Table of Contents

Volume 1

Course Introduction 1

Overview 1
Course Goal and Objectives 3
Course Flow 4
Additional References 5

Understanding the Benefits of the Cisco Unified Computing System ***1-1***

Overview 1-1
Module Objectives 1-1

Implementing Cisco Virtualized Data Centers **1-3**

Overview 1-3
Objectives 1-3
Evolution of Server Scalability 1-4
Management Challenges 1-10
Blade Challenges 1-13
Summary 1-15

The Cisco Unified Computing System **1-17**

Overview 1-17
Objectives 1-17
Logical View of UCS 1-18
I/O Consolidation 1-23
Summary 1-29
Module Summary 1-31
Module Self-Check 1-33

Understanding the Cisco Unified Computing System ***2-1***

Overview 2-1
Module Objectives 2-1

Defining Cisco UCS Hardware Components **2-3**

Overview 2-3
Objectives 2-3
Cisco UCS Hardware Overview 2-4
Cisco UCS Chassis Hardware 2-6
Chassis 2-6
Power 2-6
Cooling 2-6
Cisco UCS Feature Support 2-7
Chassis 2-8
Power 2-8
I/O Module 2-9
CMC 2-9
Cisco UCS Feature Support 2-9
The Midplane 2-10
Cisco UCS Fabric Interconnect 2-11
Fabric Interconnect 2-11
Expansion Modules 2-12
Cisco UCS Blades 2-14
Hard Disk Drives 2-20
UCS Feature Support 2-20
Mezzanine Cards 2-21
Cisco UCS VIC M81KR 2-21
Cisco UCS CNA M71KR 2-21
Cisco UCS 82598KR-CI 2-22
Comparison 2-22
Describe VMware Integration 2-24

https://t.me/learningnets

Summary 2-27
Architectural and Feature Overview 2-29
Overview 2-29
Objectives 2-29
Cisco UCS Management Architecture 2-30
Chassis Management Controller 2-33
Cisco UCS Management Assignments 2-35
Diagnostic Information 2-39
Management Solution Integration 2-43
Cisco UCS Manager High Availability and Disaster Recovery 2-47
Summary 2-49
Installing Cisco UCS Hardware 2-51
Overview 2-51
Objectives 2-51
Installation Guidelines 2-52
Chassis Weight 2-53
Module Weights 2-53
Configuring Chassis Power Supplies 2-56
Chassis Module Installation 2-59
Summary 2-64
Module Summary 2-65
Module Self-Check 2-67
Understanding Cisco UCS Connectivity 3-1
Overview 3-1
Module Objectives 3-1
Describing Unified Connectivity 3-3
Overview 3-3
Objectives 3-3
Hardware Component Connectivity 3-4
Cluster Interfaces 3-6
IOM Management 3-7
High Availability Connectivity 3-10
High Availability Requirements 3-10
Fabric A/B Datacenter Operating System (DCOS) Networks 3-11
Fabric A/B Adapter Management (AM) Networks 3-11
Fabric A/B Adapter Management (AM) Infrastructure Networks 3-11
Fabric A/B Cisco UCS Manager PXE Networks 3-12
Fabric Interconnect Network Configuration 3-14
Summary 3-23
Exploring the Unified Computing User Interfaces 3-25
Overview 3-25
Objectives 3-25
The Graphical User Interface 3-26
Command-Line Interface 3-33
The connect Commands 3-41
Summary 3-45
Describe Cisco UCS Startup and Shutdown Procedures 3-47
Overview 3-47
Objectives 3-47
Cisco UCS 6100 Startup and Shutdown 3-48
Summary 3-55

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Upgrading Software in the Cisco Unified Computing System 3-57
Overview 3-57
Objectives 3-57
Cisco UCS Software Upgrade Overview 3-58
Components to Upgrade 3-58
Methods of Upgrade 3-66
Downloading Images 3-59
Kernel 3-60
System 3-60
Backup 3-60
Startup 3-60
Ignoring Compatibility Checks 3-61
Resetting on Activation 3-61
On the File Server 3-63
From the Cisco UCS Manager CL 3-63
Upgrading Cisco UCS Fabric Interconnect Software 3-66
Upgrading the Fabric Interconnect Software using the Cisco UCS Manager GUI 3-68
Upgrading CMC Software 3-69
Upgrading BMC Software 3-72
Upgrading Adapter Software 3-76
Summary 3-81
Module Summary 3-83
Module Self-Check 3-85

