a SAN boot, because it offers the most service profile mobility within the system; moving a service profile from one server to another ensures that the new server boots from the same operating system image.

LAN Boot

Boots from a centralized provisioning server and is frequently used to install operating systems on a server from that server.

Local Disk Boot

If the server has a local drive, boots from that drive.

Virtual Media Boot

Mimics the insertion of a physical CD-ROM disk (read-only) or floppy disk (read-write) into a server, and is typically used to manually install operating systems on a server.

Chassis Discovery Policy

The discovery policy determines how the system reacts when you add a new chassis. Creating a chassis discovery policy enables the system to perform the following tasks:

- Automatically configure the chassis for the number of links between the chassis and the fabric interconnect specified in the policy
- Specify the power policy to be used by the chassis

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter—including how the adapter manages traffic—and are dependent on the operating system type. An administrator can use these policies to change default settings for the following:

- Queues
- Interrupt handling
- Performance enhancement
- Really Simple Syndication (RSS) hash
- Failover in a cluster configuration with two fabric interconnects

Host Firmware Package

This policy enables you to specify a common set of firmware versions that make up the host firmware package. The host firmware includes the following server and adapter components:

- BIOS
- Single-attachment station (SAS) controller
- Emulex option ROM
- Emulex firmware
- QLogic option ROM
- Adapter firmware

The firmware package is pushed to all servers that are associated with service profiles that include this policy. This policy ensures that the host firmware is identical on all servers that are associated with service profiles that use the same policy. Thus, the firmware versions are maintained if the service profile is moved from one server to another.

Note You must ensure that the appropriate firmware has been downloaded to the fabric interconnect.

IPMI Access Profile

This policy allows the administrator to determine whether IPMI commands can be sent directly to the server, using the IP address. It also defines the IPMI access, including a username and password that can be authenticated locally on the server, and it defines whether the access is read-only or read-write.

Configuration Policies (Cont.)

Policy	Policy Description
Local disk configuration	Optional local disk drive configuration and RAID settings
Management firmware	BMC firmware version
QoS definitions	Outgoing QoS for vNIC and vHBA (CoS, burst, rate)
Server discovery	System action upon new server discovery
Server inheritance	Inheritance from the hardware on a new server discovery
Server pool	Pool membership per specific server pool qualification policy
Server pool qualification	Server qualification per inventory rules (memory capacity, number of CPUs)
vHBA and vNIC	Connectivity and quality template for vHBA and vNIC

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-32

Local Disk Configuration Policy

This policy configures any optional SAS local drives that have been installed on a server through the onboard Redundant Array of Independent Disks (RAID) controller of the local drive. This policy enables an administrator to set the RAID mode and the way the drives are partitioned.

Management Firmware Package

This policy enables an administrator to specify a common set of firmware versions that make up the management firmware package. The management firmware includes the server controller, or BMC, on the server. The firmware package is pushed to all servers that are associated with service profiles that include this policy. This policy ensures that the BMC firmware is identical on all servers that are associated with service profiles that use the same policy. Therefore, if the service profile is moved from one server to another, the firmware versions are maintained.

Note You must ensure that the appropriate firmware has been downloaded to the fabric interconnect.

QoS Policy

This policy defines the outgoing quality of service (QoS) for a vNIC or vHBA—a system class to the traffic egress for an individual vNIC or vHBA is assigned.

Server Discovery Policy

This discovery policy determines how the system reacts when you add a new server. With a server discovery policy, an administrator can control whether the system conducts a deep discovery when a server is added to a chassis, or whether a user must first acknowledge the new server. By default, the system conducts a complete discovery. With this policy, an inventory of the server is conducted, and then server pool policy qualifications are run to determine whether the new server qualifies for one or more server pools.

Server Inheritance Policy

This policy is invoked during the server discovery process in order to create a service profile for the server. All service profiles that are created from the policy use the values that are burned into the blade at manufacture. The policy performs the following:

- Analyzes the inventory of the server
- If configured, assigns the server to the selected organization
- Creates a service profile for the server with the identity burned into the server at manufacture

Note You cannot migrate a service profile that was created with this policy to another server.

Server Pool Policy

This policy is invoked during the server discovery process. It determines what happens if server pool policy qualifications match a server to the target pool specified in the policy. If a server qualifies for more than one pool, and those pools have server pool policies, the server is added to all those pools.

Server Pool Policy Qualifications

This policy qualifies servers based on a server inventory that is conducted during the discovery process. The qualifications are individual rules that you configure in the policy to determine whether a server meets the selection criteria. An administrator can create a rule that specifies the minimum memory capacity for servers in a data center pool.

Qualifications are used in other policies to place servers, not just by the server pool policies. For example, if a server meets the criteria in a qualification policy, it can be added to one or more server pools, or it can be automatically associated with a service profile. Depending upon the implementation, you may include server pool policy qualifications in the following policies:

- Autoconfiguration policy
- Chassis discovery policy
- Server discovery policy
- Server inheritance policy
- Server pool policy

vHBA Template

This policy defines how a vHBA on a server connects to the SAN. This policy is also referred to as a vHBA SAN connectivity template.

vNIC Template

This policy defines how a vNIC on a server connects to the LAN. This policy is also referred to as a vNIC LAN connectivity template.

Operational Policies

Policy	Policy Description
Adapter collection	Adapter collection and reporting statistics intervals
Blade collection	Blade collection and reporting statistics intervals
Call Home	Call Home profiles for email notification
Chassis collection	Chassis collection and reporting statistics intervals
IPMI profile	Server IPMI capability and read-only or read/write access type
Fault collection	Fault action and clearance retention intervals
Port collection	Ports collection and reporting statistics intervals
Scrub	Defines if server state is to be kept during discovery
Serial over LAN	Serial over LAN server capabilities
Threshold	Set alarm triggers for Ethernet, Fibre Channel, adapters, blades, chassis, PSU, FEXs, FANs, and so on

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-33

The table lists and describes the available Cisco UCS operational policies.

Scrub Policy

This policy determines what happens to local data on a server during the discovery process and when the server is disassociated from a service profile. This policy can ensure that the data on local drives is erased at those times.

Serial over LAN Policy

This policy sets the configuration for the SoL connection for all servers that are associated with service profiles that use the policy. By default, the SoL connection is disabled. If this policy is implemented, it is recommended that you also create an IPMI profile.

Statistics Collection Policy

A statistics collection policy defines how frequently statistics are to be collected (collection interval), and how frequently the statistics are to be reported (reporting interval). Reporting intervals are longer than collection intervals so that multiple statistical data points can be collected during the reporting interval, which provides Cisco UCS Manager with sufficient data to calculate and report minimum, maximum, and average values. Statistics can be collected and reported for the following five functional areas of the Cisco UCS system:

- **Adapter:** Statistics that are related to the adapters in the fabric interconnect
- **Chassis:** Statistics that are related to the blade chassis
- **Host:** This policy is a placeholder for future support
- **Port:** Statistics that are related to the ports, including server ports, uplink Ethernet ports, and uplink Fibre Channel ports
- **Server:** Statistics that are related to servers

Note Cisco UCS Manager has one default statistics collection policy for each of the five functional areas. You cannot create additional statistics collection policies, and you cannot delete the existing default policies. You can only modify the default policies.

Statistics Threshold Policy

A statistics threshold policy monitors statistics about certain aspects of the system and generates an event if the threshold is crossed. An administrator can set both minimum and maximum thresholds. For example, an administrator can configure the policy to raise an alarm if the CPU temperature exceeds a certain value, or if a server is overutilized or underutilized.

These threshold policies do not control the hardware or device-level thresholds that are enforced by endpoints, such as the BMC. Those thresholds are burned in to the hardware components at manufacture. Cisco UCS enables you to configure statistics threshold policies for the following components:

- Servers and server components
- Uplink Ethernet ports
- Ethernet server ports, chassis, and fabric interconnects
- Fibre Channel ports

Note You cannot create or delete a statistics threshold policy for Ethernet server ports, uplink Ethernet ports, or uplink Fibre Channel ports. You can only configure the existing default policy.

Cisco UCS 1.4 Enhancements

This topic evaluates UCS 1.4 enhancements from the server deployment perspective.

Service Profile Deployment Scheduling

- Schedule disruptions
- Maintenance policies
 - New policy contains schedule when service profile can be changed
 - Policy can be mapped to one or more service profiles
- Benefits
 - Service profile disruptions in maintenance windows
 - Plan design changes for service profiles

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-4-35

Service profile deployment has been enhanced with scheduling, which enables the following:

- Scheduled service profile deployments
- Awareness of maintenance windows
- Resource reservations


These are the benefits of scheduled deployment and changes:

- Service profile disruptions can be kept in maintenance windows.
- Administrators can plan design changes for service profiles.

To deploy the functionality, new maintenance policies are defined which contain a schedule that governs when service profile changes can be implemented.

Cisco IMC IP Address in Service Profile

- Second IP address can be assigned for Cisco IMC
 - IP address associated with service profile
 - When service profile moves, second IP address moves to new server



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-4-36

With the release of Cisco UCS 1.4, administrators can now use a second IP address for server management—an IP address attached to the Cisco Integrated Management Controller (IMC) that follows the service profile if it is moved to another blade. The IP address is actually associated with service profile.

Server management can now be achieved using a regular system-assigned IP address and the administratively defined IP address.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The Cisco UCS server deployment model abstracts hardware resources from the server.
- Cisco UCS resource pools are collections of identities and resources.
- MAC, WWN, and UUID suffix pools are identity resource pools that are used to automatically assign identities to a server.
- Server blades can be part of multiple server pools.
- A Cisco UCS service profile is a logical representation of a server and is the basis for stateless computing.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-37

Summary (Cont.)

- A service profile can use policies and pools to configure a server when associated with a blade.
- Cisco UCS service profile templates ease deployment of many similar service profiles.
- Cisco UCS policies separate functions within a system.
- Configuration policies set the server and server components.
- Operational policies control management and monitoring aspects of a server.

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-4-38

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- A service profile defines server LAN and SAN connectivity.
- A server pool qualification policy is used to place server blades into a server pool automatically.
- Management IP address pools must be large enough to address all the server blades in the Cisco UCS solution.
- Resource pools are used to manage MAC, nWWN, and pWWN addresses, as well as UUID suffix identifiers.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which MAC address would you be able to use in Cisco UCS? (Source: Designing the Cisco UCS Server Deployment Model)
- A) any
 - B) 00:25:b5:11:22:33
 - C) 20:00:00:00:00:01
 - D) 0c:10:20:34:ab:01
- Q2) Which two policies are used for automatic server pool selection? (Choose two.) (Source: Designing the Cisco UCS Server Deployment Model)
- A) server pool qualification policy
 - B) server pool policy
 - C) chassis discovery policy
 - D) threshold policy
 - E) blade discovery policy
- Q3) What is the requirement for the management IP pool? (Source: Designing the Cisco UCS Server Deployment Model)
- A) Two IP addresses per blade are required.
 - B) The IP address pool should be from the same subnet as the cluster IP address.
 - C) Only private IPv4 addresses can be used.
 - D) IP addresses should be contiguous.

Module Self-Check Answer Key

- Q1) B
- Q2) A, B
- Q3) A

Design Cisco UCS Solution Management

Overview

This module identifies and describes how to propose a Cisco Unified Computing System (UCS) solution management design for a given environment.

Objectives

Upon completing this module, you will be able to propose a Cisco UCS management design for a given environment. This includes the ability to meet this objective:

- Identify and describe the appropriate management tools and management hierarchy for Cisco UCS setup

Examining Cisco UCS Solution Management

Overview

This lesson identifies and describes the appropriate management tools and management hierarchy for Cisco Unified Computing Solution (UCS) setup.

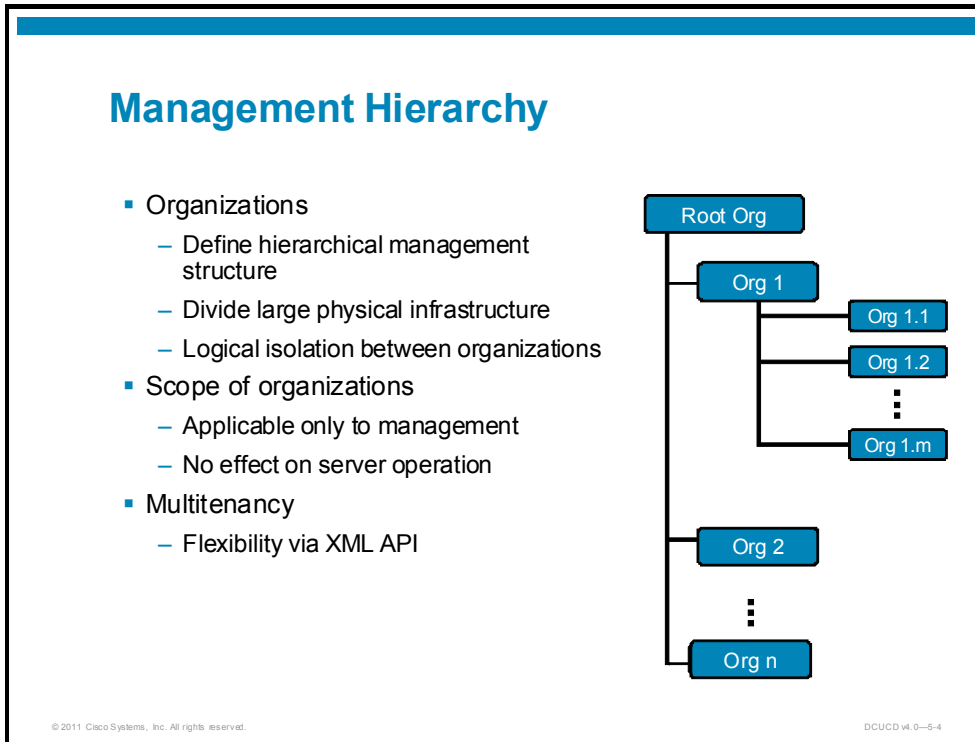
Objectives

Upon completing this lesson, you will be able to explain the management of the Cisco UCS solution. This includes the ability to meet these objectives:

- Identify solution management aspects
- Identify tools that are used to manage multiple Cisco UCS pods
- Identify Cisco UCS version 1.4 enhancements from the management aspect

Cisco UCS Management Aspects

This topic identifies Cisco UCS management aspects.



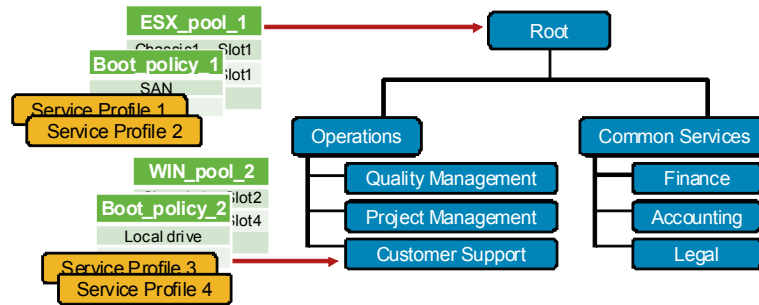
Cisco UCS enables an administrator to divide a large physical infrastructure into logical entities, providing logical isolation and eliminating the need to use a dedicated physical infrastructure for each organization. Such entities are called organizations within Cisco UCS.

Organizations are placed into a tree structure, depending on the business needs of the company. Organizations affect only the management aspect and have no impact on server operation and management.

Although the Cisco UCS Manager offers flexibility and configuration, the best flexibility is offered through a rich configuration set that is available via an XML API. The XML API is also the means through which more thorough multitenancy can be applied.

Organization Structure

- Multiple hierarchy levels
 - Root = base organization
 - Other organizations are root members
- Organization content—service profiles, policies, pools
 - Not owned by organization
- Multitenant environments

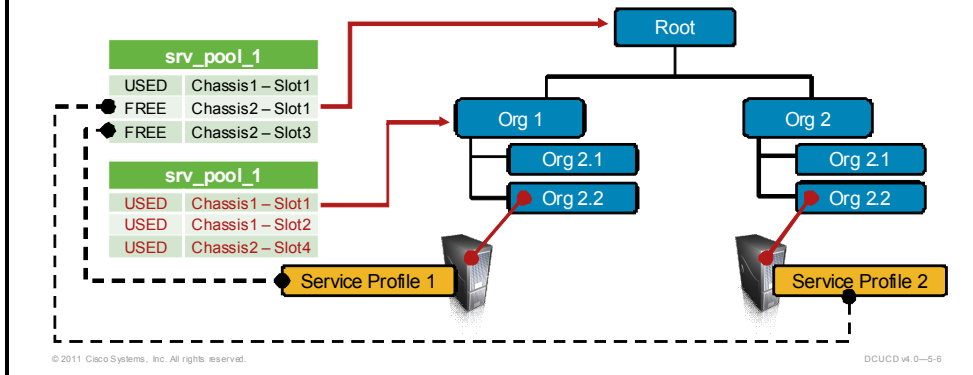


Organizations are hierarchically organized into a tree, where the top-level organization is always the root. Cisco UCS organizations enable an administrator to assign unique resources via related organizations; the resources can include policies, pools, and service profiles and templates.

The policies and pools that you create in root are systemwide and are available to all organizations in the system. However, any policies and pools that are created in other organizations are available only to organizations that are above them in the same hierarchy. For example, if a system has organizations named Finance and HR that are not in the same hierarchy, Finance cannot use any policies in the HR organization, and HR cannot access any policies in the Finance organization. However, both Finance and HR can use policies and pools in the root organization.

Organization Pool Resolution

- Service profile associated with own or parent organization pools
- Hierarchical “bottom to top” tree resolution
 - Search tree for available resource in pools with identical name
 - “First come, first serve” principle
- All pool types (blade server, MAC, WWN, UUID)



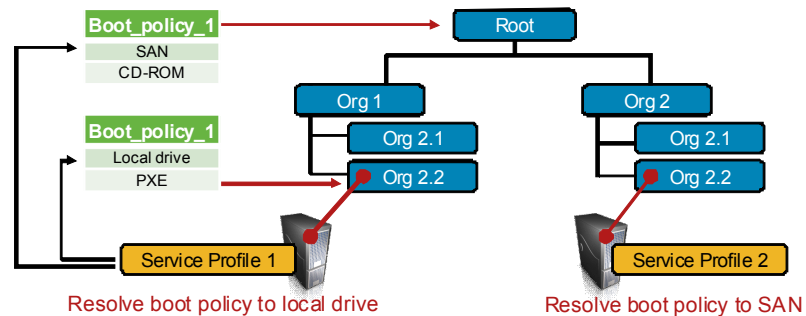
Organizations use a hierarchical pool resolution method; in other words, a pool is resolved to a service profile from the bottom to the top of the organization tree. This approach is used for any pool type—whether a blade server, MAC, world wide name (WWN), or universally unique identifier (UUID) pool.

The resolution process is performed in the following way:

- Cisco UCS Manager checks for pools with the specified name within the organization that are assigned to the service profile or policy.
- If an available resource is found inside a pool, Cisco UCS Manager uses that resource.
- If the pool does not have any available resources at the local level, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a pool with the same name. This is repeated until the search reaches the root organization.
- If the search reaches the root organization and has not found an available resource, Cisco UCS Manager returns to the local organization and begins to search for a default policy or available resource in the default pool.
- If an available resource in a default pool is found, it is used. If the pool does not have any available resources, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a default pool. Again, this step is repeated until the search reaches the root organization.
- If Cisco UCS Manager cannot find an available resource in the hierarchy, it returns an allocation error.

Organization Policy Resolution

- Service profile associated with own or parent organization policies
 - Only policies with different names are applicable
- Hierarchical “bottom to top” tree resolution
 - “First match” policy principle
- All policy types (boot, firmware, etc.)




Organizations use a hierarchical policy resolution method; in other words, a policy is resolved to a service profile from the bottom to the top of the organization tree. If there are policies of the same type using the same name but different settings, only the policy from the same organization level can be applied to a service profile.

On the other hand, a service profile on a certain organization level can be associated with policies from different, higher organizational levels (only the higher levels), as long as the policies use different names.

The resolution process is similar to the one used for pools, with the difference that the policy used is the first one in the organizational tree that matches the assigned policy name, starting from the same level as the service profile. If no policy is available, the default policy is used.

RBAC Overview

- RBAC = method of authorizing user system access
 - Based on user roles and privileges
 - Ability to grant privileges based on user responsibilities
 - Delegates management rights
- Does not control access to server-deployed operating system




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-8

Role-based access control (RBAC) is a method of restricting or authorizing system access for users, based on user roles and locales.

Roles and Privileges

- Privilege—individual rule
- Role—collection of privileges
 - Grants read-only or read/write access
- User
 - Assigned one or more roles
 - Has superset of privileges of assigned roles



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-5-9

A role defines the privileges of a user in the system. Privileges give users that are assigned to user roles access to specific system resources and permission to perform specific tasks. There are default privileges and user roles are given those privileges in Cisco UCS.

Roles can be created, modified to add new privileges or remove existing privileges, or deleted. When a role is modified, the new privileges are applied to all users that are assigned to that role. Privilege assignment is not restricted to the privileges defined for the default roles. That is, you can use a custom set of privileges to create a unique role.

A user can be assigned to one or more roles. A user that is assigned to multiple roles has the combined privileges of all assigned roles. For example, if Role 1 has storage-related privileges, and Role 2 has server-related privileges, then users who are assigned to both Role 1 and Role 2 have both storage- and server-related privileges.

All roles include read access to all configurations on the system. The difference between the read-only role and other roles is that a user who is assigned to only the read-only role cannot modify the system state. A user that is assigned to another role can modify the system state in assigned area or areas specific to that user.

Predefined Access Control Elements

- Special privileges
 - Admin—grants all actions
 - AAA—allows RBAC administration
- Admin
 - Default superuser account with full privileges (role = administrator)
 - Cannot be deleted or modified
 - Must set password upon initial setup

Role	Read/Write Access	Read-Only Access
Administrator	Entire system	Entire system
Network administrator	Switch infrastructure and network security operations	Rest of the system
AAA administrator	Users, roles, AAA configuration	Rest of the system
Server administrator	Server-related operations (manage logical servers)	Rest of the system
Storage administrator	Storage operations	Rest of the system
Read-only		Entire system
Operations	System logs (such as syslog) and faults	Rest of the system

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-10

There is a predefined set of privileges within the Cisco UCS solution. The following two privileges grant certain special actions:

- **Admin:** A user-assigned role with this privilege can perform any action in Cisco UCS within any organization.
- **Authentication, authorization, and accounting (AAA):** A user-assigned role with this privilege can perform administration of the RBAC features.