Understanding the Cisco UCS Manager 4-1

Overview 4-1
Module Objectives 4-1

Cisco UCS Manager Fundamentals 4-3

Overview 4-3
Objectives 4-3
Cisco UCS Resources 4-4
Resource Pools 4-9
Service Profiles 4-14
Service Profile Policies 4-19
Service Profiles in Basic Opt-In Model 4-21
Service Profiles in Logical Server Model 4-23
Summary 4-27

Role-Based Access Control 4-29

Overview 4-29
Objectives 4-29
Management Groups 4-30
RBAC Overview 4-37
RBAC Configuration 4-46
Remote RBAC Configuration 4-53
Multi-Tenancy Features and Options 4-57
AAA Configuration 4-60
Authentication 4-60
Authorization 4-60
Accounting 4-60
Summary 4-68
Module Summary 4-69
Module Self-Check 4-71

Establishing Compute Node Connectivity 5-1

Overview 5-1
Module Objectives 5-1

https://t.me/learningnets

Compute Node LAN Connectivity 5-3
Overview 5-3
Objectives 5-3
Server and Uplink Ports 5-4
Ethernet Host Virtualizer Mode 5-6
Ethernet Switching Mode 5-12
Describe the use of VLANs 5-15
Virtual NICs 5-21
Mezzanine Cards 5-21
Cisco UCS VIC M81KR 5-21
Cisco UCS CNA M71KR 5-21
Cisco UCS 82598KR-CI 5-22
Comparison 5-22
Port Channels 5-27
Pin Groups 5-30
LAN Tools 5-33
Summary 5-37
Compute Node SAN Connectivity 5-39
Overview 5-39
Objectives 5-39
N Port Identifier Virtualization 5-40
Server Interfaces 5-42
Border Interfaces 5-42
VSANs 5-47
Virtual HBAs 5-54
Mezzanine Cards 5-54
Cisco UCS VIC M81KR 5-54
Cisco UCS CNA M71KR 5-54
Cisco UCS 82598KR-CI 5-55
Comparison 5-55
SAN Port Connectivity 5-60
Pin Groups 5-63
Summary 5-64
Module Summary 5-65
Module Self-Check 5-67
Implementation of Server Resources 6-1
Overview 6-1
Module Objectives 6-1
Identity and Resource Pools 6-3
Overview 6-3
Objectives 6-3
Mobile Service Profiles 6-4
MAC Addresses 6-6
World Wide Names 6-10
Universally Unique Identifiers 6-18
Server Pools 6-22
Summary 6-29
Creating Service Profiles 6-31
Overview 6-31
Objectives 6-31
Service Profile Review 6-32
Opt-In Models 6-34
Service Profile Wizard 6-36
Service Profile Expert Wizard 6-40
Profile Templates and Clones 6-55
Scrub Policies 6-60
Summary 6-61

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Managing Service Profiles **6-63**
Overview 6-63
Objectives 6-63
Associating Service Profiles 6-64
Moving a Logical Server 6-70
KVM and Virtual Media 6-72
Summary 6-81
Module Summary 6-83
Module Self-Check 6-85

https://t.me/learningnets

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

DCUCI

Course Introduction

Overview

This 5-day hands-on course focuses on Cisco Unified Computing System (UCS) deployment and operations. You will learn how to configure and manage Cisco UCS servers with consolidated I/O networking for LAN and SAN connectivity, and how to virtualize server properties to enable simple and rapid mobility of server images between physical servers.

In the lab, you will practice implementing a multi-tenant management model, configuring fault tolerance, backing up and restoring system configurations, and using the built-in monitoring and troubleshooting tools in UCS Manager. You will also install and configure the Cisco Nexus 1000V Virtual Switch on the UCS infrastructure.