Apart from predefined privileges, the system also includes predefined roles:

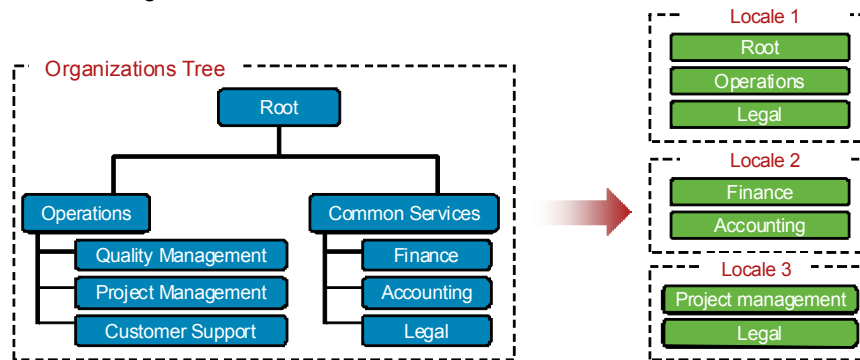
- **Administrator:** Enables complete read/write access to the entire system.
- **Network administrator:** Enables read/write access to switch infrastructure and network security operations. Enables read-only access to the rest of the system.
- **AAA administrator:** Enables read/write access to users, roles, and AAA configuration. Enables read-only access to the rest of the system.
- **Server administrator:** Enables read/write access to server-related operations. A user with this role has privileges to manage logical servers. Enables read-only access to the rest of the system.
- **Storage administrator:** Enables read/write access to storage operations. Enables read-only access to the rest of the system.
- **Read-only:** Enables read-only access to system configuration with no privileges to modify the system state.
- **Operations:** Enables read/write access to system logs, including the syslog servers, and faults. Enables read-only access to the rest of the system.

Note The default administrator account is assigned this role by default, and this assignment cannot be changed.

Note If a role is deleted after it has been assigned to users, all users with that role are assigned the default read-only role.

Locale

- Collection of organizations
- Organizations do not have to be related in hierarchy
- Used with roles
- Assigned to user



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-11


Locale defines the organizations (domains) that a user is allowed to access.

A user can be assigned one or more locales. Each locale defines one or more organizations (domains) that the user is allowed to access, and access is limited to the organizations that are specified in the locale.

An exception to this rule is a locale without any organizations, which gives unrestricted access to system resources in all organizations.

Using RBAC

- User is assigned role-based privileges upon login.
- Privilege sets for all roles are applied to the user-authorized organizations.



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-12

Users are not directly assigned privileges; rather, management of individual user privileges is simply a matter of assigning the appropriate roles and locales.

A user is granted write access to desired system resources only if the assigned role grants the access privileges and the assigned locale allows access.


For example, a user with the `srvRole` role in the Operations organization could update server configurations in the Operations organization, but that user would not be able to update server configurations in the Common Services organization unless the locale for the `srvRole` role also included that organization.

RBAC and Organizations

- Can be used independently

RBAC Without Organizations

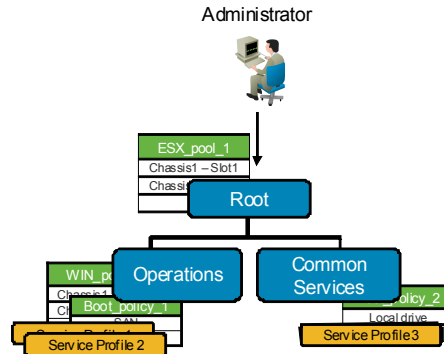
- Everything located in root organization
- Administration can be delegated



© 2011 Cisco Systems, Inc. All rights reserved.

Organizations Without RBAC

- Enforces structured management hierarchy only
- User administrator does everything



DCUCD v4.0-5-13

RBAC and organizations can be used independently—it is not necessary to use RBAC if you are managing a hierarchical set of organizations, and vice versa.

Using RBAC Without Organizations

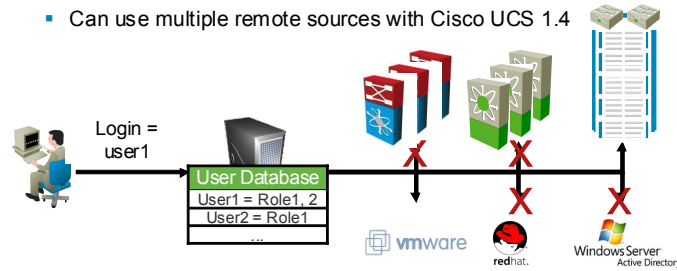
In cases where RBAC is used without organizations—in other words, everything is in the root organization—the advantages of the RBAC style of delegated management can still be used. Users can be segregated based on their responsibilities. There can be users that have control over the logical server structure and other users that have control over the border (LAN and SAN) configuration, for example.

Using Organizations Without RBAC

In cases where organizations are used without RBAC, the organization structural management hierarchy can be utilized to better manage large amounts of equipment and large numbers of logical objects. In addition, some of the pool inheritance features can scale the design logic.

User Authentication

- User authentication used in conjunction with RBAC
 - Applies administrative control
- Centralized authentication
 - Simplifies administration of multiple RBAC schemes
 - Cisco UCS, Nexus 7000, MDS 9000, etc.
- Cisco UCS authentication methods
 - Local to Cisco UCS
 - Remote via protocol (LDAP, RADIUS, TACACS+)
 - Can use multiple remote sources with Cisco UCS 1.4



User authentication is used to apply administrative control. It can be applied in conjunction with RBAC. When a user logs into the Cisco UCS Manager, he or she is authenticated against the appropriate database and assigned privileges that are based on his or her roles.

Centralized authentication is typically used to simplify administration of multiple RBAC schemes.

Cisco UCS supports two methods to authenticate user logins:

- Local authentication with a user database deployed locally to the Cisco UCS Manager
- Remote authentication by using one of the following protocols:
 - Lightweight Directory Access Protocol (LDAP)
 - RADIUS
 - TACACS+

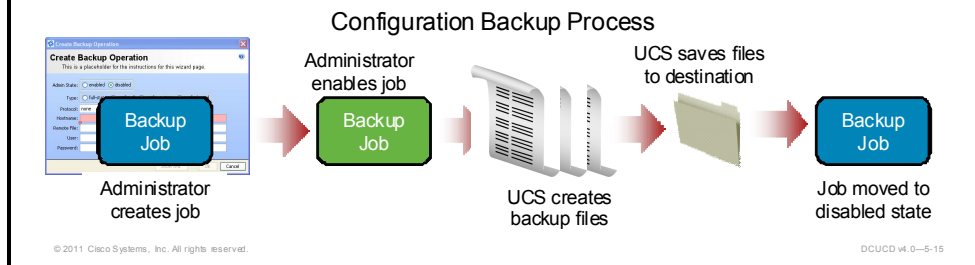
The operation of the user authentication occurs as follows:

- User logs into the Cisco UCS system.
- The Cisco UCS system queries the authentication server for the user approval.
- The authentication server grants or denies user access based on these credentials. If the user is granted access, the authentication server passes the approved role or roles to the Cisco UCS system.
- Cisco UCS grants or denies user access based on the authentication server response.

Cisco UCS version 1.4 introduces several enhancements, among which is the ability to define multiple remote sources for the authentication.

Configuration Management

- Back up or restore Cisco UCS configuration
 - All or partial system configuration snapshot saved in a file
 - Does not back up data on server
- Backup or import job
 - Managed Cisco UCS object—“enabled/disabled” administrative state
 - Run at any time by setting state to “enabled”
 - Protocol options—FTP, TFTP, SCP, SFTP



Cisco UCS allows administrators to manage system configuration—that is, to perform backup and restore operations. Backup and restore operations are performed via Cisco UCS managed objects—backup and import jobs.

The backup or import job has an administrative state. When the state is set to “enabled,” the backup/import job is performed. Once the job ends, the state is changed to “disabled.”

Cisco UCS backup and import jobs can utilize FTP, TFTP, secure copy (SCP), or SFTP protocols to perform the task. Note that you must choose a protocol that is running on the backup server and that you must provide user credentials that are valid on the backup server.

Configuration Backup Options

- Full-state backup
 - System snapshot—complete Cisco UCS database
 - Saved as gzip tarball
 - Cisco UCS upgrades or disaster recovery
- Configuration backup options
 - Config-system—all system configuration settings
 - Config-logical—all logical configuration settings
 - Config-all—all system and logical configuration settings
 - Saved as XML file

Full State

- All configurations
- All runtime states and statuses

System Configuration

- RBAC, AAA configuration
- Usernames, roles, locales

Logical Configuration

- Service profiles, templates
- VLANs, VSANs, pools, policies

All Configuration

- System configuration settings
- Logical configuration settings

© 2011 Cisco Systems, Inc. All rights reserved.

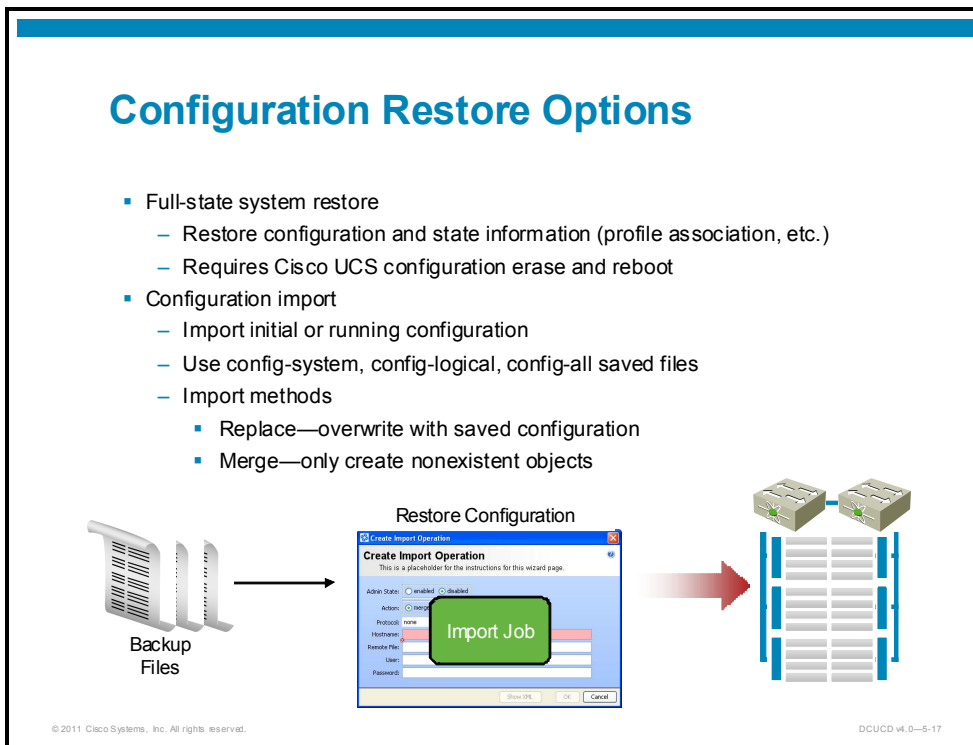
DCUCD v4.0—5-16

Cisco UCS Manager can perform the following types of backups:

- **Full-state:** Includes a snapshot of the entire system. You can use the file that is generated from this backup for disaster recovery if you need to recreate every configuration on a switch or rebuild a switch.
- **All configuration:** Includes all system and logical configuration settings.
- **System configuration:** Includes all system configuration settings such as usernames, roles, and locales.
- **Logical configuration:** Includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies.

Configuration Restore Options

- Full-state system restore
 - Restore configuration and state information (profile association, etc.)
 - Requires Cisco UCS configuration erase and reboot
- Configuration import
 - Import initial or running configuration
 - Use config-system, config-logical, config-all saved files
 - Import methods
 - Replace—overwrite with saved configuration
 - Merge—only create nonexistent objects



Full-State System Restore

An administrator can restore a system configuration from any full-state backup file that was exported from Cisco UCS Manager. The file does not need to have been previously exported from the same Cisco UCS Manager.

The full-state system restore can be performed only in the initial system setup. This function can be used for disaster recovery.

Configuration Import

The import function is available for all configuration, system configuration, and logical configuration files. The import can be performed upon initial configuration or while the system is operating.

An import operation modifies information on the management plane only. Some modifications that are caused by an import operation, such as a change to a virtual network interface card (vNIC) assigned to a server, can cause a server reboot or other operations that disrupt traffic.

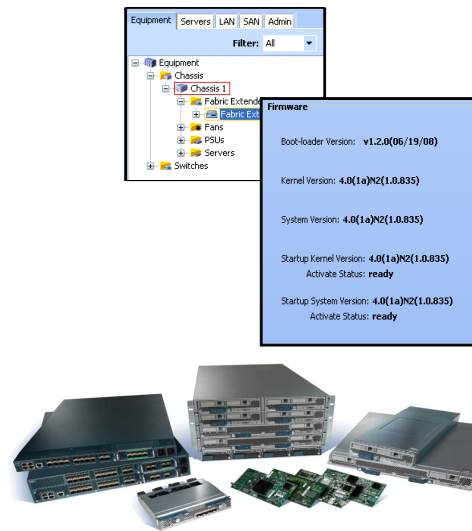
Import Methods

An administrator can use one of the following methods to import and update a system configuration through Cisco UCS Manager:

- The information in the imported configuration file is compared with the existing configuration information. If there are conflicts, the import operation overwrites the information on the Cisco UCS instance with the information in the import configuration file.
- The current configuration information is replaced with the information in the imported configuration file one object at a time.

Firmware Management

- Components
 - Servers—BIOS, disk controller, Cisco IMC
 - NIC and HBA adapters
 - IOMs
 - Fabric interconnects
 - Cisco UCS Manager
- Firmware update packages
 - Downloaded to fabric interconnect
 - Component vs. bundle image



Cisco UCS uses firmware that is obtained from and certified by Cisco to upgrade firmware on the following components:

- Servers, including the BIOS, storage controller, and server baseboard management controller (BMC)
- Adapters, including network interface card (NIC) and host bus adapter (HBA) firmware, and option ROM (where applicable)
- I/O modules
- Fabric interconnects
- Cisco UCS Manager

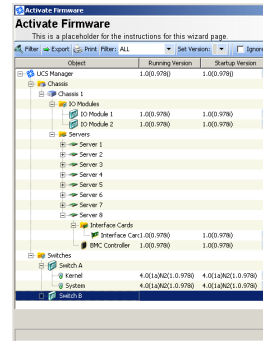
All firmware updates or packages for the Cisco UCS components are in images, which can be either of the following:

- **Component image:** Contains the firmware for one component
- **Bundle image:** A collection of component images

The update packages are downloaded and saved to the Cisco UCS fabric interconnects.

Updating Firmware

- Direct component update
 - Fabric interconnects, extenders, adapters, Cisco UCS Manager
- Server component update via service profile
 - Host firmware package policy
 - Management firmware package policy
- Update via CLI or GUI



Firmware Version	Cisco IMC, Fabric Extenders, Adapters (Active and Backup Firmware Flash Slots)	Fabric Interconnects, Cisco UCS Manager
Running	Active and in use; in active slot	Active and in use
Startup	Use upon next boot; in active slot	Use upon next boot
Backup	Not in use; in backup slot	-

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-19

The Cisco UCS firmware can be updated via any Cisco UCS management interface—for example, the Cisco UCS Manager GUI or command-line interface (CLI). The firmware image can be applied by the following methods:

- Direct update at the Cisco UCS components
- Updates to server components through service profiles that include a host firmware package policy and a management firmware package policy

The firmware update process with UCS provides greater flexibility when compared to other vendor solutions—the individual server firmware (BIOS, Cisco Integrated Management Controller [IMC], and so on) can be either updated manually or in real time via service profiles.

Managing Multiple Cisco UCS Pods

This topic identifies applications and tools that are used to manage multiple Cisco UCS pods.

Cisco UCS Manager Challenges

No sharing across Cisco UCS systems

- Templates
- Available MAC addresses scoped per Cisco UCS Manager
- Available WWN addresses scoped per Cisco UCS Manager
- Available UUIDs scoped per Cisco UCS Manager
- RBAC per Cisco UCS Manager
- Service profiles

Overlapping addresses

- Must control manually

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-21

Cisco UCS Manager is limited to managing a single Cisco UCS pod. This means that the configuration and parameters are not unique between the different pods. Templates, MAC addresses, WWN addresses, UUID identifiers, and even access control is applied singly per pod.

When multiple pods are deployed, the architect and the implementation engineer must manually ensure that there is no address overlapping. Otherwise, problems will occur.

Using XML API

Management flexibility and control

- Can be used to implement proper multipod management

Implement sharing across Cisco UCS systems

- Template management
- Global management of MAC, WWN, and UUID across Cisco UCS systems
- Global RBAC
- Audit and logging of operational activities

© 2011 Cisco Systems, Inc. All rights reserved.

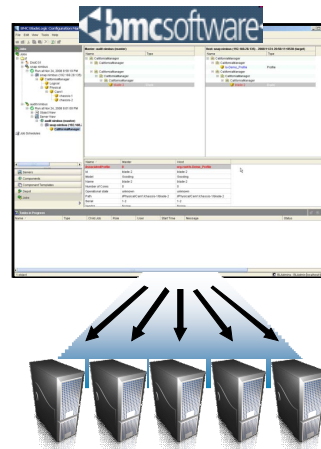
DCUCD v4.0-5-22

The answer to managing multiple UCS pods in the proper manner is to use the Cisco UCS XML application programming interface (API). This is a powerful and flexible tool which can be used to implement multipod environments, so that configuration is shared between the pods and overlapping verification is performed automatically.

The XML API is the interface that ecosystem partner management solutions use to implement and integrate Cisco UCS pod management into their tools. Such tools are BMC BladeLogic and EMC Ionix.

BMC BladeLogic Overview

- End-to-end provisioning tool
- Addresses multiple management aspects
 - Integrated Cisco UCS management
 - Operating system provisioning
 - Installation
 - Boot
 - Application installation
 - Patch installation
 - Configuration and compliance management
 - Business service deployment orchestration



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-23

BMC BladeLogic is a purpose-built automation solution for the Cisco Unified Computing System. It exploits the Cisco UCS XML API to provide automation intelligence.

BMC BladeLogic integrates and supports Cisco UCS management by leveraging unique Cisco UCS features such as policies and service profiles. The integration allows administrators to manage and deploy servers, and to perform compliance operations on Cisco UCS. BMC BladeLogic enables administrators to perform day-to-day Cisco UCS system management without using native Cisco UCS Manager command-line interface (CLI) or GUI management interfaces.

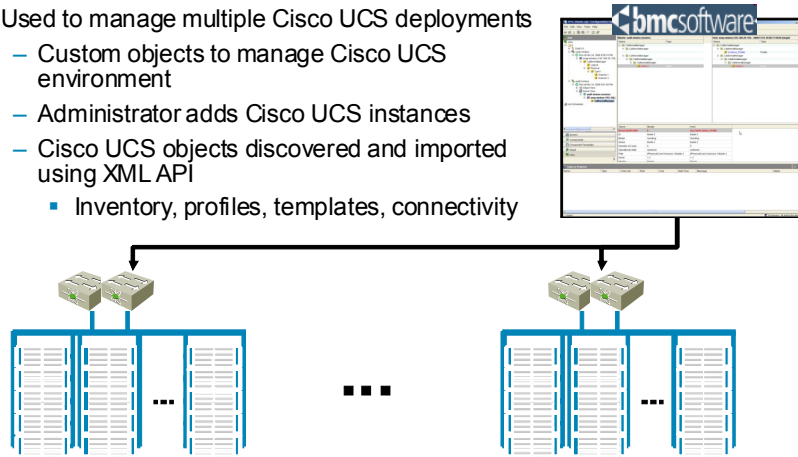
BCM BladeLogic manages across physical Cisco UCS systems, enabling system administrators to define a template once (for example, for an Oracle Database server) and then replicate that template across the farm of Cisco UCS.

BMC BladeLogic also manages unique IDs such as MAC addresses and world wide name addresses for the entire farm, ensuring that an administrator never assigns an address that is already in use.

BMC BladeLogic reaches beyond Cisco UCS into the server software, including the hypervisor layer, guest operating system, and applications.

BMC BladeLogic Integration with Cisco UCS

- Integrated with Cisco UCS
 - Leverages service profiles and policies
- Used to manage multiple Cisco UCS deployments
 - Custom objects to manage Cisco UCS environment
 - Administrator adds Cisco UCS instances
 - Cisco UCS objects discovered and imported using XML API
 - Inventory, profiles, templates, connectivity



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-24

BMC BladeLogic enables administrators in a multiple Cisco UCS deployment scenario to manage all Cisco UCS instances from a single application. BMC BladeLogic uses custom objects to manage the Cisco UCS environment. The discovery and import of Cisco UCS objects is performed using the XML API. The objects that are imported include general inventory, profiles, templates, relationships, topology, connectivity, and policy definitions.

BMC BladeLogic stores and edits the hardware templates that are initially created by the Cisco UCS administrator. Combined with software profiles also stored in BladeLogic, administrators can provision and re-provision an entire technology stack consisting of the hardware resource, virtualization layer, operating system, and business applications.

Cisco UCS Management Actions

- Cisco UCS management actions integrated in BMC BladeLogic

Cisco UCS Aspect	Description
Organization	Create, delete, modify, browse
ID pool	Create, delete, modify, browse
Service profile	Create, delete, associate, disassociate, move, modify, browse, start and stop power actions
KVM	Launch KVM to server
Service profile template	Create, delete, modify, browse
RBAC	Manage the multiple Cisco UCS deployment RBAC
Boot	SAN boot definitions

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-25

BMC BladeLogic creates and maintains its own definitions of service profiles and templates. When an administrator uses the application to associate a service profile with a blade, BladeLogic automatically creates an identical service profile in the local Cisco UCS system. Upon disassociating the service profile, the service profile is also removed from the local Cisco UCS system.

Identity pools (including WWN, MAC, and UUID pools) are also locally created and maintained from within BladeLogic.

Since BladeLogic owns service profile definitions and identities, it can manage multiple Cisco UCS systems and provide service profile mobility across Cisco UCS systems.

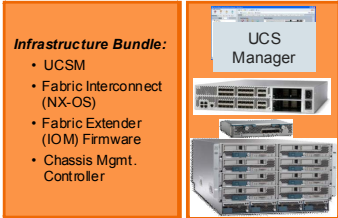
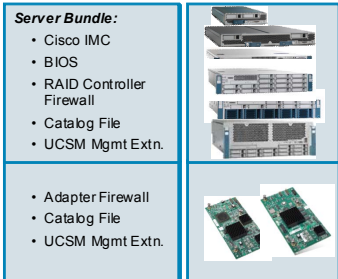
Blade server pools are created and maintained within Cisco UCS and then consumed by BladeLogic. Likewise, the policies are locally maintained and managed from within the Cisco UCS, but are utilized by the service profiles that are defined by BladeLogic.

Cisco UCS 1.4 Enhancements

This topic identifies Cisco UCS 1.4 enhancements from the management aspect.

Cisco UCS Firmware Bundles

- Unbundled UCS software
 - Infrastructure bundle
 - Server bundle

Infrastructure Bundle: <ul style="list-style-type: none">• UCSM• Fabric Interconnect (NX-OS)• Fabric Extender (IOM) Firmware• Chassis Mgmt. Controller 	Server Bundle: <ul style="list-style-type: none">• Cisco IMC• BIOS• RAID Controller Firewall• Catalog File• UCSM Mgmt Extn. 
--	---

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-5-27

The UCS software has now been unbundled—that is, separate infrastructure and server bundles have been created to ease the upgrade effort, with the following features:

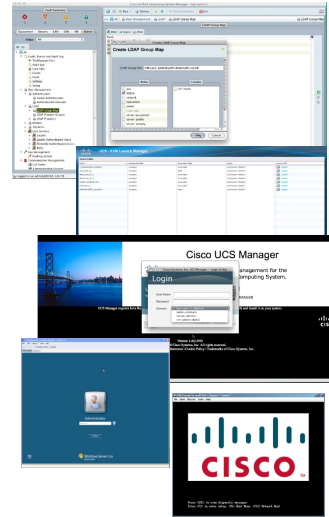
- Server and adapter packs are aimed at new server and adapter hardware deployment.
- Off-cycle introduction of server and adapter packs are synchronized with new hardware releases.
- Next major software releases will incorporate previous modular server and adapter packs.

The benefits are as follows:

- Management is simplified, easier to scale to larger environments.
- Maintenance is simplified (infrastructure and servers have simplified dependencies).
- Code can be updated as needed to take advantage of new software features.
- New server and adapters are introduced rapidly, without the need to wait for the next major software release.

Authentication and Security

- Active Directory group integration
- Multiple authentication options allowed
 - Allow for multiple Active Directory domains
 - LDAP, RADIUS, TACACS used simultaneously
- KVM security enhancement
 - First user receives full KVM access
 - Subsequent users get read-only access
- SNMP MIB support—same information as XML API



SNMP = Simple Network Management Protocol
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—5-28

The Microsoft Active Directory group integration provides administrators with the following features:


- Administrators have complete control over users and their privileges.
- The users and the groups that they belong are created within Active Directory.
- The groups are mapped to roles within Cisco UCS Manager, and the roles themselves are defined within Cisco UCS Manager.

When users log in, they are authenticated against Active Directory, their group membership is looked up, and they are provided with the appropriate role mapping based on their group membership.

The authentication can now be configured such that multiple authentication sources can be used, even mixed ones.

Management Interface Failover

- Cluster IP failover to secondary Cisco UCS 6100 on active management port failure
- KVM, IPMI, SSH IP fails over to the available fabric



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-29

With the previous version of Cisco UCS, the cluster IP address of the management interface did not fail over to the secondary Cisco UCS 6120XP. The recovery process usually meant manual failover using the CLI of the secondary UCS 6100XP.

With Cisco UCS 1.4, the VIP failover to the secondary UCS 6100XP occurs upon active management port failure. However, the managing instance is not stateful, and thus requires relogin for cluster IP and KVM, IPMI, or SSH IP.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco UCS uses organizations to divide system administration into responsible teams.
- RBAC provides granular, per-user permissions configuration.
- A locale is a collection of multiple organizations.
- User authentication configuration can be simplified by centralized authentication servers like RADIUS or TACACS+.
- BMC BladeLogic enables easy management of multiple instances of Cisco UCS.

Designing Cisco UCS Solution Management

Overview

This lesson identifies the appropriate management tools and management hierarchy for Cisco Unified Computing System (UCS) setup.

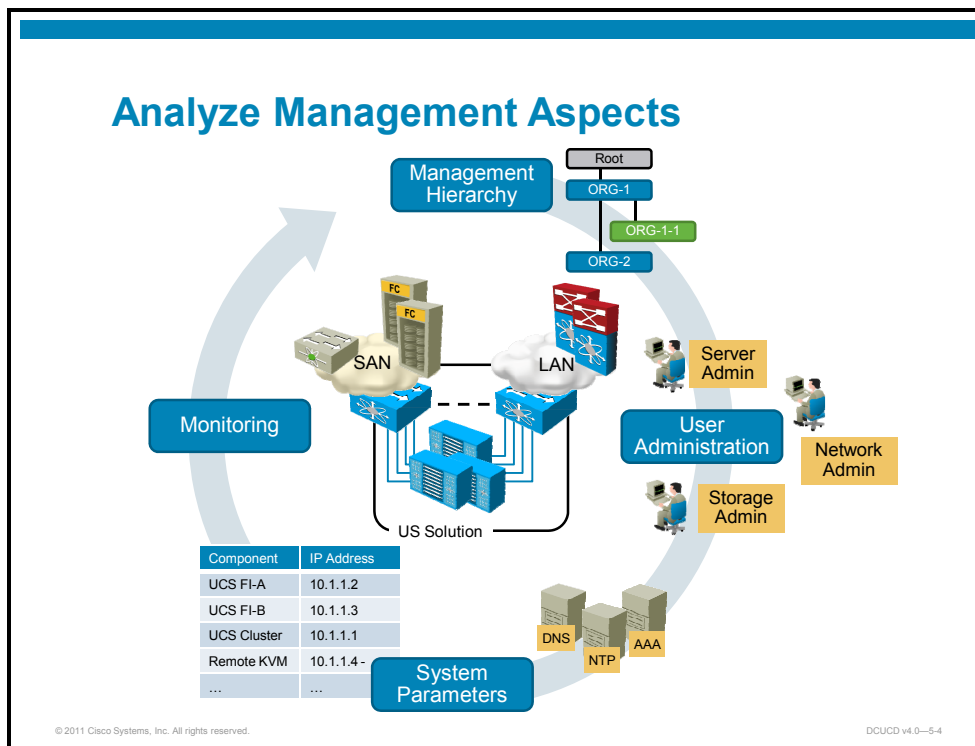
Objectives

Upon completing this lesson, you will be able to describe the management design for the Cisco Data Center Unified Computing solution. This ability includes being able to meet these objectives:

- Identify solution management aspects
- Propose a Cisco UCS management design

Analyzing Requirements

This topic discusses solution management requirements.



The management aspects of Cisco UCS consist of:

- Management hierarchy with organizations
- User management access
- System-related policies and parameters
- Monitoring policies

The first three are necessary for Cisco UCS deployment and need to be agreed with the customer especially when the equipment is to be deployed in the existing environment.

Management Hierarchy and Control

Organizational structure


- Map to departments

Administrative access

- Define roles—server, LAN, SAN
- Locales—group organizations
- Define users
 - Apply role(s)
 - Apply locale(s)

Integrate with external user manager

- RADIUS, TACACS+, LDAP



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-5-5

Organizations

Multitenancy allows you to divide up the large physical infrastructure of an instance into logical entities that are known as organizations. With organizations, you can achieve a logical isolation between organizations without providing a dedicated physical infrastructure for each organization.

The designer must (with the help of the customer) analyze and define the organizational structure, that is, the hierarchy that will be implemented in Cisco UCS.

If you set up a multitenant environment, all organizations are hierarchical. The top-level organization is always root. The policies and pools that you create in root are systemwide and are available to all organizations in the system. However, any policies and pools that are created in other organizations are only available to organizations that are above them in the same hierarchy.

Administrative Access

You can assign unique resources to each tenant through the related organization, in the multitenant environment. These resources can include different policies, pools, and quality of service definitions. You can also implement locales to assign or restrict user privileges and roles by organization, if you do not want all users to have access to all organizations.

Cisco UCS supports two methods to authenticate user logins:

- Local database
- Remote authentication using one of these protocols:
 - Lightweight Directory Access Protocol (LDAP)
 - RADIUS
 - TACACS+

Cisco UCS Naming Convention

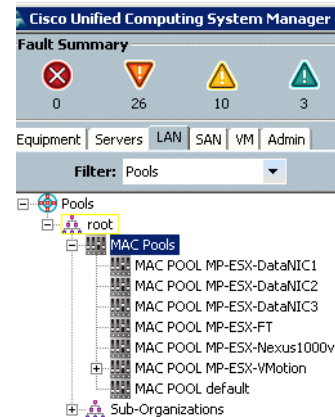
“Names that speak”

Entities

- Service profiles and templates
- MAC, WWN, UUID, server pools
- Organizations, locales, roles
- Management IP pool

Numbering schema

- **Pod ID**—UCS system ID
- **Server type**—Windows, Linux, VMware
- **NIC, HBA ID**—adapter ID



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-6

The first item to define for the deployment is the naming convention for all the elements that will be created during the Cisco UCS configuration. Although this is not necessary, it is recommended to ease not only the deployment but also management and possible troubleshooting.

The naming convention must take care of how the various pools, policies, service profiles, templates, and so on will be named.

When deploying multiple Cisco UCS systems on multiple locations, each system—that is, each Cisco UCS pod—should have a name.

For example, a naming convention can be built using a numbering convention, as follows:

- Each site is given a site ID—a numeric identifier.
- Each pod (Cisco UCS system) is given a pod ID—a numeric identifier.

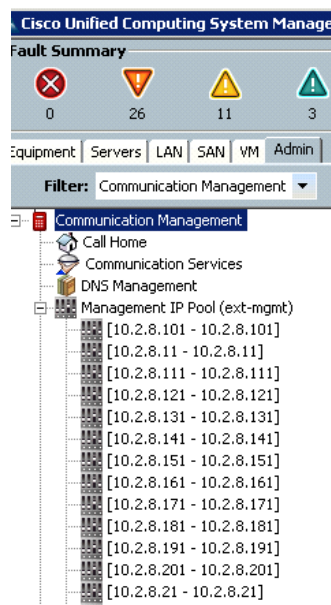
The Cisco UCS pod could be named by using the location name and the pod ID.

Systemwide Parameters

Cluster and fabric interconnect IP addressing

Management IP pool for blade KVM access

Admin user credentials



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-5-7

Management IP Pool

Each Cisco UCS cluster also requires a pool of IP addresses that are used for management access to the individual server blade. An individual Cisco UCS cluster can have up to 320 server blades—each requiring an IP address, a cluster IP address, and the addresses of the fabric interconnect switches A and B. Thus, a maximum of 323 IP addresses are required for a cluster.

System Parameters

Cisco UCS provides several diagnostic tools to aid in troubleshooting and monitoring the environment. These tools include command-line debug statements, syslog, and Simple Network Management Protocol (SNMP). (Call Home is another diagnostic and reporting tool.) Syslog is the mechanism for processes and scripts to write log entries. The callers can fully specify all of the characteristics of the log entries.

Network Time Protocol (NTP) is important from a configuration and maintenance perspective.

Cisco UCS can be accessed via different management interfaces, including SSH, HTTPS, and Telnet. The design must specify which interface should be enabled and used.

Designing Cisco UCS Management

This topic discusses Cisco UCS management design.

Management IP Address Pool Example

UCS IP management pool—dedicated subnet

- **Pod IP address**—one IP in a subnet
- **Fabric interconnect A, B** IP addresses—additional two IPs in a subnet
- **Server blade** IP addresses for **KVM access** (two per blade in 1.4—static and dynamic)
- Default gateway

Description	IP Addresses
IP subnet	192.168.1.0/24
Pod	192.168.1.10
Fabric Interconnect A	192.168.1.11
Fabric Interconnect B	192.168.1.12
Eight server blades	192.168.1.21 – 192.168.1.28
Default gateway	192.168.1.1

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-9

Management Pool Design Example

During the Cisco UCS sizing process, the following quantities and server classes that you will need to deploy have been established:

- 18x PS-class1 server blade requiring redundant LAN and SAN connectivity
- 17x PS-class2 server blades requiring redundant LAN and SAN connectivity, with SAN boot functionality
- 15x VS-class1 server blades requiring redundant LAN and SAN connectivity, with twice the number of LAN interfaces

Since the servers will be deployed in a single data center, and there is no special requirement that the servers need to be identified by the MAC or WWN addresses, you can create a simple naming convention and pools.

For the site ID, you will use 01, in case the implementation grows in the future.

The name of the Cisco UCS cluster will be set to sanjose_01, where 01 stands for the first Cisco UCS cluster in the San Jose data center. This gives the flexibility of introducing new Cisco UCS clusters, if needed, while still being able to identify the cluster from the name.

For the management IP pool, you need 53 IP addresses in all:

- $18+17+15 = 50$ IP addresses for the server blades
- Three IP addresses for the Cisco UCS fabric interconnects and cluster IP

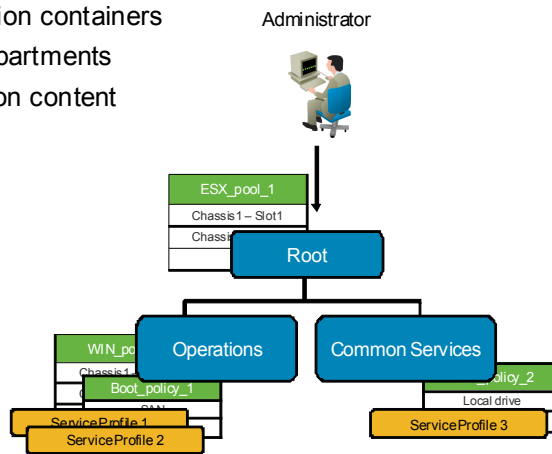
The 192.168.1.0/24 subnet will be used, which is wide enough for future growth. From this management IP address pool, the IP addresses will be given in the following fashion:

- **Cluster IP:** 192.168.1.10
- **Switch A:** 192.168.1.11
- **Switch B:** 192.168.1.12
- **Blade management IP addresses:** 192.168.1.21–192.168.1.28
- **Gateway for the segment:** 192.168.1.1

Design Management Hierarchy

Define organizations

- Create configuration containers
- Set naming to departments
- Define organization content



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-10

Each organization can include at least one suborganization. For example, a company could create organizations for each major division in the company and for subdivisions of those divisions. In this configuration, each organization has access to its local policies and resource pools and to the resource pools in the parent hierarchy.

In the example, the root organization includes two suborganizations: Operations and Common Services. A service profile in the Operations organization is configured to use servers from the OP server pool.

Plan Administrative Access


RBAC

Define administrative roles

- “Superuser” only
- Server, storage, network role
- Define privileges

Define users

- Assign roles
- Internal vs. external database
- Integrate with AAA servers



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-11

The Role-Based Access Control (RBAC) function of the Cisco Unified Computing System allows you to control user access to the actions and resources in the UCS.

A Cisco UCS can consist of up to 40 chassis and as many as 320 blades. This makes the Cisco UCS well suited for multiple administrators in the system. Most companies have Server, Network, and SAN administrators. There is also a fourth distinct role that is the role of Cisco UCS Administrator. These roles already have privileges applied to them, but you can create custom roles with different privilege sets.

The Cisco UCS RBAC allows access to be controlled based on the roles that individual users are assigned. This section describes the elements of the Cisco UCS RBAC model:

- **Roles:** A role is a job function within the context of Locale, along with the authority and responsibility conferred on the user that is assigned to the role.
- **Users:** A user is a person that uses Cisco UCS. Users are assigned to one or more roles.
- **Resources:** A resource is anything in Cisco UCS that is subject to access control.
- **Actions:** An action is any task that a user can perform in Cisco UCS that is subject to access control. An action is performed on a resource.
- **Privileges:** Roles are granted or denied permission to perform an action.
- **Locale:** A locale is a logical object that is created to manage organizations and determine which users have privileges to use the resources in organizations.


Define Strict Management Control

Combine organizations with RBAC

- Assign a user set of organizations with read/write privileges

Define locales

- Specify members—organizations



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-5-12

These roles can also be divided further to have different people responsible for different aspects of a given system. For example, a role for LAN can be to manage quality of service (QoS) parameters. The GUI is organized so that each of the tabs in the navigation pane contains the elements of responsibilities for each admin on the system.

A locale is a logical object that is created to manage organizations and determine which users have privileges to use the resources in a given organization. A locale can be as small as a single organization or as large as several organizations. The locale simply allows you to precisely define the scope of management.

You must have at least one locale, and a locale must contain at least one organization. Once a user is assigned a role in a locale, that user cannot be assigned a role in another locale. Within a given locale, different users can have different roles and different privileges as far as what they can do within the locale.

Strict Management Control

Role

Privileges

- service-profile-config
- service-profile-config-policy
- service-profile-ext-access
- service-profile-network
- service-profile-network-policy
- service-profile-qos
- service-profile-qos-policy
- service-profile-security
- service-profile-security-policy
- service-profile-server
- service-profile-server-policy
- service-profile-storage
- service-profile-storage-policy
- operations
- server-equipment
- server-maintenance
- server-policy
- server-security

User

Properties

Login ID: HR-serveradmin

First Name:

Last Name:

Email:

Phone:

Password:

Confirm Password:

Password Expires:

Roles

- aaa
- admin
- network
- operations
- read-only
- server-equipment
- server-profile
- server-security
- serveradmin
- storage

Locales

- LOC-Finance
- LOC-HR
- LOC-NILHC

Locale and Organization

Name: LOC-NILHC

Organizations

Filter | Export | Print

Name	Org
NILHC-test	org-root/org-NILHC-test

© 2011 Cisco Systems, Inc. All rights reserved.
DCUCD v4.0-5-13

The figure shows an example of strict management access control, with a custom role created for the server admin that is applied to a new user, which can exercise the privileges only in the assigned locale LOC-NILHC.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Management hierarchy with organizations allows more control over how Cisco UCS is configured and how resources are utilized.
- Access control can combine RBAC, which enforces proper user authentication, with locales that allow even more granular management of user access control.

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- Organizations are used to implement hierarchy in Cisco UCS.
- Cisco UCS users can have multiple roles.
- Locales are used to assign read/write management access to Cisco UCS users.
- The config-all backup option is used to save and export all system and configuration Cisco UCS settings in an XML file.
- Multiple Cisco UCS clusters can be managed from a single management instance using the BMC BladeLogic application.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which management application is used to manage the Cisco Nexus 7000? (Source: Examining Cisco UCS Solution Management)
- A) Cisco UCS Manager
 - B) Fabric Manager
 - C) Cisco DCNM
 - D) Navisphere
- Q2) Which management option is used to define hierarchical management structure in Cisco UCS? (Source: Designing Cisco UCS Solution Management)
- A) organizations
 - B) roles
 - C) privileges
 - D) locales
- Q3) Which statement is true if RBAC without organizations is used? (Source: Designing Cisco UCS Solution Management)
- A) Administration can be delegated.
 - B) Structured management hierarchy is enforced.
 - C) Only the administrator user is defined.
 - D) The root organization must be deleted.
- Q4) Which backup option is used to store Cisco UCS service profile settings? (Source: Designing Cisco UCS Solution Management)
- A) full-state
 - B) config-system
 - C) config-logical
 - D) config-all

Module Self-Check Answer Key

- Q1) C
- Q2) A
- Q3) A
- Q4) C

Design Advanced Server Deployment

Overview

This module evaluates how to propose a design plan for various server deployment options.

Objectives

Upon completing this module, you will be able to describe the advanced Cisco Unified Computing System (UCS) server deployment model. This includes the ability to meet these objectives:

- Evaluate Cisco UCS deployment with Microsoft Hyper-V R2
- Evaluate Cisco UCS integration with VMware vSphere
- Evaluate Cisco UCS integration with VMware vSphere and Nexus 1000V

Evaluating Cisco UCS Deployment with Microsoft Hyper-V

Overview

This lesson evaluates Cisco UCS deployment with Microsoft Hyper-V R2.

Objectives

Upon completing this lesson, you will be able to identify where and how Cisco UCS can be used. This ability includes being able to meet these objectives:

- Evaluate Microsoft Hyper-V R2 deployment requirements
- Propose a design for deploying Microsoft Hyper-V R2

Assessing Microsoft Hyper-V R2 Requirements

This topic describes Microsoft Hyper-V R2 deployment requirements.


Microsoft Hyper-V R2 Deployment

Operating systems

- Microsoft Windows 2008 R2
- Hyper-V role

General requirements

- Jumbo frames
- BIOS
 - Hardware DEP bit enabled
 - Intel VT bit enabled
- Minimum 2 GB RAM
- Live migration support
 - Windows failover clustering feature
 - CSVs for VHD storage



The diagram illustrates the deployment of Microsoft Hyper-V R2. It features a central server rack labeled 'Windows Server 2008 Standard'. Above the server is a computer monitor icon. A red box labeled 'Microsoft System Center VMM' is positioned at the top. A large grey circular arrow surrounds the server rack, indicating a cycle or process.