https://t.me/learningnets

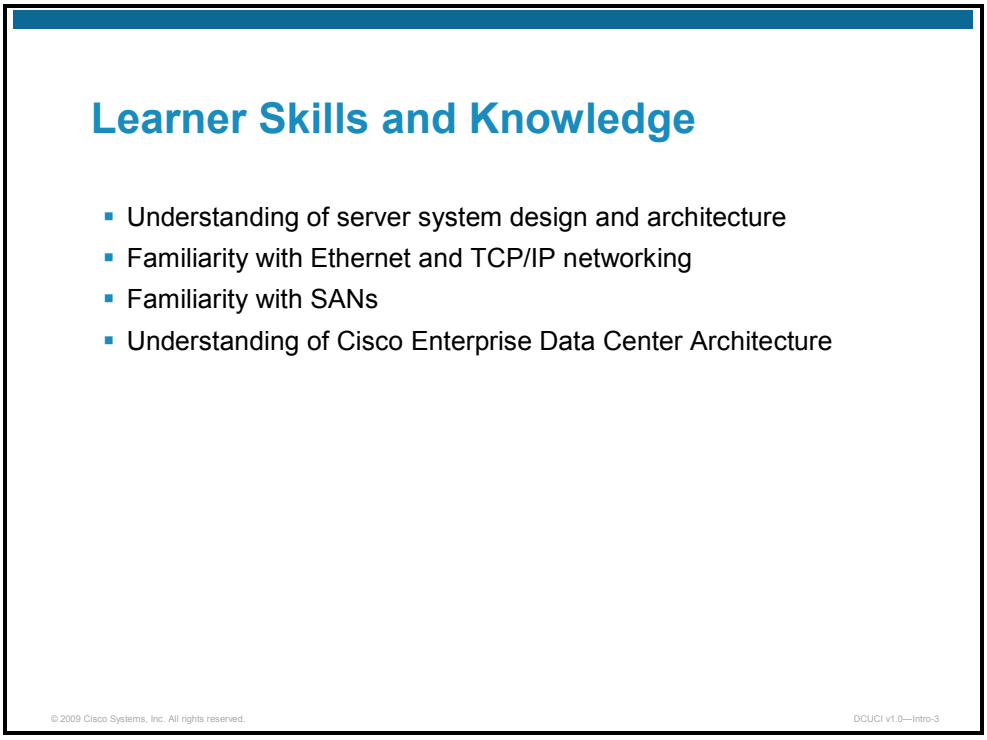

The figure lists the skills and knowledge that you should possess to benefit fully from the course.

https://t.me/learningnets

Course Goal and Objectives

This topic describes the course goal and objectives.

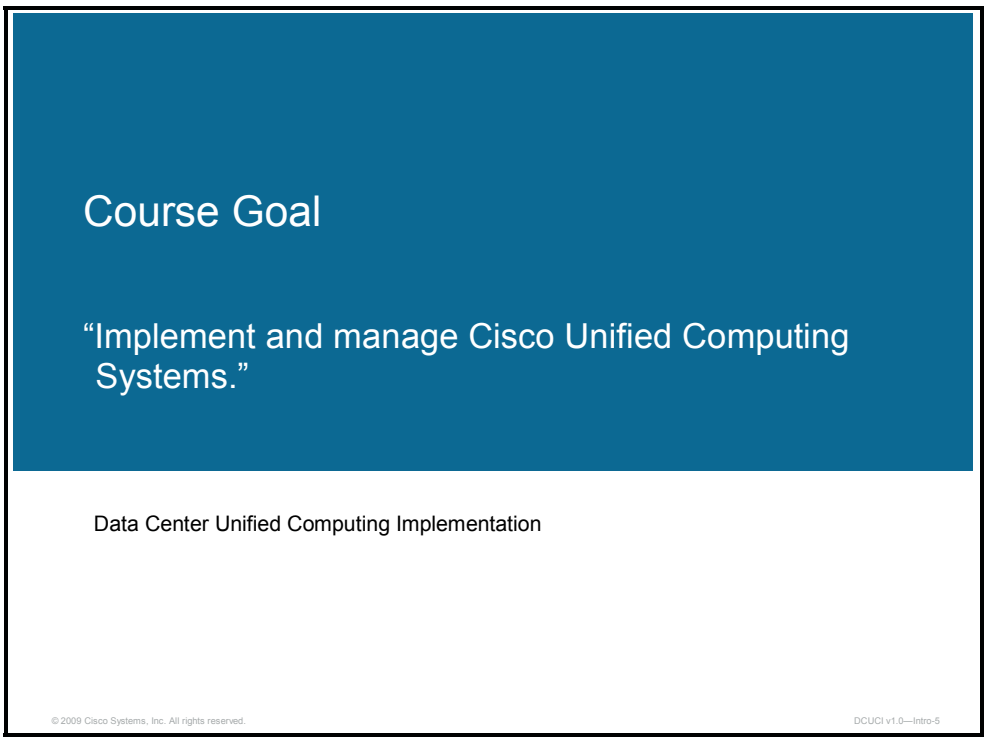

Upon completing this course, you will be able to implement and manage a Cisco UCS deployment.

https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.