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6-4

Overview

Microsoft Hyper-V is offered as a server role that is packaged into the Windows Server 2008 R2 installation or as a standalone server. In either case, it is a hypervisor-based virtualization technology for x64 versions of Windows Server 2008. The hypervisor is a processor-specific virtualization platform that allows multiple isolated operating systems to share a single hardware platform. In order to run Hyper-V virtualization, the following system requirements must be met:

- An x86-64-capable processor running an x64 version of Windows Server 2008 Standard, Windows Server 2008 Enterprise, or Windows Server 2008 Datacenter.
- Hardware-assisted virtualization, which is available in processors that include a virtualization option—specifically Intel VT, which is an extension of the x86 architecture. The processor extension allows a virtual machine (VM) hypervisor to run an unmodified operating system without incurring significant performance penalties within operating system emulation.
- A CPU that is compatible with the no-execute (NX) bit must be available, and the Hardware Data Execution Prevention (DEP) bit must be enabled in the BIOS. For the Cisco UCS system, these are offered and enabled by default.
- At least 2 GB of memory must be available, and more memory may be needed based on the virtual operating system and application requirements. The standalone Hyper-V Server does not require an existing installation of Windows Server 2008, and minimum requirements are 1 GB of memory and 2 GB of disk space.

Hyper-V isolates operating systems that are running on the virtual machines from each other through partitioning or logical isolation by the hypervisor. Each hypervisor instance has at least one parent partition that runs Windows Server 2008. The parent partition houses the virtualization stack, which has direct access to hardware devices such as network interface cards (NICs) and is responsible for creating the child partitions that host the guest operating systems. The parent partition creates these child partitions using the hypercall API, an application programming interface that is exposed by Hyper-V.

Virtual Partitions

A virtualized partition does not have access to the physical processor, nor does it manage its real interrupts. Instead, it has a virtual view of the processor and runs in a guest virtual address space, which (depending on the configuration of the hypervisor) might not necessarily be the entire virtual address space. A hypervisor could choose to expose only a subset of the processors to each partition. The hypervisor intercepts the interrupts to the processor and redirects them to the respective partition using a logical synthetic interrupt controller (SynIC). Hyper-V can hardware accelerate the address translation between various guest virtual address spaces by using an I/O memory management unit (IOMMU), which operates independently of the memory management hardware that is used by the CPU.

Child partitions do not have direct access to hardware resources, but instead have a virtual view of the resources, in terms of virtual devices. Any request to the virtual devices is redirected via the VMBus to the devices in the parent partition, which manage the requests. The VMBus is a logical channel that enables interpartition communication. The response is also redirected via the VMBus. If the devices in the parent partition are also virtual devices, the response is redirected further until it reaches the parent partition, where it gains access to the physical devices.

Parent partitions run a virtualization service provider (VSP), which connects to the VMBus and processes device access requests from child partitions. Child partition virtual devices internally run a virtualization service client (VSC), which redirects the request to VSPs in the parent partition via the VMBus. This entire process is transparent to the guest operating system.

Cluster Shared Volumes

With Windows Server 2008 R2, Hyper-V uses Cluster Shared Volumes (CSV) storage to support live migration of Hyper-V virtual machines from one Hyper-V server to another. CSV enables multiple Windows servers to access SAN storage using a single consistent namespace for all volumes on all hosts.

Multiple hosts can access the same logical unit number (LUN) on SAN storage so that virtual hard disk (VHD) files can be migrated from one hosting Hyper-V server to another. CSV is available as part of the failover clustering feature of Windows Server 2008 R2.

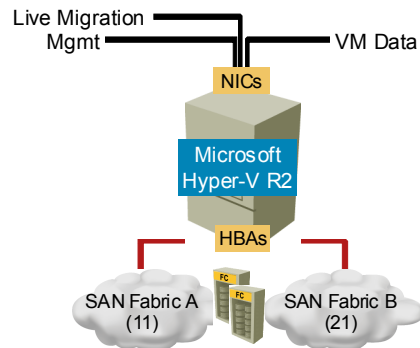
Connectivity Requirements

Separate LAN segments

- **VM Data** (11, 12)—multiple 10 Gb uplinks
- **Management** (99)—redundant 1 Gb uplinks
- **Live Migration** (299)—redundant 10 Gb uplinks

High availability

- LAN connectivity redundancy
- SAN connectivity redundancy



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-5

Live Migration Feature

Hyper-V that is built into Windows Server 2008 R2 supports live migration with the use of CSVs. This allows for an individual VM that is still online and actively supporting user sessions to be moved seamlessly to a different physical server (Hyper-V server) without disrupting user services.

Hyper-V live migration moves running VMs with no impact on VM availability to users. By precopying the memory of the migrating VM to the destination physical host, live migration minimizes the transfer time of the VM. A live migration is deterministic, meaning that the administrator, or script, that initiates the live migration can control which computer will be the destination for the live migration. The guest operating system of the migrating VM is unaware that the migration is happening, so no special configuration for the guest operating system is needed.

During the live migration setup stage, the source physical host creates a TCP connection with the destination physical host. This connection transfers the VM configuration data to the destination physical host. A skeleton VM is set up on the destination physical host, and memory is allocated to the destination VM.

In the second stage of a live migration, the memory that is assigned to the migrating VM is copied over the network to the destination physical host. This memory is referred to as the working set of the migrating VM. A page of memory is 4 KB.

In addition to copying the working set of the migrating VM to the destination physical host, Hyper-V on the source physical host monitors the pages in the working set for that migrating VM. As memory pages are modified by the migrating VM, they are tracked and marked as being modified. The list of modified pages is simply the list of memory pages that the migrating VM has modified after the copy of its working set has begun.

During this phase of the migration, the migrating VM continues to run. Hyper-V iterates the memory copy process several times, so that each time a smaller number of modified pages need to be copied to the destination physical computer. After the working set is copied to the destination physical host, the next stage of the live migration begins.

The next stage is a memory copy process that duplicates the remaining modified memory pages for the migrating VM to the destination physical host. The source physical host transfers the register and device state of the VM to the destination physical host.

During this stage, the network bandwidth that is available between the source and destination physical hosts is critical to the speed of the live migration. It is important to use a 1 Gigabit or faster Ethernet connection. In fact, Microsoft recommends that a 1 Gigabit Ethernet connection is dedicated to the live migration network between cluster nodes to transfer the large number of memory pages that is typical for a virtual machine.

The faster that the source physical host transfers the modified pages from the working set of the migrating VM, the more quickly the live migration completes.

LAN Connectivity Requirements

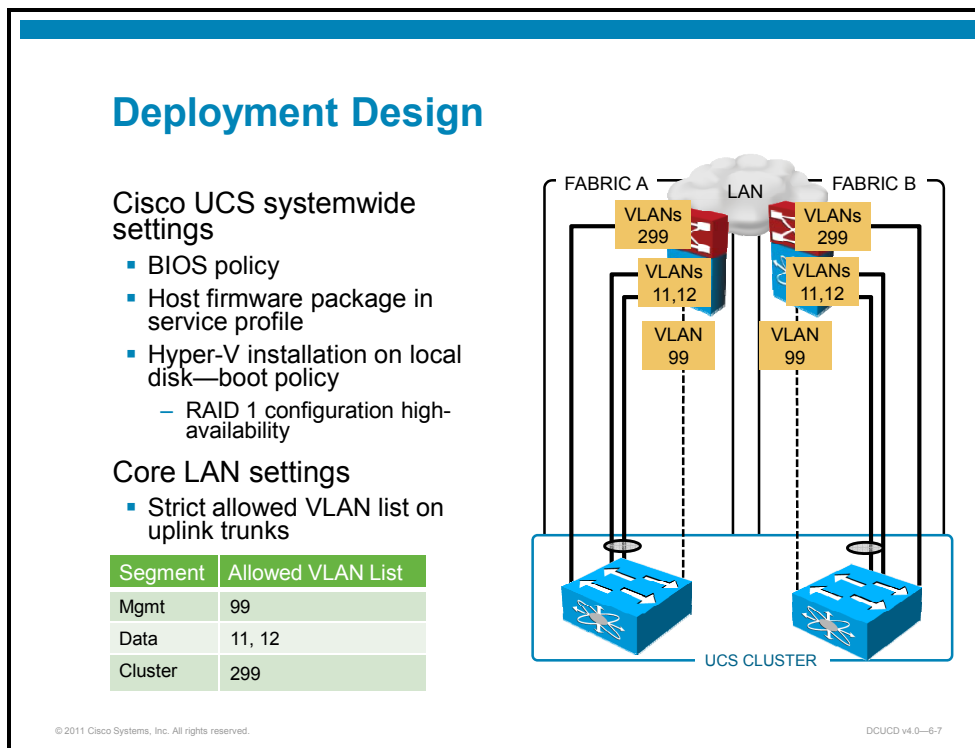
Microsoft recommends assigning dedicated links to various traffic types that are critical to the operation of certain applications and network services.

Some of the traffic types involve communication to the Hyper-V server (parent operating system), and some involve traffic up to the virtual guest operating systems themselves:

- Live migration
- VM management by System Center Operations Manager Virtual Machine Manager and Hyper-V Manager
- VM access

Designing Microsoft Hyper-V R2 Deployment

This topic describes a proposed design for deploying Microsoft Hyper-V R2.



Cisco UCS Policies

To adhere to the Hyper-V R2 deployment requirements, some policies that will be used by service profiles need to be defined:

- **BIOS policy:** Make sure that the DEP and Intel VT bits are set to On.
- **Host firmware package:** Ensure that certified firmware is used with hardware (not governed only by Cisco UCS and Hyper-V, but also disk array type and vendors such as EMC Clariion).
- **Boot policy:** Boot order will be set to CD/DVD proceeding with local disk. Hyper-V will be installed on the local disk.
- **Local disk policy:** To achieve higher availability for Hyper-V the hypervisor, the local disks will be set to Mirror.

LAN Segments

The connectivity requirements mean that more than one segment is used to ensure enough bandwidth to certain features (like live migration), as well as to separate management and control traffic from VM user traffic. Thus, the following segments will be used:

- **Management segment:** This segment connects the Cisco UCS cluster 1 Gigabit uplink in fabric A and B to the core LAN with a single 1 Gigabit Ethernet uplink. On the core LAN side, the trunks will be configured with the strict VLAN allowed list—allowing only VLAN 99 to be used for management.
- **Data segment:** This segment connects the Cisco UCS cluster with a 10 Gigabit EtherChannel uplink (two or more 10 Gigabit Ethernet interfaces) in fabric A and B to the core LAN. On the core LAN side, the allowed VLAN list will allow all the VM data VLANs that are necessary. Currently these are VLANs 11 and 12.
- **Cluster segment:** This segment connects the Cisco UCS cluster with a 10 Gigabit uplink per fabric to core LAN. On the core LAN side, the allowed VLAN list will allow live migration to VLAN 299.

Naming Convention and Pools

Identifier	Description
OStype	Denotes OS types
LANsegment	Denotes LAN segments
SANsegment	Denotes SAN segments
segmentID	Denotes subsegment ID

ID	Description
mmm	LAN segment ID
xx	Host ID
sss	SAN segment ID

LAN	ID
Mgmt	099
VM Data1	011
VM Data 2	012
Live migration	299

Type	Naming Convention
MAC	MP-LANsegment
nWWN	WNP-poolNum
pWWN	WPP-SANsegment
UUID	USP-poolNum
VLAN	LANsegment segmentID
VSAN	SANsegment segmentID
vNICs	NIC-LANsegment
vHBAs	HBA-SANsegment
Pin Group	PG-Segment

SAN	ID	FCoE VLAN
SAN-A	011	1011
SAN-B	021	1021

Pool	Address Template
MAC	00:25:b5:8m:mm:xx
nWWN	20:11:00:25:b5:00:80:xx
pWWN	20:00:00:25:b5:8s:ss:xx
UUID	Prefix = 8000000-0000-00xx Master suffix = 8000-1000000000xx

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6.8


To ease the deployment effort, as well as any troubleshooting that might be needed later, the naming convention will incorporate segment information into the addresses and identifiers. This is indicated in the tables in the figure.

Note that this is only a suggestion of a good design practice, and is not enforced by Cisco UCS or by the Cisco UCS Manager.

LAN Connectivity Design

- EHV mode
- Global VLANs
- Dual physical fabric design
- Three physical segments

Segment	Hyper-V Segments	VLAN
Mgmt	Mgmt	99
Data	10 Gb EtherChannel for VM Data1 and 2	11, 12
Cluster	Live migration	299



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0—6-9

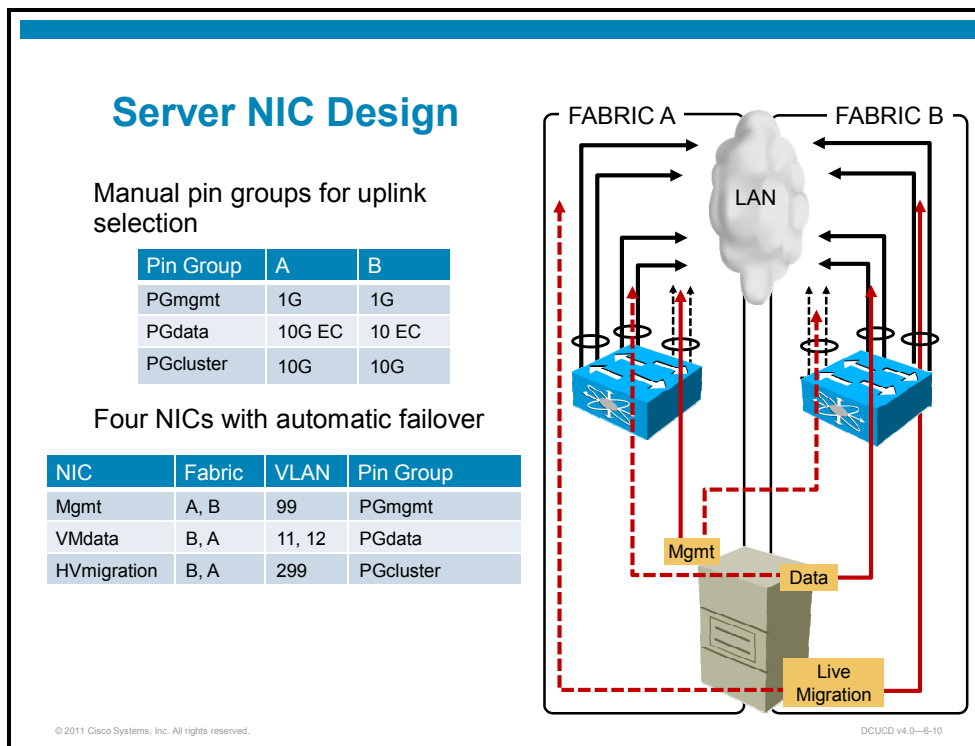
The Hyper-V deployment will be physically connected to a single physical LAN core; that is, the LAN core is not disjointed (formed from two or more disconnected domains).

For this reason—as well as the Cisco UCS level redundancy used to achieve failover between the fabric and manual pinning for the uplink selection—the Cisco UCS cluster will operate in Ethernet host virtualizer (EHV) mode for Ethernet. This also eliminates Spanning Tree Protocol (STP) in the access layer.

LAN Fabric Design

The “physical” LAN fabric will consist of two fabrics—fabric A (left) and fabric B (right). The fabrics are physically separated on the Cisco UCS fabric interconnect level, but can be (and typically will be) on the same devices in the LAN core. Since the VLANs used will have the same IDs in both fabrics, the VLANs will be defined globally in Cisco UCS.

As mentioned for the LAN core, there will be three distinct segments in each fabric to segregate management, control, and VM data traffic. The control traffic from live migration will use the same physical uplink, but will use different VLANs. Note that the design of two physical fabrics allows live migration to use a separate fabric when all the uplinks are operating.



Uplink Selection

Since the LAN connectivity requires three separate segments, administrative pinning will be used to select proper uplinks:

- **PGmgmt pin group:** Uses 1 Gigabit Ethernet uplink in fabric A and 1 Gigabit Ethernet uplink in fabric B.
- **PGdata pin group:** Uses 10 Gigabit EtherChannel (comprising two or more 10 Gigabit Ethernet uplinks) in fabric A and fabric B.
- **PGcluster pin group:** Uses a single 10 Gigabit Ethernet uplink in fabric A and fabric B (this is a different interface from the one that is used in the PGdata pin group).

Hyper-V Host NICs

The individual Hyper-V host will have four NICs to separate traffic on the hypervisor level:


- **Management NIC:** Used for host and VM management by the Microsoft System Center management application.
 - VLAN 99
 - NIC will be configured for failover with fabric A as the primary fabric
 - PGmgmt pin group will be used for uplink selection
- **VMdata NIC:** Used for VM data VLANs 11 and 12 to transfer “user” traffic to and from VM servers. The NIC will be configured as a trunk and both VLANs will be tagged. In the event that additional VM data VLANs are needed, they can easily be added on the service profile level.
 - VLANs 11 and 12
 - NIC will be configured for failover with fabric B as the primary fabric
 - PGdata pin group will be used for uplink selection

- **HVmigration NIC:** Used for Hyper-V live migration segment, where communication and data that are required to transfer a VM from one Hyper-V host to another is exchanged.
 - VLAN 299
 - NIC will be configured for failover with fabric B as the primary fabric to ensure that live migration traffic is using a different uplink (fabric B) than the cluster interconnect (fabric A), even though they are pinned to the same uplinks. (This is only true when all the uplinks are operational.)
 - PGcluster pin group will be used for uplink selection

SAN Connectivity Design

- Cluster SAN design
 - Dual physical fabric design
 - Per-fabric VSAN
 - Two or more uplinks per switch
- Hyper-V SAN design
 - Single HBA per each fabric
 - Operating system-level multipathing

Segment	Hyper-V HBA	VLAN
Fabric A	HBA-A	11
Fabric B	HBA-B	12



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-11

SAN connectivity is important due to shared volumes that need to be accessed by the Hyper-V hosts in order to carry VM migrations, achieve hypervisor high availability, and so on. The VMs are stored in a form of VHD.

The shared volumes reside on SAN attached storage drive arrays and are maintained by the Microsoft Windows clustering functionality, which incorporates the CSV clustering file system.

Since the SAN path selection depends primarily on the parties that are involved in communication—that is, the Hyper-V host and drive array—the high availability is achieved on the hypervisor-disk array level with proper drivers. From the Cisco UCS perspective, the design has to cater to two HBAs that are redundant, that is, connected to separate physical fabrics. Thus, the individual Hyper-V host will be configured with:

- **HBA-A HBA:** Connected to SAN fabric A, which corresponds to VSAN 11
- **HBA-B HBA:** Connected to SAN fabric B, which corresponds to VSAN 12

Apart from the host-level redundancy, Cisco UCS will also incorporate redundancy with multiple Fibre Channel uplinks (two or more) from individual fabric interconnects. This will also provide more throughput for the storage traffic between the Hyper-V hosts and drive array.

The design will allow the system to dynamically pin HBAs to different uplinks in order to share the load between the uplinks. This not only simplifies design and deployment, but also ensures that in a case of individual Fibre Channel uplink failure, Cisco UCS autonomously repins HBAs from the failed uplink to a new uplink, thus achieving better high availability for SAN.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Hyper-V is a role of Microsoft Windows 2008 R2.
- High availability requires Windows failover cluster feature and access to CSV.
- Network connectivity redundancy is best achieved with LAN failover capability at the Cisco UCS level.

Evaluating Cisco UCS Integration with VMware vSphere

Overview

This lesson evaluates Cisco Unified Computing System (UCS) integration with VMware vSphere.

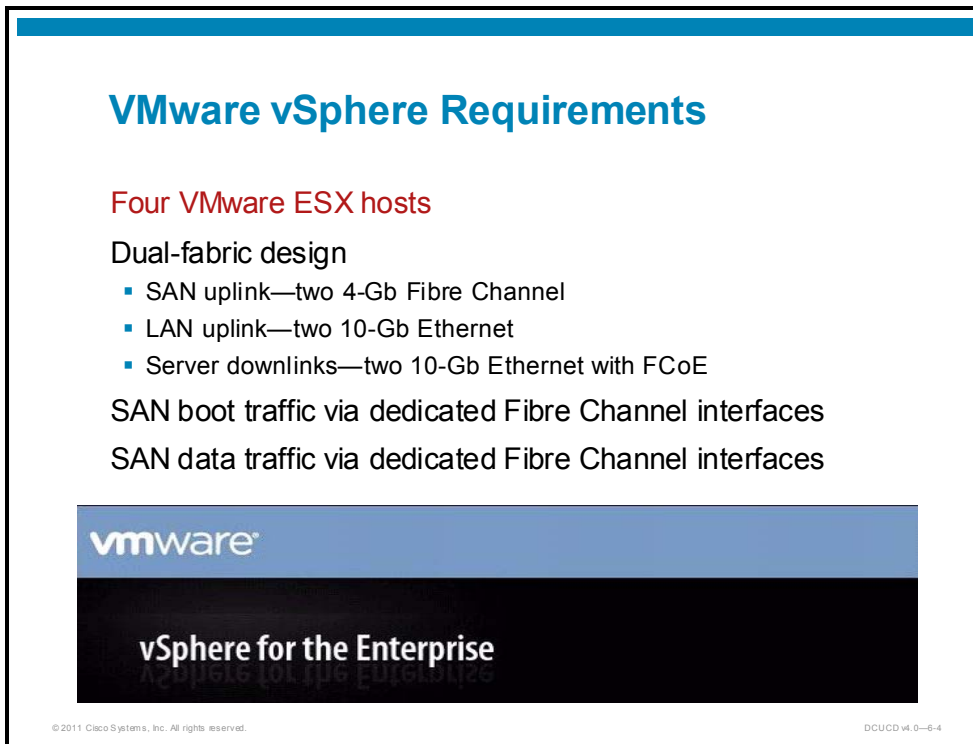
Objectives

Upon completing this lesson, you will be able to deploy VMware vSphere on Cisco UCS. This includes the ability to meet these objectives:

- Evaluate VMware vSphere deployment requirements
- Propose a design for deploying VMware vSphere

Assessing VMware vSphere Requirements

This topic describes examples of VMware vSphere deployment requirements.



The slide features a blue header bar at the top. Below it, the title "VMware vSphere Requirements" is displayed in blue. The main content is in red and black text, listing requirements for four ESX hosts, including dual-fabric design with SAN and LAN uplinks and server downlinks. It also specifies SAN boot and data traffic paths. At the bottom, there is a VMware logo and the text "vSphere for the Enterprise". Small copyright and version information are visible at the very bottom of the slide.

VMware vSphere Requirements

Four VMware ESX hosts

Dual-fabric design

- SAN uplink—two 4-Gb Fibre Channel
- LAN uplink—two 10-Gb Ethernet
- Server downlinks—two 10-Gb Ethernet with FCoE

SAN boot traffic via dedicated Fibre Channel interfaces

SAN data traffic via dedicated Fibre Channel interfaces

vmware

vSphere for the Enterprise

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6.4

In this example, the VMware vSphere deployment is performed with four ESX hosts.

LAN Connectivity Requirements

The physical connectivity requirements dictate that more than two vSphere hosts should be connected to two separate fabrics for each LAN segment that requires redundancy.

Thus, there will be two separate physical LAN fabrics on the Cisco UCS cluster level, which converge in the LAN core on the same set of devices. It has been determined that two 10 Gigabit Ethernet uplinks have enough throughput for all the traffic.

SAN Connectivity Requirements

Likewise, the SAN connectivity should be redundant for each segment or host bus adapter (HBA) that connects vSphere hosts to the SAN fabric. Further, you must use separate segments for accessing the boot and data partitions of the vSphere hosts and environment.

Thus, from the Cisco UCS cluster perspective, the two separate physical SAN fabrics will be used, connected to upstream Fibre Channel switches with the N_Port ID Virtualization (NPIV) feature enabled. Boot and data traffic will be separated through zoning in the SAN and logical unit number (LUN) masking on the drive array side.

ESX Host Requirements

Half-width blades

Memory = minimum 4 GB at 1067 MHz

Processor = minimum 4-core at 2.266 GHz

Boot from SAN

Local disk = striped (RAID 0) for local datastore

Adapter firmware = 1.3(1c)




© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-6-5

Each ESX host should be deployed using hardware with these minimum characteristics:

- **Blade form factor:** Half-width blade
- **Memory:** At least 4 GB at 1067 MHz
- **CPU:** At least 4-core at 2.266 GHz
- **Boot policy:** ESX should be booted from SAN to support fast blade replacement and complete stateless configuration
- **Local disk:** Used by the local datastore to store ISO images and nonmovable VMs (like Cisco Visual Switch Manager [VSM] for Cisco Nexus 1000V)
- **Adapter firmware:** Cisco UCS Manager version 1.3(1c) should be used since it is certified to be used with the selected drive array

ESX Host Connectivity Requirements



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-6

The individual vSphere host has the following connectivity requirements:

- **Management segment:** This segment connects the host to a management segment that is used for management, configuration, and state maintenance, as well as for communication with the vSphere vCenter server. For this example, the default VLAN 1 will be used.
- **ESX vMotion segment:** This segment connects the vSphere hosts in a cluster to enable migration of VMs from one host to another, as well as to enable advanced VMware vSphere high availability and resource management functionalities like High Availability (HA), Fault Tolerance (FT), Distributed Resource Scheduler (DRS), and Distributed Power Management (DPM). For this example, VLAN 997 will be used.
- **VM data segment:** This segment is used for communication between VMs that reside in the private sector of the data center. VLANs 11 and 12 have been currently identified as required, but others can be added later. The segment thus must be configured with trunking capability.
- **VM DMZ segment:** This segment is used for communication between VMs that reside in less trusted networks. VLANs 21 and 22 have been currently identified as required, but others can be added later. Like the private VM data segment, this segment must be configured with trunking capability.

ESX Host Identity Requirements

Four LAN segments = four vNICs

Four SAN segments = four vHBAs

UUID

Identity assignment

- Derived vs. manual vs. **automatic from pool**
- To abstract hardware from personality



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-6-7

Due to the connectivity requirements, the vSphere host must be equipped with the following:

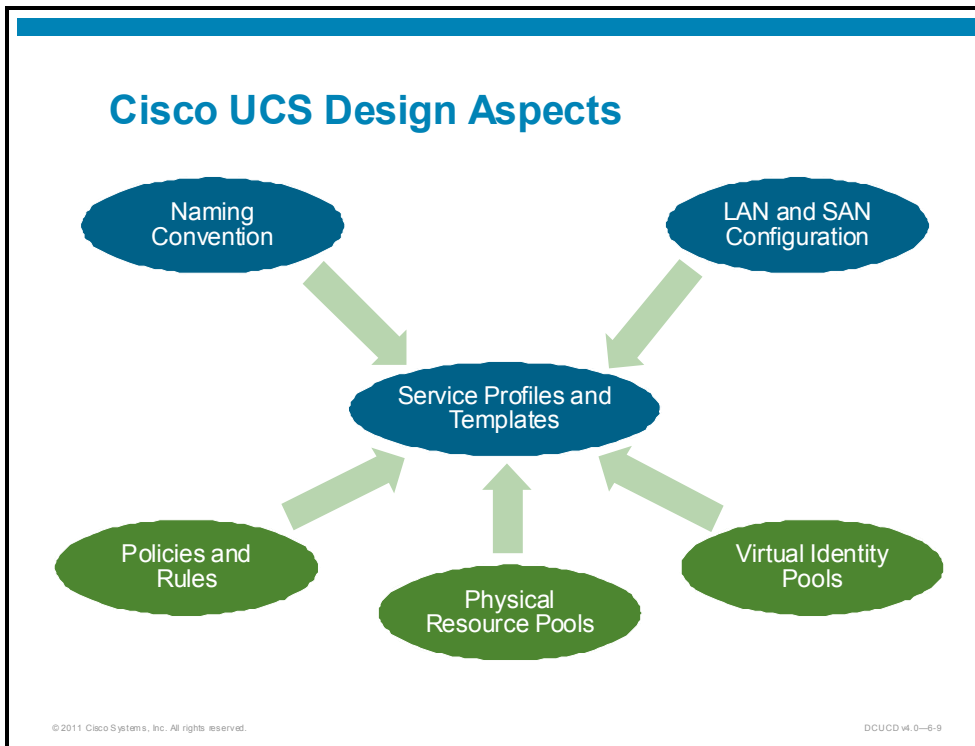
- Four Ethernet NICs to connect to different LAN segments
- Four Fibre Channel HBAs to connect to different SAN segments

The vSphere host also requires the following:

- UUID identifier per vSphere host
- Automatic server assignment from pool, based on server hardware requirements

Designing VMware vSphere Deployment

This topic describes an example design for deploying VMware vSphere.



VMware deployment on Cisco UCS has the aspects that are listed on the figure. It is not necessary to follow such a scheme, but from a long-term perspective, the suggested design provides order that is necessary to control growth and track the environment.

Naming Convention and Segments

Identifier	Description	Type	Naming
ESX	Denotes ESX hypervisor	MAC	MP-ESX <i>LANseg</i>
SubsegID	Denotes subsegment ID (for example, VLAN on physical segment)	nWWN	WNP-ESX <i>poolID</i>
LANseg	LAN segment	pWWN	WPP-ESX <i>SANseg</i>
SANseg	SAN segment	UUID	USP-ESX <i>poolNum</i>
		VLAN	<i>LANsegSubsegID</i>
		VSAN	<i>SANsegSubsegID</i>
		vNICs	NIC- <i>LANseg</i>
		vHBAs	HBA- <i>SANseg</i>
		Server pool	SRVP-ESX <i>poolID</i>
		Server pool policy	SRVPP-ESX <i>poolID</i>
		Server qualification	SPPQ-ESX <i>poolID</i>
		Host firewall package	HFP-ESX <i>poolID</i>

LAN Segment	ID	SAN Segment
Mgmt	001	SAN-FabricA
VMdata	010	SAN-FabricB
VMdmz	020	
ESXvMotion	997	

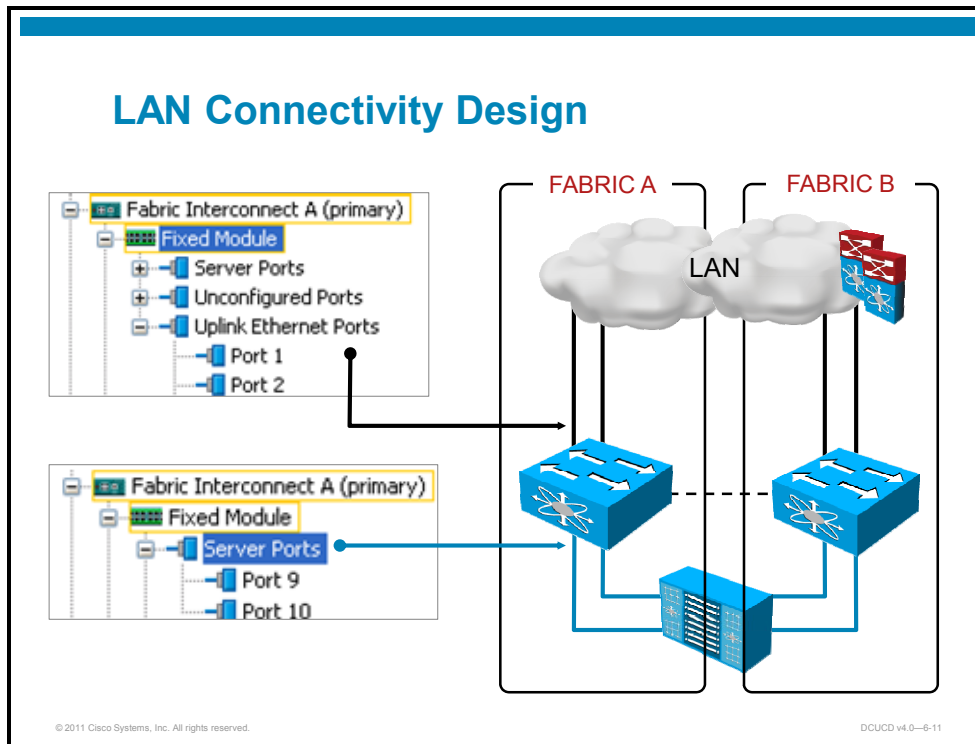
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-10

To ease the deployment effort, as well as any troubleshooting that might be needed later, the naming convention will incorporate segment information into the addresses and identifiers. This is indicated in the tables in the figure.

Note that this is only a suggestion of a good design practice, and is not enforced by Cisco UCS or by the Cisco UCS Manager.

LAN Connectivity Design



The physical LAN connectivity uses multiple 10 Gigabit Ethernet uplinks in each fabric—in this example, two per fabric. There are 10 Gigabit Ethernet uplinks for connectivity to the LAN core, and 10 Gigabit Ethernet Fibre Channel over Ethernet (FCoE) enabled links to connect the fabric interconnects and chassis.

LAN Segments and VLAN Design

VLAN scheme

- Fabric **global** (same VLAN ID in fabric A and B)
- VLAN IDs—**must be unique** (at least per fabric)

Operational mode

- **EHV** (default)
- Switching mode

VLAN	ID	Fabric
Mgmt	1	A, B
VMdata1	11	A, B
VMdata2	12	A, B
VMdmz1	21	A, B
VMdmz2	22	A, B
ESXvMotion	997	A, B
N1kv-control	999	A, B
N1kv-packet	998	A, B

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-12

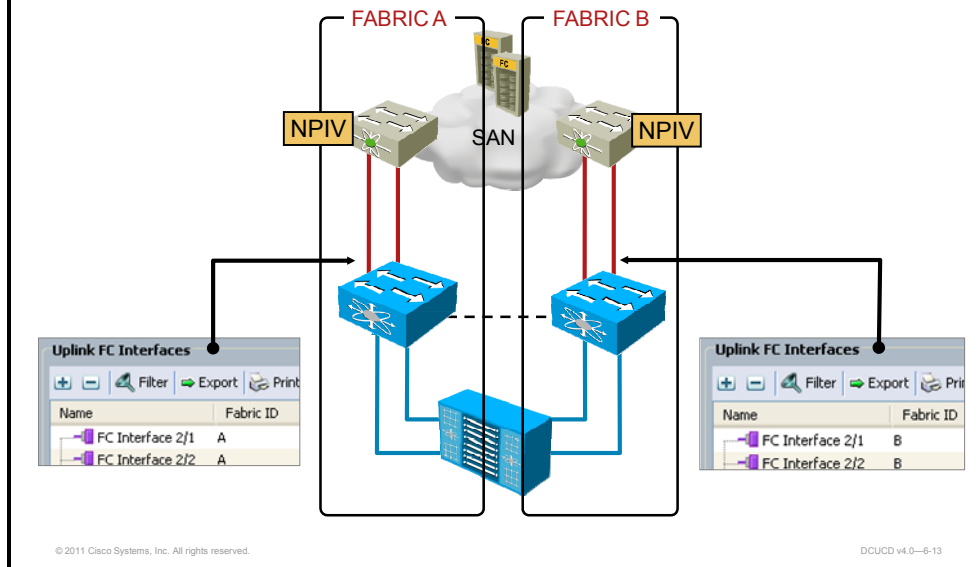
The VMware vSphere deployment will be physically connected to a single physical LAN core; that is, the LAN core is not disjointed (formed from two or more disconnected domains).

For this reason—as well as the Cisco UCS level redundancy used to achieve failover between the fabric and manual pinning for the uplink selection—the Cisco UCS cluster will operate in Ethernet host virtualizer (EHV) mode for Ethernet. This also eliminates Spanning Tree Protocol (STP) in the access layer.

VLAN Design

The VLANs that are used will have the same IDs in both fabrics, so they will be defined globally in Cisco UCS with the VLAN IDs that are specified in the table in the figure.

SAN Physical Connectivity



The physical SAN connectivity uses multiple Fibre Channel uplinks in each fabric—in this example, two per fabric. The Fibre Channel uplinks in fabric A are set to VSAN 10, and the Fibre Channel uplinks in fabric B are set to VSAN 20.

SAN Design

Must be defined before Fibre Channel uplink configuration

VSAN scheme

- Fabric A or B
- VSAN IDs—**must be unique** (at least per fabric)
- FCoE VLAN ID—**must be unique** (at least per fabric)

Operational mode is NPV Edge

- Upstream Fibre Channel switches must be enabled with NPIV

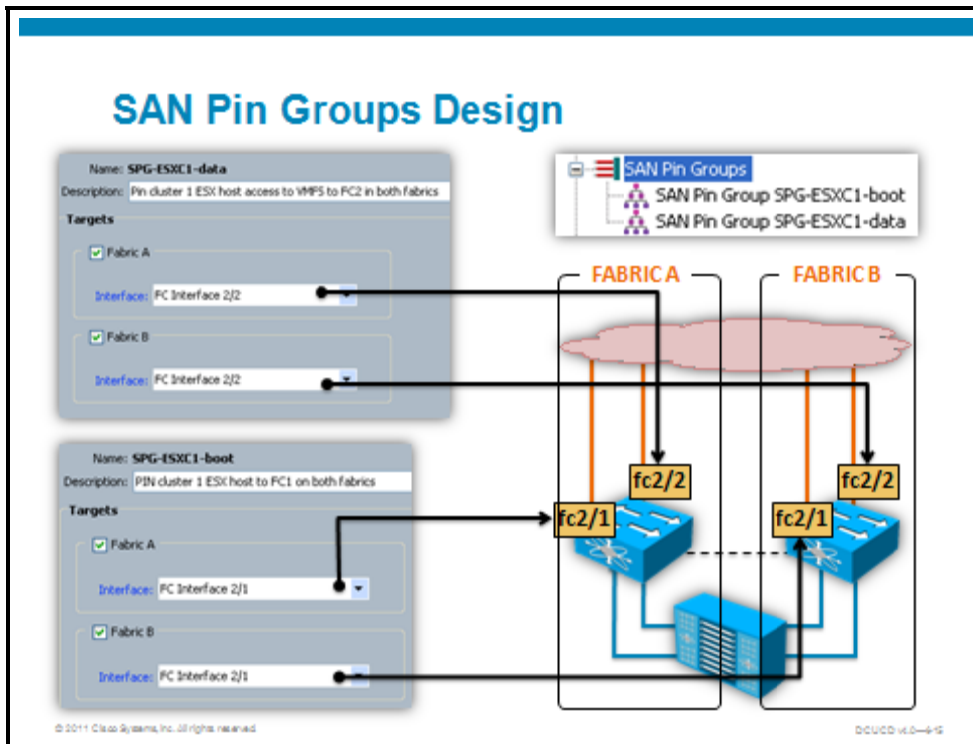
VSAN	VSAN ID	Fabric	FCoE VLAN ID
SAN-FabricA	10	A	1010
SAN-FabricB	20	B	1020

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-14

The SAN design is simple—two separate physical fabrics using the VSAN IDs, as specified in the table. Since FCoE is used, VLAN FCoE IDs have to be specified—also indicated in the table. A good design practice is to add a prefix to the VSAN ID in order to get a FCoE VLAN ID if the same VLAN ID is already in use (for example, adding 10 in front of VSAN to get the FCoE VLAN ID).

In the core SAN fabric, the core Fibre Channel switches must be enabled with the NPIV feature. Also, for the Cisco Multilayer Director Switch (MDS), the Cisco UCS attached interfaces must be part of the same VSAN (that is, in access mode). Thus, from a design perspective, the Cisco UCS interfaces that are attached to the MDS must have the same configuration as the Cisco UCS Fibre Channel interfaces.



In the example, the deployment requirement is also to separate SAN traffic for boot LUN access and data LUN—VMFS access. This is achieved with multiple HBAs on the service profile level, which have manual pinning that defines which uplink will be selected.

In the example, there are two uplinks per fabric, which enables separation of traffic mentioned. Thus, to be able to apply manual traffic engineering, SAN pin groups have to be defined:

- **SPG-ESXC1-data pin group:** Used to select the uplinks for accessing the VMFS datastore of a disk array, and will use interface Fibre Channel 2/2 in both fabrics (that is, on both fabric interconnects).
- **SPG-ESXC1-boot pin group:** Used to select the uplinks for accessing the vSphere host boot LUN of a disk array, and will use interface Fibre Channel 2/1 in both fabrics

Virtual Identity Pools

Pool	Address Template
MAC	00:25:b5:3m:mm:xx
nWWN	20:11:00:25:b5:00:30:xx
pWWN	20:00:00:25:b5:3s:ss:xx
UUID	Prefix = 30010000-0000-0000 Master suffix = 3000-1000000000xx

ID	Description	LAN Segment	ID	SAN Segment	ID
mmm	LAN segment ID	Mgmt	001	Boot SAN-A	019
xx	Host ID	VMdata	010	Boot SAN-B	029
sss	SAN segment ID	VMdmz	020	Data SAN-A	011
ESX OS	3	ESXvMotion	997	Data SAN-B	021
ESX pool num	001				

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-16

To ease the deployment effort and troubleshooting in the future, and to have better control of the environment when it scales, the virtual identity pools with addresses and identifiers will be defined. This includes MAC, UUID, nWWN, and pWWN pools and addresses.

The addresses and identifiers will have segments and IDs encoded to further ease deployment and troubleshooting using the following fields:

- **mmm:** Denotes the LAN segment ID and encodes the VLAN number in the MAC address. For the VM data segment, since there will be a single NIC used, the identifier will be the least common denominator; in other words, the last part of the VLAN ID will be omitted. The same is true for the VM DMZ segment.
- **xx:** Denotes the host ID, and is picked and determined by Cisco UCS itself.
- **sss:** Denotes the SAN segment ID and encodes the VSAN number that is added by the suffix, where 1 means data segment and 9 means boot segment.
- **ESX pool number 001:** To anticipate future addition of vSphere pools, this first ESX pool will be denoted with 001, which will be encoded into the addresses.
- **ESX OS ID 3:** Since Cisco UCS enables deployment of multiple different operating systems and hypervisors, information about the operating system and hypervisor type will also be encoded into the addresses and identifiers. In this case, the vSphere hypervisor will have an ID of 3.

The table at the top of the figure lists the pool types and also specifies a template for creating addresses and identifiers in each pool that have these fields encoded.

MAC Address Pools

MAC		Address
MP-ESXmgmt	First	00:25:b5:30:01:01
	Last	00:25:b5:30:01:04
MP-VMdata	First	00:25:b5:30:10:01
	Last	00:25:b5:30:10:04
MP-VMdmz	First	00:25:b5:30:20:01
	Last	00:25:b5:30:20:04
MP-ESXvMotion	First	00:25:b5:39:97:01
	Last	00:25:b5:39:97:04

LAN Segment	ID
Mgmt	001
VMdata	010
VMdmz	020
ESXvMotion	997

Easy vmnic
Identification →

Device	MAC Address
Cisco Systems Inc VIC Ethernet NIC	
vmnic3	00:25:b5:30:20:02
vmnic2	00:25:b5:39:97:02
vmnic1	00:25:b5:30:10:02
vmnic0	00:25:b5:30:01:03

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-17

MAC address pools are designed for each NIC that will be available to the vSphere host. It has been determined that individual host will have six NICs, thus six MAC pools are required.

- **MP-ESXmgmt pool:** MAC addresses that are used by the NIC that is connected to the management segment. Encodes VLAN 1, which is used for the management segment.
- **MP-VMdata pool:** MAC addresses that are used by the NIC that is connected to the VM data segment, which combines the VMdata1 and VMdata2 subsegments (VLANs 11 and 12). Thus, the encoded field is 10 (the least common denominator).
- **MP-VMdmz pool:** MAC addresses that are used by the NIC that is connected to the VM DMZ segment, which combines the VMdmz1 and VMdmz2 subsegments (VLANs 21 and 22). Like the VM data pool, this one encodes field value 20, which is the least common denominator for the given segments.
- **MP-ESXvMotion pool:** MAC addresses that are used by the NIC that is connected to the vMotion vSphere segment, which is utilized by the VM migration as well as advanced functionalities like high availability, FT, DRS, and DPM.

You can see from the tables in the figure that the MAC addresses in an individual pool are created based on the MAC address template that was specified earlier.

Notice the NIC appearance in the vSphere vCenter configuration—there is no similarity between the name of the vmnic (the vSphere host physical NIC) and the NIC defined by Cisco UCS. However, since the MAC addresses have encoded fields, it is easy to determine which vmnic should be used for which segment. Thus, deployment is easier.

WWN Address Pools

nWWN		Address
WNP-ESXC1	First	20:11:00:25:b5:00:30:01
	Last	20:11:00:25:b5:00:30:04

pWWN		Address
WPP-ESX-BootA	First	20:00:00:25:b5:30:19:01
	Last	20:00:00:25:b5:30:19:04
WPP-ESX-BootB	First	20:00:00:25:b5:30:29:01
	Last	20:00:00:25:b5:30:29:04
WPP-ESX-DataA	First	20:00:00:25:b5:30:11:01
	Last	20:00:00:25:b5:30:11:04
WPP-ESX-DataB	First	20:00:00:25:b5:30:21:01
	Last	20:00:00:25:b5:30:21:04

SAN Segment	ID
Boot SAN-A	019
Boot SAN-B	029
Data SAN-A	011
Data SAN-B	021

Easy vmhba
Identification



Storage Adapters		ESX Host
Device	WWN	
vmhba5	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:19:03	
vmhba6	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:29:03	
vmhba7	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:11:03	
vmhba8	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:21:03	

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-18

nWWN pools are designed for each vSphere host that will be connected to the SAN attached storage. Since there are only four hosts initially, the WNP-ESXC1 nWWN pool is designed with only four nWWN addresses. However, the addresses have an encoded operating system type (3 for vSphere hypervisor) and host ID. The VSAN segment is not encoded since the host has only one nWWN address in the fabric.

The pWWN address pools are designed for each HBA that will be available to the vSphere host. It has been determined that individual host will have four HBAs, so four pWWN pools are required.

- **WPP-ESX-BootA pool:** pWWN addresses that are used by the HBA that is connected to the boot segment in fabric A. Uses field value 19 to encode VSAN 10 and identify the boot HBA.
- **WPP-ESX-BootB pool:** pWWN addresses that are used by the HBA that is connected to the boot segment in fabric B. Uses field value 29 to encode VSAN 20 and identify the boot HBA.
- **WPP-ESX-DataA pool:** pWWN addresses that are used by the HBA that is connected to the data segment in fabric A (the VMFS). Uses field value 11 to encode VSAN 10 and identify the data HBA.
- **WPP-ESX-DataB pool:** pWWN addresses that are used by the HBA that is connected to the data segment in fabric B (the same VMFS as fabric A). Uses field value 21 to encode VSAN 20 and identify the boot HBA.

You can see from the tables in the figure that the pWWN addresses in an individual pool are created based on the pWWN address template that was specified earlier.

Notice the HBA appearance in the vSphere vCenter configuration—there is no similarity between the name of the vmhba (the vSphere host physical HBA) and the HBA defined by Cisco UCS. However, since the pWWN addresses have encoded fields, it is easy to determine which vmhba should be used for which segment. Thus, deployment is easier.

UUID Pool

ID	Description
ESX OS	3
ESX pool num	001

UUID	Prefix		Suffix
USP-ESXC1	30010000-0000-0000	First	3000-1000000000001
		Last	3000-1000000000004



UUID Assignment		
Name	Assigned	Assigned To
3000-0000000000001	yes	org-root/org-NILHC-test/ls-NILHC-ESX4
3000-0000000000002	yes	org-root/org-NILHC-test/ls-NILHC-ESX3
3000-0000000000003	yes	org-root/org-NILHC-test/ls-NILHC-ESX2
3000-0000000000004	yes	org-root/org-NILHC-test/ls-NILHC-ESX1

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-19

The UUID pool and addresses are defined per the UUID template that was defined earlier—they encode the operating system type (3 for vSphere host) and vSphere cluster ID (001 for the first one).

Like the MAC and pWWN addresses, the NICs, and the HBAs, the vCenter configuration is easier to read due to encoded fields.

Server Assignment


Manual vs. **automatic**

Automate server blade assignment

- Gather servers with similar characteristics

Design aspects

1. Define server pool policy qualification
2. Define server pools
3. Define server pool policies



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-20

vSphere hosts in this example have minimum requirements for the server hardware. When the service profile is created, the blade can be assigned manually or automatically from the pool. In this example, since the minimum hardware characteristics are identified, the automatic assignment with policy enforcement can be used—that is, the server blade will be automatically assigned from the pool to the service profile upon creation.

For this purpose, the design will specify the following objects upon which the server blade will be automatically assigned to the service profile upon creation:

- **Server pool policy qualification:** This object specifies the hardware characteristics against which server blades are matched when discovered by the system.
- **Server pool:** This object groups the blades of similar characteristics (in this example).
- **Server pool policy:** This object binds together qualification policy and server pool—that is, it populates the server pool with blades matched against criteria that are defined by the qualification policy.

Server Pool Policy Qualification

Qualification	Value
Adapter	Virtualized Ethernet and FC
Memory	Min 4 GB at 1067 MHz
Processor	Min 4-core at 2.266 GHz
Blade model	N20-B6620-1 (B200-M1)

Properties

Name: **SPPQ-ESXC1**

Description: Qualification criteria for ESX cluster 1 servers

Adapter Qualifications

- virtualized-eth-if
- virtualized-fc-if

Properties for: Memory qualification

General Events

Clock (MHz): 1067

Min Cap (MB): 4096

Width: 64

Properties for: Server Model Qualification

General Events

Model (Regex): N20-B6620-1

Properties for: Processor qualification

General Events

Processor Architecture: Xeon

Min Number of Cores: 4

Min Number of Threads: unspecified

CPU Speed (MHz): 2266.0

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-21

In this example, the qualification policy criteria are the following:

- **Adapter** that must be virtualized with Ethernet and Fibre Channel support
- **Memory** that is at least 4 GB in size and at least at 1067 MHz
- **CPU** with a minimum of 4 cores running at 2.266 GHz
- **Blade size** should be half width.

The name of the qualification policy for the vSphere cluster hosts is **SPPQ-ESXC1**.

Server Pool and Pool Policy

Used only on blade discovery

The diagram illustrates the configuration and discovery process for a server pool. It consists of three main components:

- Pool Policy Properties:** A box titled "Properties" for "SRVPP-ESXC1". It includes:
 - Name: SRVPP-ESXC1
 - Description: Put server blades for ESX cluster 1 into pool
 - Target Pool: Server Pool SRVP-ESXC1
 - Qualification: SPPQ-ESXC1-CPU
- Server Pool Summary:** A box for "SRVP-ESXC1" with:
 - Name: SRVP-ESXC1
 - Description: Server blades for ESX cluster 1
 - Size: 4
 - Assigned: 4
- Server List Table:** A table with columns "Name", "Chassis ID", and "Slot ID". It lists four servers:

Name	Chassis ID	Slot ID
Server 1/4	1	4
Server 1/5	1	5
Server 1/6	1	6
Server 1/7	1	7


Red arrows indicate the flow: a downward arrow from the Pool Policy to the Server Pool, and an L-shaped arrow pointing from the Server List up to the Server Pool.

The server pool that will be used by service profiles for auto blade assignment will be named SRVP-ESXC1, and will be bound to the SRVPP-ESXC1 pool policy that combines it with the qualification policy mentioned earlier.

The design has to include a note that the automatic pool population is done only upon server blade discovery. In other words, if the server blades are already present and the pool, pool policy, and qualification policy are defined later, the implementation engineer should start blade rediscovery for all the blades that are not used. This way the pool is populated.


ESX Host Common Policies

BIOS defaults



Local storage policy—**LDCPol-ESXC1**

- Striped (RAID 0) for local datastore



Apart from pools and related policies, there are also other policies that need to be defined.

The first one is the BIOS policy, or BIOS defaults for the whole system, with which it is ensured that the necessary settings are applied (for example, Intel VT bit).

Second is the local storage policy—the requirement is to use the local disk as a local datastore in VMware and to use the entire space. Thus, the LDCPol-ESXC1 is defined with the Redundant Array of Independent Disks (RAID) stripe disk configuration option.

ESX Host Common Policies (Cont.)

Host firmware package—**HFP-ESXC1**

- Adapter firmware = 1.3(1c)

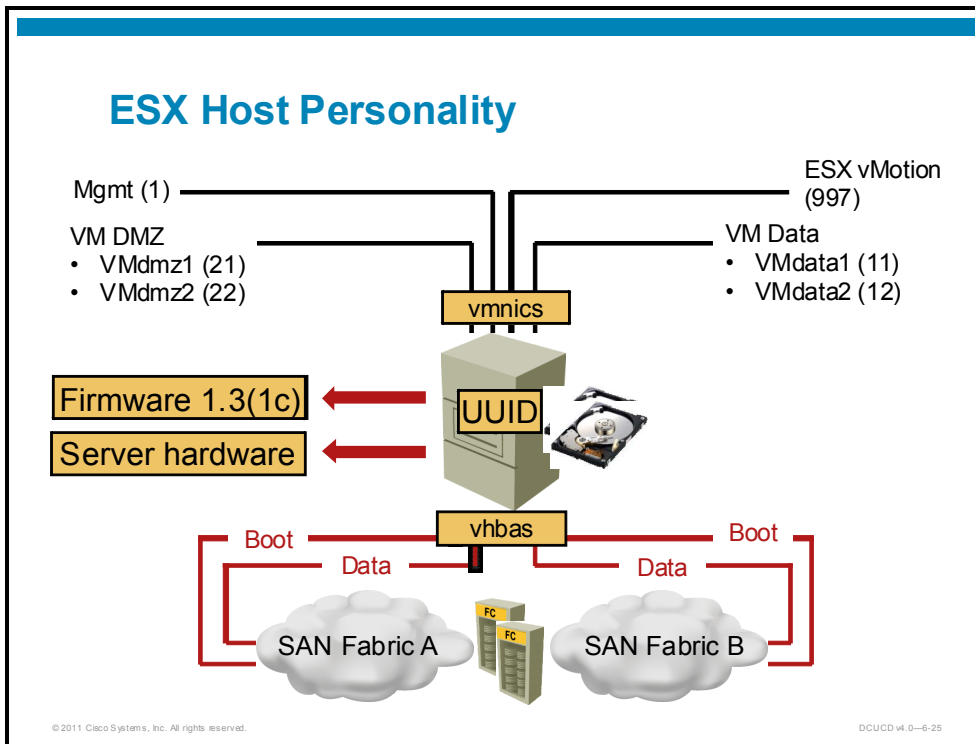
The screenshot displays the configuration page for the host firmware package 'HFP-ESXC1'. It includes a 'General' tab, a 'Name' field with the value 'HFP-ESXC1', and a 'Description' field. Below these are two expandable sections: 'Adapter Firmware Packages' and 'Storage Controller Firmware Package'. To the right, there is a table with columns for Type, Vendor, Model, Presence, and Version. The table contains one entry: Adapter, Cisco Systems Inc, N20-AC0002, present, 1.3(1c). Above the table are buttons for Filter, Export, and Print.

Type	Vendor	Model	Presence	Version
Adapter	Cisco Systems Inc	N20-AC0002	present	1.3(1c)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-24

Likewise, the host firmware package policy named HFP-ESXC1 is defined to ensure that the vSphere host service profiles (and consequently, blades) will use proper firmware for the adapters, specifically for the HBAs. This is important from the drive array vendor certification and support aspect.



Combining all the requirements, the vSphere host personality consists of the following:

- LAN connectivity definition with four NICs
- SAN connectivity definition with four HBAs
- Host firmware package with 1.3(1c) firmware version
- Distinct server hardware (that is, certain physical resources have to be available)
- Two local disks for local datastore
- UUID identifier to distinguish between the vSphere hosts

Service Profiles

Service template

- For fast, massive server deployment
- Initial vs. updating

Define how service profiles are created

- Manually created
- Cloned from existing
- **Derived from template**

Define service profile parameters

- Basic vs. **advanced** service profiles



Service profile
UUID, MAC, WWN
Boot info
LAN, SAN config
LAN, SAN counts
Firmware

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-26

The server personality is described with the service profile. In cases where there are multiple similar service profiles, the deployment can be eased by templates—for service profiles, for NICs, for HBAs, and so on. This not only eases initial deployment of four vSphere hosts, like in the example, but also provides a configure-once-use-many-times template for each subsequent vSphere host deployment.

The next decision has to be made on the template type—either initial or updating.

The initial template is safer in terms of applying the changes, since the service profile configuration is changed per individual service profile. This way, if any intrusive configuration is applied, a single host would be affected.

The updating template is recommended for test and proof-of-concept environments, since it provides a single point of reconfiguration for all the service profiles and other objects that are connected to the template. This makes the configuration changes quick. This option is less attractive for the production environment, since the changes affect all the profiles or objects bound to the template, which could result in reboot of all the hosts.

In this example, the initial template is used, since the design is well-defined and the configuration does not need to be changed later.

As for the service profile template definition and the other template definition, the design is similar to the one where no Cisco Nexus 1000V was deployed with VMware vSphere.

Template—Identity

- Name—**SPT-ESXC1**
- UUID assignment from **USP-ESXC1**

Create Service Profile Template

1. ✓ **Identify Service Profile Template**
2. Storage
3. Networking
4. vNIC/vHBA Placement
5. Server Boot Order
6. Server Assignment
7. Operational Policies

Identify Service Profile Template

You must enter a name for the service profile template and assigned to this template and enter a description.

Name: **SPT-ESXC1**

The template will be created in the following organization. Its name must be

Where: **org-root/org-NILHC-test**

The template will be created in the following organization. Its name must be

Type: Initial Template Updating Template

Specify how the UUID will be assigned to the server associated with the se

UUID

UUID Assignment: **USP-ESXC1(0/4)**

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6-27

Once all the pools are defined, the service profile template can be designed.

The identity comes from the UUID identifiers that are specified by the USP-ESXC1 UUID pool.

In this case, the initial template is used. Since the design is well-defined there is no need to later change the configuration of all the hosts at once. By using the initial template, you are protected from unintentional configuration changes and all-at-once host reboots.

Template—Storage

- Use policy **LDCPol-ESXC1** for local storage
- SAN connectivity configuration expert mode
- nWWN assigned from **WNP-ESXC1**

The screenshot shows the 'Storage' configuration page. On the left, a navigation pane lists steps: 1. Identify Service Profile Template, 2. Storage (selected), 3. Networking, 4. vNIC/vHBA Placement, 5. Server Boot Order, 6. Server Assignment, and 7. Operational Policies. The main area is titled 'Storage' and contains the following fields and options:

- Local Storage: LDCPol-ESXC1 (dropdown menu)
- Mode: RAID 0 Striped
- Protect Configuration: yes
- How would you like to configure SAN connectivity? Simple (radio button), Expert (radio button, selected)
- World Wide Node Name: WNP-ESXC1(0/4) (dropdown menu)

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0—6-28

The next step of the service profile design is to configure storage.

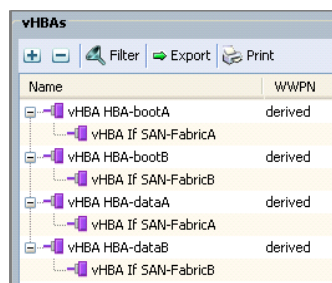
The local storage policy LDCPol-ESXC1 is used—it specifies that disks are used in the RAID stripe manner.

SAN connectivity is defined by expert mode, since more than two HBAs are needed.

The nWWN address assignment simply uses the pool that was defined in advance.

Template—vHBAs

- One per SAN segment



Name	WWPN
vHBA HBA-bootA	derived
vHBA If SAN-FabricA	
vHBA HBA-bootB	derived
vHBA If SAN-FabricB	
vHBA HBA-dataA	derived
vHBA If SAN-FabricA	
vHBA HBA-dataB	derived
vHBA If SAN-FabricB	

vHBA	Address Pool	SAN Segment	Fabric	VSAN
HBA-bootA	WPP-ESX-BootA	Boot SAN-A	A	SAN-FabricA (10)
HBA-bootB	WPP-ESX-BootB	Boot SAN-B	B	SAN-FabricB (20)
HBA-dataA	WPP-ESX-DataA	Data SAN-A	A	SAN-FabricA (10)
HBA-dataB	WPP-ESX-DataB	Data SAN-B	B	SAN-FabricB (20)

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-29

The HBAs can be defined by using the information that is listed in the table.

The table specifies the HBA name, the pWWN address pool from which the address should be assigned, SAN segment assignment (that is, VSAN), and fabric ID.

As required, there will be a single SAN per fabric segment, resulting in two HBAs per SAN functional segment (data or boot).

Template—Networking

- LAN connectivity configuration expert mode
- One vNIC per LAN segment

VLAN	ID
Mgmt	1
VMdata1	11
VMdata2	12
VMdmz1	21
VMdmz2	22
ESXvMotion	997

vNIC	MAC Pool	LAN Seg.	Fabric	VLAN	Trunk
NIC-mgmt	MP-ESXmgmt	Mgmt	A-B	Mgmt (native)	No
NIC-VMdata	MP-VMdata	VMdata	A-B	VMdata1, VMdata2	Yes
NIC-VMdmz	MP-VMdmz	VMdmz	A-B	ESXvMotion	Yes
NIC-ESXvMotion	MP-ESXvMotion	ESXvMotion	B-A	VMdmz1, VMdmz2	Yes

Name	MAC Address	Fabric ID	Native VLAN
vNIC NICmgmt	derived	A-B	
Network default			
vNIC NICVMdata	derived	A-B	
Network VMdata1			
Network VMdata2			

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-30

For the network connectivity—that is, NIC definition—the example also uses NIC templates. This approach speeds up deployment as well as design.

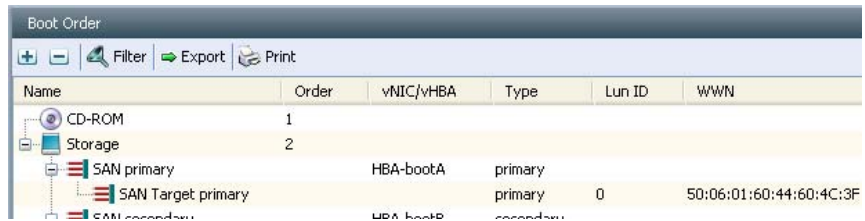
The table in the figure specifies the NIC template characteristics—the MAC pools to get the address from, the LAN segment to be connected to, primary and failover fabric selection, and trunking configuration.

All the segments of the host will be connected with a single NIC that will be configured with Cisco UCS level failover—this achieves proper redundancy for all the segments.

Template—Boot Order

Create **BP-ESXC1-SANboot** policy with boot order

- CD-ROM
- Storage
 - SAN primary—HBA-bootA
 - SAN secondary—HBA-bootB
 - Specify drive array WWN (target WWN)



Name	Order	vNIC/vHBA	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		HBA-bootA	primary		
SAN Target primary		HBA-bootA	primary	0	50:06:01:60:44:60:4C:3F
SAN secondary		HBA-bootB	secondary		

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-31

The vSphere hosts should be booted off SAN via the dedicated HBAs—two are used for redundancy purposes. HBA-bootA is the primary and HBA-bootB is the secondary.

The boot policy design specifies the following:

- **Boot order:** CD/DVD drive, then LUN SAN target. It is necessary to specify the CD/DVD option in order to be able to install the hypervisor.
- **SAN boot devices:** HBA-bootA as the primary and HBA-bootB as the secondary.
- **SAN boot target devices:** One for each HBA specified above. The pWWN used in this part depends on the drive array—that is, it is the pWWN of the service profile port through which a boot LUN is visible.

Template—Server Assignment

Automatic selection from pool **SRVP-ESXC1**

Use **HFP-ESXC1** policy for host firmware

Set power state to **Down**

The screenshot shows two configuration panels from the Cisco UCS Manager interface. The top panel, titled 'Server Assignment', is part of a 'Create Service Profile Template' wizard. It shows a progress list on the left with 'Server Assignment' selected. The main area contains instructions to optionally specify a server pool and a dropdown menu for 'Pool Assignment' set to 'SRVP-ESXC1'. To the right, there are radio buttons for 'Down' (selected) and 'Up'. The bottom panel, titled 'Firmware Management (BIOS, Disk Controller, Adapter)', contains instructions about host or management firmware policy and a dropdown menu for 'Host Firmware' set to 'HFP-ESXC1', with a 'Create Host Firmware Package' button next to it.

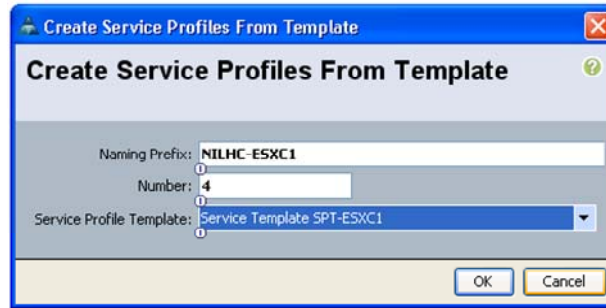
Server assignment is straightforward. Use the server pool SRVP-ESXC1 that has been populated with the pool policy.

Similarly, the predefined firmware package is used in the service profile template.

ESX Service Profiles Design

Derive service profiles from template **SPT-ESXC1**

- Service profile prefix = **NILHC-ESXC1**
- Quantity = 4



Once you have defined the service profile template, the design must specify how many service profiles should be created out of a template and what the service profile prefix name should be.

In the example, the prefix for the name is NILHC-ESXC1, and four service profiles are specified to be created for the four vSphere hosts required.

ESX Service Profiles

The screenshot displays the ESX Service Profiles interface. On the left, a tree view shows service profiles: Service Profile NILHC-ESX1, Service Profile NILHC-ESX2, Service Profile NILHC-ESX3, and Service Profile NILHC-ESX4. Each profile is associated with vHBAs, vNICs, and a Server Pool SRVP-ESXC1. A red arrow points from the 'Service Profile NILHC-ESX4' entry in the tree to a table on the right. Above the table is a blue box labeled 'Associated Server Blades'. The table lists the following data:

Name	Profile	Assoc State	Chas
Server 4	org-root/org-NILHC-test/ls-NILHC-ESX4	associated	1
Server 5	org-root/org-NILHC-test/ls-NILHC-ESX1	associated	1
Server 6	org-root/org-NILHC-test/ls-NILHC-ESX2	associated	1
Server 7	org-root/org-NILHC-test/ls-NILHC-ESX3	associated	1

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6-34

If the service profiles are created from a template that automatically assigns blades from the server pool, the result is four service profiles that are associated with physical blades once the implementation engineer starts the service profile creation.

Since the NILHC-ESX prefix was used for the service profile, we have service profiles named NILHC-ESX1, NILHC-ESX2, NILHC-ESX3, and NILHC-ESX4.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Advanced service profile deployment gives the maximum flexibility for deploying VMware vSphere.
- Cisco UCS LAN and SAN configuration depend on VMware vSphere requirements.
- Use addresses that "speak" to ease deployment and troubleshooting.
- Service profile templates and server assignment from pool speed the deployment.

Evaluating Cisco UCS and Cisco Nexus 1000V Integration with VMware vSphere

Overview

This lesson evaluates Cisco Unified Computing System (UCS) integration with VMware vSphere and Cisco Nexus 1000V.

Objectives

Upon completing this lesson, you will be able to integrate Cisco UCS, Cisco Nexus 1000V, and VMware vSphere. This includes the ability to meet these objectives:

- Evaluate VMware vSphere and Cisco Nexus 1000V deployment requirements
- Propose a design for deploying VMware vSphere with Cisco Nexus 1000V

Assessing VMware vSphere with Cisco Nexus 1000V Requirements

This topic evaluates VMware vSphere and Cisco Nexus 1000V deployment requirements for a given example.

Nexus 1000V Requirements

ESX cluster = single Cisco Nexus 1000V switch

- Two VSM for HA
- One VEM per ESX host

License per host CPU for each VEM

- Installed on the VSM
- One or more license files can be installed
- Installation is non-disruptive to operation
- Dual supervisor environment runs licensed software on both

The diagram illustrates the Nexus 1000V architecture. At the top, two blue blocks represent the Primary VSM and Secondary VSM, connected by a yellow bar. Below them, a yellow bar labeled 'ESX Cluster' contains six blue blocks representing VEMs, arranged in two rows of three. A yellow bar labeled 'VEMs' is positioned below the ESX Cluster.

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6.4

Cisco Nexus 1000V Overview

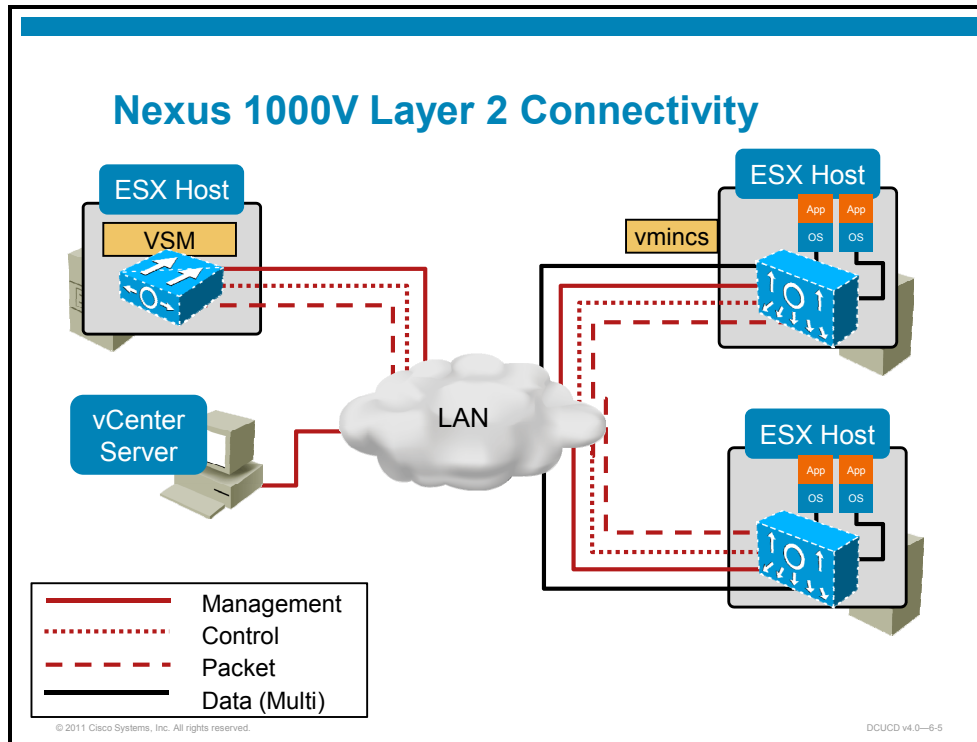
The Cisco Nexus 1000V uses distributed architecture. This architecture separates control plane and data plane functionality. The control plane functionality is represented by Cisco Visual Switch Manager (VSM), which manages multiple distributed data planes (Virtual Ethernet Module [VEM] in each ESX host). Thus, the VSM acts as a supervisor module for remote VEMs.

Cisco Nexus 1000V VSM

All configuration and supervisor functions are managed by the VSM. Using Console, Telnet, or SSH, an administrator makes all configuration changes on the VSM. When a change is made on the VSM, the configuration is passed to vCenter and the changes are made on the distributed virtual switch (DVS). DVS changes are passed down to the corresponding VEM.

Cisco Nexus 1000V VEM

Each VEM will act as a module on the VSM, and the VSM or VEM will appear as a single switch to Cisco Discovery Protocol neighbors. The VSM does not reside in the data path and therefore cannot directly receive or respond to Cisco Discovery Protocol messages. Cisco Discovery Protocol and other network management packets are transferred between the VEM and the VSM on one of three required VLANs, known as the packet VLAN.



The VSM, VEM, vCenter, and VM connectivity use dedicated VLANs—management, control, packet, and one or more data networks.

Management VLAN

Each VMware ESX host, the VSM, and the vCenter server must all reside in the same management network and be part of the same Layer 2 domain.

Domain ID

A single Cisco Nexus 1000V instance, including dual redundant VSMs and managed VEMs, forms a switch domain. Each Cisco Nexus 1000V domain within a VMware vCenter Server needs to be distinguished by a unique integer that is called the domain identifier.

Control and Packet VLANs

Control and packet VLANs from the VSM must be accessible by uplink profiles on each VEM. The control VLAN and the packet VLAN are used for communication between the VSM and the VEMs within a switch domain.

The packet VLAN is used by protocols such as Cisco Discovery Protocol, Link Aggregation Control Protocol (LACP), and Internet Group Management Protocol (IGMP).

The control VLAN is used for the following:

- VSM configuration commands to each VEM, and their responses
- VEM notifications to the VSM—for example, a VEM notifies the VSM of the attachment or detachment of ports to the DVS
- VEM NetFlow exports are sent to the VSM, where they are then forwarded to a NetFlow collector.

Data VLANs

The data networks carry VM packet traffic (server data). One or more data VLANs are defined for this purpose. Data traffic from the VM is not sent to the VSM and the VSM does not require access to the data VLANs. All VSM management is out of band and switching decisions do not rely on the VSM.

Note It is recommended to set up the control VLAN and packet VLAN as separate VLANs, and to put them on separate VLANs from those that carry data.

VMware vSphere Requirements

Four VMware ESX hosts

Dual-fabric design

- SAN uplink—two 4-Gb Fibre Channel
- LAN uplink—two 10-Gb Ethernet
- Server downlinks—two 10-Gb Ethernet with FCoE

SAN boot traffic via dedicated Fibre Channel interfaces

SAN data traffic via dedicated Fibre Channel interfaces



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v1.0-6-6

In this example, the VMware vSphere deployment is performed with four ESX hosts, like in the earlier example.

LAN Connectivity Requirements

The physical connectivity requirements dictate that more than two vSphere hosts should be connected to two separate fabrics for each LAN segment that requires redundancy.

Thus, there will be two separate physical LAN fabrics on the Cisco UCS cluster level, which converge in the LAN core on the same set of devices. It has been determined that two 10 Gigabit Ethernet uplinks have enough throughput for all the traffic.

SAN Connectivity Requirements

Likewise, the SAN connectivity should be redundant for each segment or host bus adapter (HBA) that connects vSphere hosts to the SAN fabric. Further, you must use separate segments for accessing the boot and data partitions of the vSphere hosts and environment.

Thus, from the Cisco UCS cluster perspective, the two separate physical SAN fabrics will be used, connected to upstream Fibre Channel switches with the N_Port ID Virtualization (NPIV) feature enabled. Boot and data traffic will be separated through zoning in the SAN and logical unit number (LUN) masking on the drive array side.

ESX Host Requirements

Half-width blades

Memory = minimum 4 GB at 1067 MHz

Processor = minimum 4-core at 2.266 GHz

Boot from SAN

Local disk = striped (RAID 0) for local datastore

Adapter firmware = 1.3(1c)



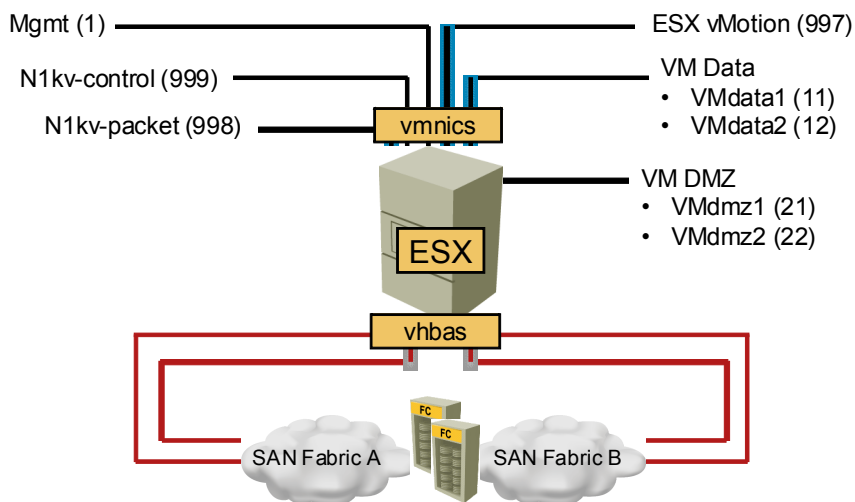
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-7

Each ESX host should be deployed using hardware with these minimum characteristics:

- **Blade form factor:** Half-width blade
- **Memory:** At least 4 GB at 1067 MHz
- **CPU:** At least 4-core at 2.266 GHz
- **Boot policy:** ESX should be booted from SAN to support fast blade replacement and complete stateless configuration
- **Local disk:** Used by the local datastore to store ISO images and nonmovable VMs (like VSM for Cisco Nexus 1000V)
- **Adapter firmware:** Cisco UCS Manager version 1.3(1c) should be used since it is certified to be used with the selected drive array

ESX Host Connectivity Requirements



The individual vSphere host has the following connectivity requirements:

- **Mgmt segment:** This segment connects the host to a management segment that is used for management, configuration, and state maintenance, as well as for communication with the vSphere vCenter server. For this example, the default VLAN 1 will be used.
- **ESX vMotion segment:** This segment connects the vSphere hosts in a cluster to enable migration of VMs from one host to another, as well as to enable advanced VMware vSphere high availability and resource management functionalities like High Availability (HA), Fault Tolerance (FT), Distributed Resource Scheduler (DRS), and Distributed Power Management (DPM). For this example, VLAN 997 will be used.
- **N1kv-control segment:** This segment enables control communication between Cisco Nexus 1000V components—VSMs and VEMs. For this example, VLAN 999 will be used.
- **N1kv-packet segment:** This segment enables other control and management communication between Cisco Nexus 1000V components (for instance, Cisco Discovery Protocol, IGMP, and so on). For this example, VLAN 998 will be used.
- **VM data segment:** This segment is used for communication between VMs that reside in the private sector of the Data Center. VLANs 11 and 12 have been currently identified as required, but others can be added later. The segment thus must be configured with trunking capability.
- **VM dmz segment:** This segment is used for communication between VMs that reside in less trusted networks. VLANs 21 and 22 have been currently identified as required, but others can be added later. Like the private VM data segment, this segment must be configured with trunking capability.

ESX Host Identity Requirements

Six LAN segments = six vNICs

Four SAN segments = four vHBAs

UUID

Identity assignment

- Derived vs. manual vs. **automatic from pool**
- To abstract hardware from personality



© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-9

Due to the connectivity requirements, the vSphere host must be equipped with the following:

- Six Ethernet NICs to connect to different LAN segments
- Four Fibre Channel HBAs to connect to different SAN segments

The vSphere host also requires the following:

- UUID identifier per vSphere host
- Automatic server assignment from pool, based on server hardware requirements

Designing VMware vSphere with Cisco Nexus 1000V Deployment

This topic proposes a design for deploying VMware vSphere with Cisco Nexus 1000V for the given example.

Nexus 1000V Design

Dedicated VLANs

- Management, control, packet

HA solution

- Primary and secondary VSM on different hosts

© 2011 Cisco Systems, Inc. All rights reserved. DCUCD v4.0-6-11

Planning the Cisco Nexus 1000V deployment must take into consideration Cisco Nexus 1000V, VMware, and uplink switch aspects.

From the Cisco Nexus 1000V perspective, the following must be addressed:

- **Licensing**—Cisco Nexus 1000V is licensed per server CPU. The designer must thus know how many ESX hosts will be initially used and how many will be used in the future in order to plan and select the proper licensing pack.
- **VLAN scheme**—For the Cisco Nexus 1000V, a deployment minimum of three VLANs is required—management, control, and packet. The designer must reserve and assign these VLAN IDs from the free VLAN ID pool. Although the VLAN IDs could be any of those already used, it is recommended to separate VLAN IDs.
- **Design VSM deployment**

Naming Convention and Segments

Identifier	Description	Type	Naming
ESX	Denotes ESX hypervisor	MAC	MP-ESX <i>LANseg</i>
SubsegID	Denotes subsegment ID (for example, VLAN on physical segment)	nWWN	WNP-ESX <i>poolID</i>
LANseg	LAN segment	pWWN	WPP-ESX <i>SANseg</i>
SANseg	SAN segment	UUID	USP-ESX <i>poolNum</i>
		VLAN	<i>LANsegSubsegID</i>
		VSAN	<i>SANsegSubsegID</i>
		vNICs	NIC- <i>LANseg</i>
		vHBAs	HBA- <i>SANseg</i>
		Server pool	SRVP-ESX <i>poolID</i>
		Server pool policy	SRVPP-ESX <i>poolID</i>
		Server qualification	SPPQ-ESX <i>poolID</i>
		Host firewall package	HFP-ESX <i>poolID</i>

LAN Segment	ID	SAN Segment
Mgmt	001	SAN-FabricA
VMdata	010	SAN-FabricB
VMdmz	020	
ESXvMotion	997	
N1kv-packet	998	
N1kv-control	999	

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-12

To ease the deployment effort, as well as any troubleshooting that might be needed later, the naming convention will incorporate segment information into the addresses and identifiers. This is indicated in the tables in the figure.

Note that this is only a suggestion of a good design practice, and is not enforced by Cisco UCS or by the Cisco UCS Manager.

LAN Segments and VLAN Design

VLAN scheme

- Fabric **global** (same VLAN ID in fabric A and B)
- VLAN IDs—**must be unique** (at least per fabric)

Operational mode

- **EHV** (default)
- Switching mode

VLAN	ID	Fabric
Mgmt	1	A, B
VMdata1	11	A, B
VMdata2	12	A, B
VMdmz1	21	A, B
VMdmz2	22	A, B
ESXvMotion	997	A, B
N1kv-control	999	A, B
N1kv-packet	998	A, B

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-13

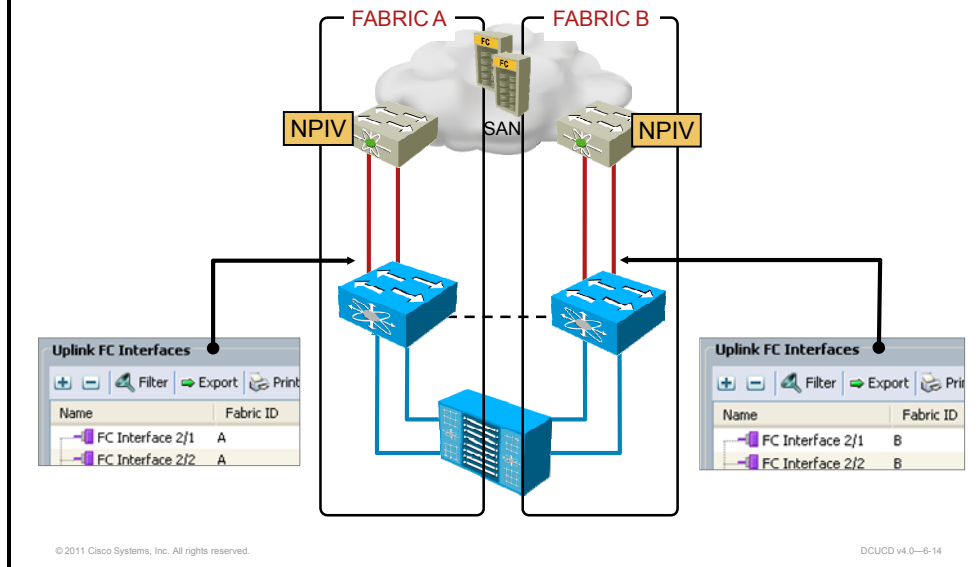
The VMware vSphere with Nexus 1000V deployment will be physically connected to a single physical LAN core; that is, the LAN core is not disjointed (formed from two or more disconnected domains).

For this reason—as well as the Cisco UCS level redundancy used to achieve failover between the fabric and manual pinning for the uplink selection—the Cisco UCS cluster will operate in Ethernet host virtualizer (EHV) mode for Ethernet. This also eliminates Spanning Tree Protocol (STP) in the access layer.

VLAN Design

The VLANs that are used will have the same IDs in both fabrics, so they will be defined globally in Cisco UCS with the VLAN IDs that are specified in the table in the figure.

SAN Physical Connectivity



The physical SAN connectivity uses multiple Fibre Channel uplinks in each fabric—in this example, two per fabric. The Fibre Channel uplinks in fabric A are set to VSAN 10, and the Fibre Channel uplinks in fabric B are set to VSAN 20.

SAN Configuration

Must be defined before Fibre Channel uplink configuration

VSAN scheme

- Fabric A or B
- VSAN IDs—**must be unique** (at least per fabric)
- FCoE VLAN ID—**must be unique** (at least per fabric)

Operational mode is NPV Edge

- Upstream Fibre Channel switches must be enabled with NPIV

VSAN	VSAN ID	Fabric	FCoE VLAN ID
SAN-FabricA	10	A	1010
SAN-FabricB	20	B	1020

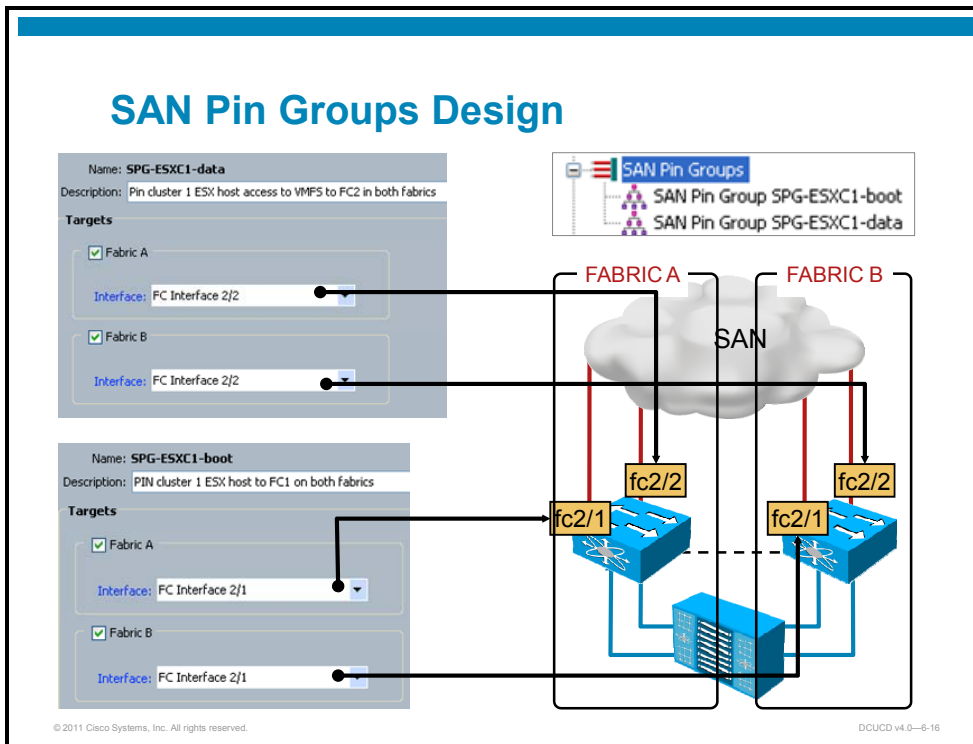
© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-15

The SAN design is simple—two separate physical fabrics using the VSAN IDs, as specified in the table. Since FCoE is used, VLAN FCoE IDs have to be specified—also indicated in the table. A good design practice is to add a prefix to the VSAN ID in order to get a FCoE VLAN ID if the same VLAN ID is already in use (for example, adding 10 in front of VSAN to get the FCoE VLAN ID).

In the core SAN fabric, the core Fibre Channel switches must be enabled with the NPIV feature. Also, for the Cisco Multilayer Director Switch (MDS), the Cisco UCS attached interfaces must be part of the same VSAN (that is, in access mode). Thus, from a design perspective, the Cisco UCS interfaces that are attached to the MDS must have the same configuration as the Cisco UCS Fibre Channel interfaces.

SAN Pin Groups Design



In the example, the deployment requirement is also to separate SAN traffic for boot LUN access and data LUN—VMFS access. This is achieved with multiple HBAs on the service profile level, which have manual pinning that defines which uplink will be selected.

In the example, there are two uplinks per fabric, which enables separation of traffic mentioned. Thus, to be able to apply manual traffic engineering, SAN pin groups have to be defined:

- **SPG-ESXC1-data pin group:** Used to select the uplinks for accessing the VMFS datastore of a disk array, and will use interface Fibre Channel 2/2 in both fabrics (that is, on both fabric interconnects).
- **SPG-ESXC1-boot pin group:** Used to select the uplinks for accessing the vSphere host boot LUN of a disk array, and will use interface Fibre Channel 2/1 in both fabrics

Virtual Identity Pools

Pool	Address Template
MAC	00:25:b5:3m:mm:xx
nWWN	20:11:00:25:b5:00:30:xx
pWWN	20:00:00:25:b5:3s:ss:xx
UUID	Prefix = 30010000-0000-0000 Master suffix = 3000-1000000000xx

ID	Description	LAN Segment	ID	SAN Segment	ID
mmm	LAN segment ID	Mgmt	001	Boot SAN-A	019
xx	Host ID	VMdata	010	Boot SAN-B	029
sss	SAN segment ID	VMdmz	020	Data SAN-A	011
ESX OS	3	ESXvMotion	997	Data SAN-B	021
ESX pool num	001	N1kv-packet	998		
		N1kv-control	999		

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-17

To ease the deployment effort and troubleshooting in the future, and to have better control of the environment when it scales, the virtual identity pools with addresses and identifiers will be defined. This includes MAC, UUID, nWWN, and pWWN pools and addresses.

The addresses and identifiers will have segments and IDs encoded to further ease deployment and troubleshooting using the following fields:

- **mmm:** Denotes the LAN segment ID and encodes the VLAN number in the MAC address. For the VM data segment, since there will be a single NIC used, the identifier will be the least common denominator; in other words, the last part of the VLAN ID will be omitted. The same goes for the VM DMZ segment.
- **xx:** Denotes the host ID, and is picked and determined by Cisco UCS itself.
- **sss:** Denotes the SAN segment ID and encodes the VSAN number that is added by the suffix, where 1 means data segment and 9 means boot segment.
- **ESX pool number 001:** To anticipate future addition of vSphere pools, this first ESX pool will be denoted with 001, which will be encoded into the addresses.
- **ESX OS ID 3:** Since Cisco UCS enables deployment of multiple different operating systems and hypervisors, information about the operating system and hypervisor type will also be encoded into the addresses and identifiers. In this case, the vSphere hypervisor will have an ID of 3.

The table at the top of the figure lists the pool types and also specifies a template for creating addresses and identifiers in each pool that have these fields encoded.

MAC Address Pools

MAC		Address
MP-ESXmgmt	First	00:25:b5:30:01:01
	Last	00:25:b5:30:01:04
MP-VMdata	First	00:25:b5:30:10:01
	Last	00:25:b5:30:10:04
MP-VMdmz	First	00:25:b5:30:20:01
	Last	00:25:b5:30:20:04
MP-ESXvMotion	First	00:25:b5:39:97:01
	Last	00:25:b5:39:97:04
MP-N1kv-packet	First	00:25:b5:39:98:01
	Last	00:25:b5:39:98:04
MP-N1kv-control	First	00:25:b5:39:99:01
	Last	00:25:b5:39:99:04

LAN Segment	ID
Mgmt	001
VMdata	010
VMdmz	020
ESXvMotion	997
N1kv-packet	998
N1kv-control	999

Easy vmnic Identification →

ESX Host

Configuration

Network Adapters

Device	MAC Address
Cisco Systems Inc VIC Ethernet NIC	
vmnic5	00:25:b5:39:98:02
vmnic4	00:25:b5:39:99:02
vmnic3	00:25:b5:30:20:02
vmnic2	00:25:b5:39:97:02
vmnic1	00:25:b5:30:10:02
vmnic0	00:25:b5:30:01:03

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-18

MAC address pools are designed for each NIC that will be available to the vSphere host. It has been determined that individual host will have six NICs, thus six MAC pools are required.

- **MP-ESXmgmt pool:** MAC addresses that are used by the NIC that is connected to the management segment. Encodes VLAN 1, which is used for the management segment.
- **MP-VMdata pool:** MAC addresses that are used by the NIC that is connected to the VM data segment, which combines the VMdata1 and VMdata2 subsegments (VLANs 11 and 12). Thus, the encoded field is 10 (the least common denominator).
- **MP-VMdmz pool:** MAC addresses that are used by the NIC that is connected to the VM DMZ segment, which combines the VMdmz1 and VMdmz2 subsegments (VLANs 21 and 22). Like the VM data pool, this one encodes field value 20, which is the least common denominator for the given segments.
- **MP-ESXvMotion pool:** MAC addresses that are used by the NIC that is connected to the vMotion vSphere segment, which is utilized by the VM migration as well as advanced functionalities like high availability, FT, DRS, and DPM.
- **MP-N1kv-packet pool:** MAC addresses that are used by the NIC connecting the N1kv packet VLAN, which is used for some control protocol communication between N1kv components (Cisco Discovery Protocol, IGMP, and so on).
- **MP-N1kv-control pool:** MAC addresses that are used by the NIC connecting the N1kv control VLAN, which is used for control and management communication between N1kv VSM and VEM components.

You can see from the tables in the figure that the MAC addresses in an individual pool are created based on the MAC address template that was specified earlier.

Notice the NIC appearance in the vSphere vCenter configuration—there is no similarity between the name of the vmnic (the vSphere host physical NIC) and the NIC defined by Cisco UCS. However, since the MAC addresses have encoded fields, it is easy to determine which vmnic should be used for which segment. Thus, deployment is easier.

WWN Address Pools

nWWN		Address
WNP-ESXC1	First	20:11:00:25:b5:00:30:01
	Last	20:11:00:25:b5:00:30:04

pWWN		Address
WPP-ESX-BootA	First	20:00:00:25:b5:30:19:01
	Last	20:00:00:25:b5:30:19:04
WPP-ESX-BootB	First	20:00:00:25:b5:30:29:01
	Last	20:00:00:25:b5:30:29:04
WPP-ESX-DataA	First	20:00:00:25:b5:30:11:01
	Last	20:00:00:25:b5:30:11:04
WPP-ESX-DataB	First	20:00:00:25:b5:30:21:01
	Last	20:00:00:25:b5:30:21:04

SAN Segment	ID
Boot SAN-A	019
Boot SAN-B	029
Data SAN-A	011
Data SAN-B	021

Easy vmhba
Identification



Storage Adapters		ESX Host
Device	WWN	
vmhba5	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:19:03	
vmhba6	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:29:03	
vmhba7	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:11:03	
vmhba8	20:11:00:25:b5:00:30:03 20:00:00:25:b5:30:21:03	

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-19

nWWN pools are designed for each vSphere host that will be connected to the SAN attached storage. Since there are only four hosts initially, the WNP-ESXC1 nWWN pool is designed with only four nWWN addresses. However, the addresses have an encoded operating system type (3 for vSphere hypervisor) and host ID. The VSAN segment is not encoded since the host has only one nWWN address in the fabric.

The pWWN address pools are designed for each HBA that will be available to the vSphere host. It has been determined that individual host will have four HBAs, so four pWWN pools are required.

- **WPP-ESX-BootA pool:** pWWN addresses that are used by the HBA that is connected to the boot segment in fabric A. Uses field value 19 to encode VSAN 10 and identify the boot HBA.
- **WPP-ESX-BootB pool:** pWWN addresses that are used by the HBA that is connected to the boot segment in fabric B. Uses field value 29 to encode VSAN 20 and identify the boot HBA.
- **WPP-ESX-DataA pool:** pWWN addresses that are used by the HBA that is connected to the data segment in fabric A (the VMFS). Uses field value 11 to encode VSAN 10 and identify the data HBA.
- **WPP-ESX-DataB pool:** pWWN addresses that are used by the HBA that is connected to the data segment in fabric B (the same VMFS as fabric A). Uses field value 21 to encode VSAN 20 and identify the boot HBA.

You can see from the tables in the figure that the pWWN addresses in an individual pool are created based on the pWWN address template that was specified earlier.

Notice the HBA appearance in the vSphere vCenter configuration—there is no similarity between the name of the vmhba (the vSphere host physical HBA) and the HBA defined by Cisco UCS. However, since the pWWN addresses have encoded fields, it is easy to determine which vmhba should be used for which segment. Thus, deployment is easier.

UUID Pool

ID	Description
ESX OS	3
ESX pool num	001

UUID	Prefix		Suffix
USP-ESXC1	30010000-0000-0000	First	3000-1000000000001
		Last	3000-1000000000004



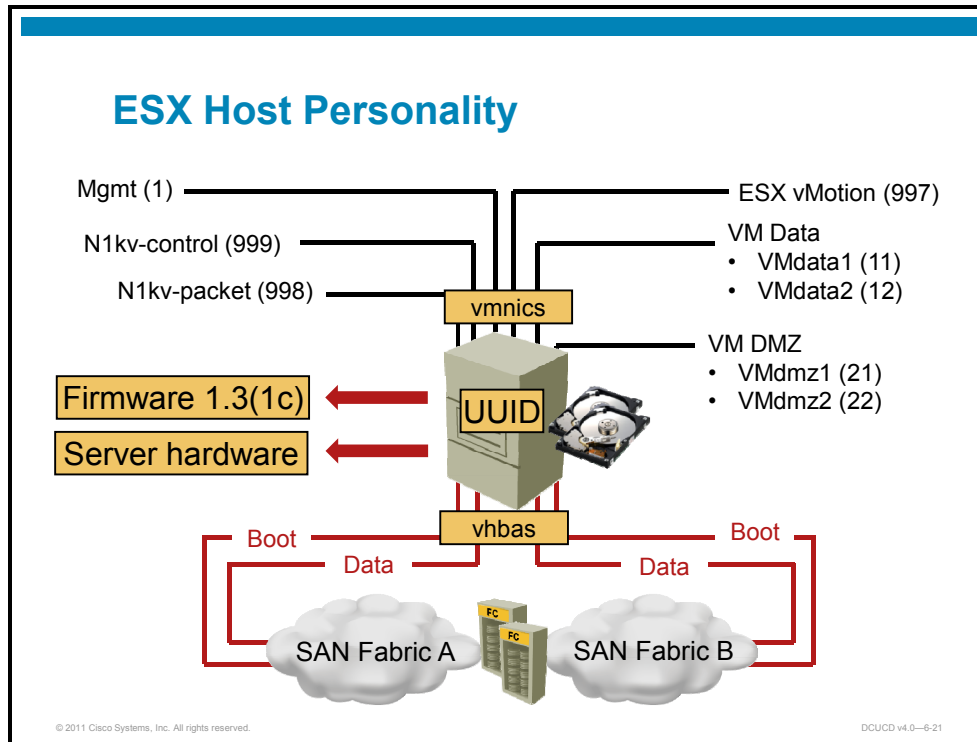
UUID Assignment		
Name	Assigned	Assigned To
3000-0000000000001	yes	org-root/org-NILHC-test/ls-NILHC-ESX4
3000-0000000000002	yes	org-root/org-NILHC-test/ls-NILHC-ESX3
3000-0000000000003	yes	org-root/org-NILHC-test/ls-NILHC-ESX2
3000-0000000000004	yes	org-root/org-NILHC-test/ls-NILHC-ESX1

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-20

The UUID pool and addresses are defined per the UUID template that was defined earlier—they encode the operating system type (3 for vSphere host) and vSphere cluster ID (001 for the first one).

Like the MAC and pWWN addresses, the NICs, and the HBAs, the vCenter configuration is easier to read due to encoded fields.



Combining all the requirements, the vSphere host personality consists of the following:

- LAN connectivity definition with six NICs
- SAN connectivity definition with four HBAs
- Host firmware package with 1.3(1c) firmware version
- Distinct server hardware (that is, certain physical resources have to be available)
- Two local disks for local datastore
- UUID identifier to distinguish between the vSphere hosts

Service Profiles

Service template

- For fast, massive server deployment
- Initial vs. updating

Define how service profiles are created

- Manually created
- Cloned from existing
- **Derived from template**

Define service profile parameters

- Basic vs. **advanced** service profiles



Service profile
UUID, MAC,WWN
Boot info
LAN, SAN config
LAN, SAN counts
Firmware

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-22

The server personality is described with the service profile. In cases where there are multiple similar service profiles, the deployment can be eased by templates—for service profiles, for NICs, for HBAs, and so on. This not only eases initial deployment of four vSphere hosts, like in the example, but also provides a “configure-once-use-many-times” template for each subsequent vSphere host deployment.

The next decision has to be made on the template type—either initial or updating.

The initial template is safer in terms of applying the changes, since the service profile configuration is changed per individual service profile. This way, if any intrusive configuration is applied, a single host would be affected.

The updating template is recommended for test and proof-of-concept environments, since it provides a single point of reconfiguration for all the service profiles and other objects that are connected to the template. This makes the configuration changes quick. This option is less attractive for the production environment, since the changes affect all the profiles or objects bound to the template, which could result in reboot of all the hosts.

In this example, the initial template is used, since the design is well-defined and the configuration does not need to be changed later.

As for the service profile template definition and the other template definition, the design is similar to the one where no Cisco Nexus 1000V was deployed with VMware vSphere.

Template—Networking

- LAN connectivity configuration expert mode
- One vNIC per LAN segment

VLAN	ID
Mgmt	1
VMdata1	11
VMdata2	12
VMdmz1	21
VMdmz2	22
N1kv-control	999
N1kv-packet	998
ESXvMotion	997

vNIC05	MAC Pool	LAN Seg.	Fabric	VLAN	Trunk
NIC-mgmt	MP-ESXmgmt	Mgmt	A-B	Mgmt (<i>native</i>)	No
NIC-VMdata	MP-VMdata	VMdata	A-B	VMdata1, VMdata2	Yes
NIC-VMdmz	MP-VMdmz	VMdmz	A-B	ESXvMotion	Yes
NIC-ESXvMotion	MP-ESXvMotion	ESXvMotion	B-A	VMdmz1, VMdmz2	Yes
NIC-N1kv-packet	MP-N1kv-packet	N1kv-packet	B-A	N1kv-packet	Yes
NIC-N1kv-control	MP-N1kv-control	N1kv-control	B-A	N1kv-control	Yes

© 2011 Cisco Systems, Inc. All rights reserved.

DCUCD v4.0-6-23

For the network connectivity—that is, NIC definition—the example also uses NIC templates. This approach speeds up deployment as well as design.

The table in the figure specifies the NIC template characteristics—the MAC pools to get the address from, the LAN segment to be connected to, primary and failover fabric selection, and trunking configuration.

All the segments of the host will be connected with a single NIC that will be configured with Cisco UCS level failover—this achieves proper redundancy for all the segments.

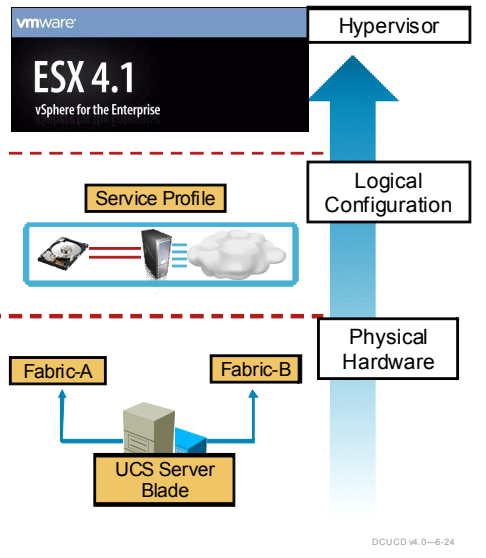
VMware vSphere Setup

Install individual VMware ESX hosts

- Open KVM console for service profile
- Add and used VMware ISO under virtual media
- Start the installation

Install vCenter Server

- Add VMware hosts



From the VMware and Nexus 1000V perspective, the Cisco UCS design defines all that is necessary for the deployment—the physical hardware layout and resources as well as logical server configuration. However, with the Cisco Nexus 1000V, there are some other aspects that also need to be addressed.

VSM Deployment Design

VSM deployment design must address the following Cisco Nexus 1000V aspects:

- Select the ESX host where the VSM appliance will be running—the host should have sufficient resources for the VSM VM.
- Define management parameters like management IP address and login credentials.
- Define the domain ID parameters for multiple Cisco Nexus 1000V deployments. A unique domain ID per Cisco Nexus 1000V domain is recommended.
- If you must use the same control and packet VLAN pair for multiple domains, you must ensure that their domain identifiers are different.
- Define VSM restrictions. For the VSM appliance, vMotion should be prohibited, and DRS and FT should be disabled.
- Define the system uplink port profile for the control and packet VLANs.
- Define the redundancy scheme for the VSM—standalone or active-standby (primary-secondary). Select the ESX host for the secondary VSM appliance. The selected host should be different from the one selected for the primary VSM.

Once the Cisco Nexus 1000V deployment is designed, the VM deployment can be planned. This includes Cisco Nexus 1000V, VMware, and upstream switch settings.

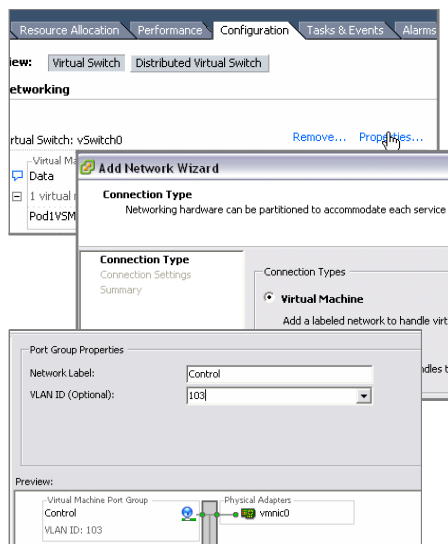
Prepare VMware Hosts for Nexus 1000V

Define management, control, packet VLANs

- Required when VSM is running as VM

Create VM port groups for each VLAN on VSM hosts

VLAN	ID	Port Group
Mgmt	1	Management
Control	999	Control
Packet	998	Packet



The Cisco UCS service profiles have been designed with proper NICs and VLANs for N1kv segments. The Cisco Nexus 1000V now has to be designed with the same information—the proper NICs defined in Cisco UCS must be used with the proper VLANs. Although the Cisco UCS design already contains this information, the overall design for VMware, Cisco Nexus 1000V, and Cisco UCS should also have a separate segment where this information is repeated for the Nexus 1000V.

Apart from that, the Cisco Nexus 1000V design also needs to define the VM port profiles that will be utilized by the VMs to connect to the network.

Cisco Nexus 1000V VM Deployment Design

The design should define the following:

- VLAN scheme for data traffic from different VMs. The scheme would typically define multiple VLANs for VM connectivity.
- VM port profiles for VM connectivity. A common VM profile for applications with the same connectivity requirements should be defined.
- Per VM port profile policy, which includes quality of service (QoS), security, and other settings.

VMware Design for VM Deployment

The design should specify the port groups derived from port profiles for different VMs. The port group description should specify the connectivity policy, which is derived from the port profiles.

Upstream Switch Design

The design should specify the interface settings of the upstream switch where the ESX server is connected for the vmnic carrying data VLANs.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco Nexus 1000V is a replacement for the VMware native distributed virtual switch.
- Cisco Nexus 1000V comprises VSMS (two for redundancy) and VEMs.
- Cisco UCS LAN and SAN configuration depend on VMware vSphere and Nexus 1000V requirements.

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- Network connectivity redundancy is best achieved with LAN failover capability at the Cisco UCS level.
- Cisco UCS LAN and SAN configuration depend on VMware vSphere and Nexus 1000V requirements.
- The Cisco Nexus 1000V is a replacement for the VMware native distributed virtual switch.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which Cisco UCS feature can be used to make the Hyper-V network connectivity highly available? (Source: Evaluating Cisco UCS Deployment with Microsoft Hyper-V)
- A) service profile
 - B) NIC hardware-based failover
 - C) vNIC template
 - D) Fibre Channel hardware-based failover
 - E) scrub policy
- Q2) What should be used to achieve traffic separation for Hyper-V Live Migration and data segments on the Cisco UCS 6120XP uplinks? (Source: Evaluating Cisco UCS Deployment with Microsoft Hyper-V)
- A) allowed VLAN list
 - B) global VLAN definition
 - C) LAN pin groups
 - D) separate vNIC for each segment
 - E) additional port license
- Q3) Which mechanism can be used to scale the bandwidth requirement for Hyper-V deployment from the Cisco UCS 6140XP perspective? (Source: Evaluating Cisco UCS Deployment with Microsoft Hyper-V)
- A) 40 Gb Ethernet interfaces
 - B) additional vNICs per Hyper-V host
 - C) LAN uplinks connected to separate core LAN switches
 - D) LAN uplinks configured in a channel
- Q4) What is considered a good practice when defining the MAC addresses for ESXi deployment? (Source: Evaluating Cisco UCS Integration with VMware vSphere)
- A) Manually assign MAC address per each vNIC.
 - B) Use derived MAC addresses.
 - C) Use a default MAC pool for MAC assignment.
 - D) Randomize MAC addresses for better security.
 - E) Incorporate administratively defined identifiers in MAC addresses.
- Q5) How can the administrator automate ESXi server deployment when using service profiles? (Source: Evaluating Cisco UCS Integration with VMware vSphere)
- A) Clone the first ESXi service profile.
 - B) Use hardware-based addresses for vNICs and vHBAs.
 - C) Populate server pools using qualification policy.
 - D) Use a scrub policy.
 - E) Use the basic service profile deployment model.

- Q6) Which two VLANs must be defined for proper VSM-to-VEM communication?
(Choose two.) (Source: Evaluating Cisco UCS and Nexus 1000V Integration with VMware vSphere)
- A) service console
 - B) vMotion
 - C) packet
 - D) FT logging
 - E) control
- Q7) How can a deployment of 50 VMware ESXi hosts with similar characteristics be optimized from the perspective of deployment time? (Source: Evaluating Cisco UCS and Nexus 1000V Integration with VMware vSphere)
- A) deriving service profiles from service profile template
 - B) creating base ESXi service profile and cloning the profile
 - C) running multiple Cisco UCS Manager instances for configuration implementation
 - D) applying single service profile to server pool
 - E) prebooting server blades with ESXi
- Q8) What should be implemented in order to separate Fibre Channel traffic for accessing boot and data LUNs ? (Source: Evaluating Cisco UCS and Nexus 1000V Integration with VMware vSphere)
- A) Use VSAN trunking on vHBA.
 - B) Use Fibre Channel port channeling on Fibre Channel uplink interfaces.
 - C) Implement SAN pin groups.
 - D) Use vHBA templates.
 - E) Use operating system-level multipathing software.

Module Self-Check Answer Key

- Q1) B
- Q2) C
- Q3) D
- Q4) E
- Q5) C
- Q6) C, E
- Q7) A
- Q8) C