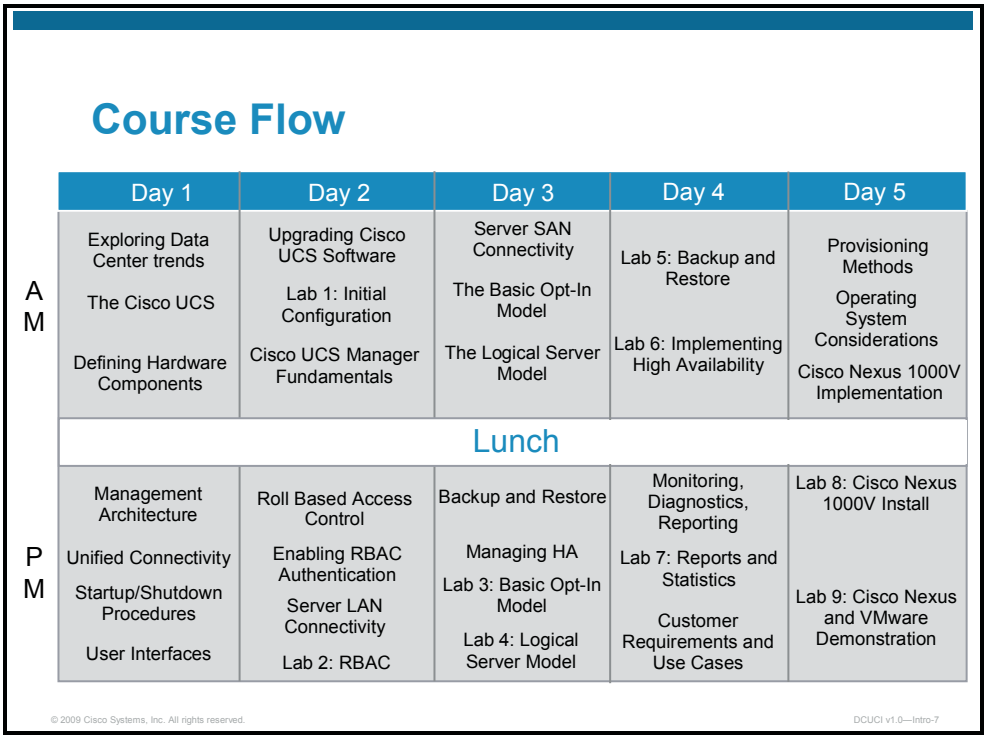

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Additional References

This topic presents the Cisco icons and symbols that are used in this course.

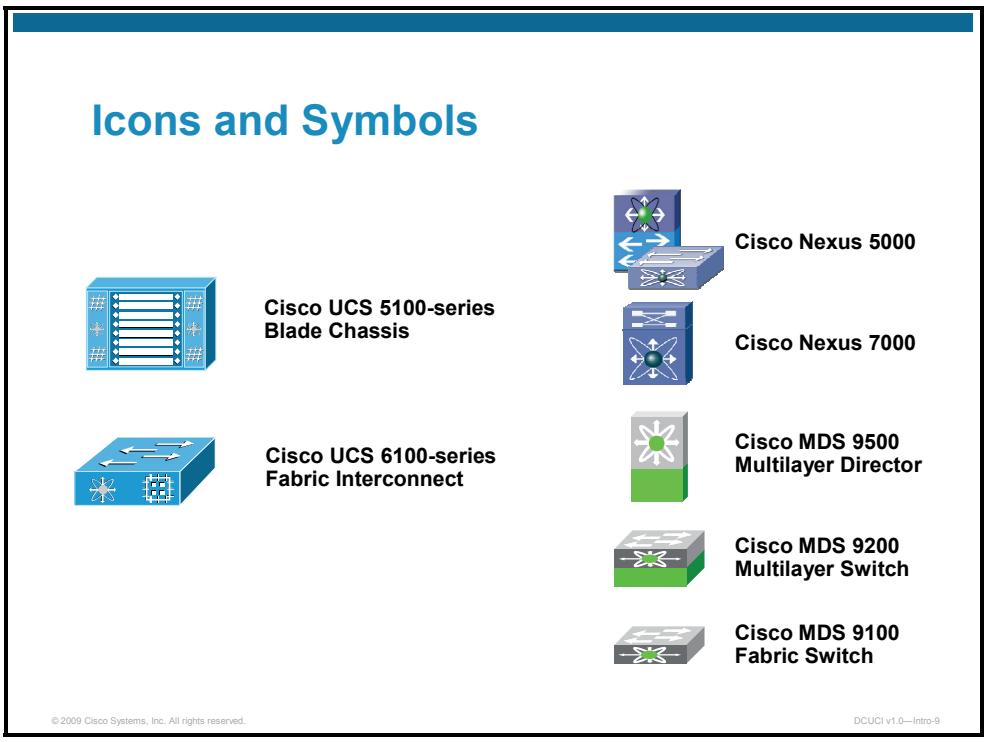

These are the icons and symbols that you will see throughout this course.

https://t.me/learningnets

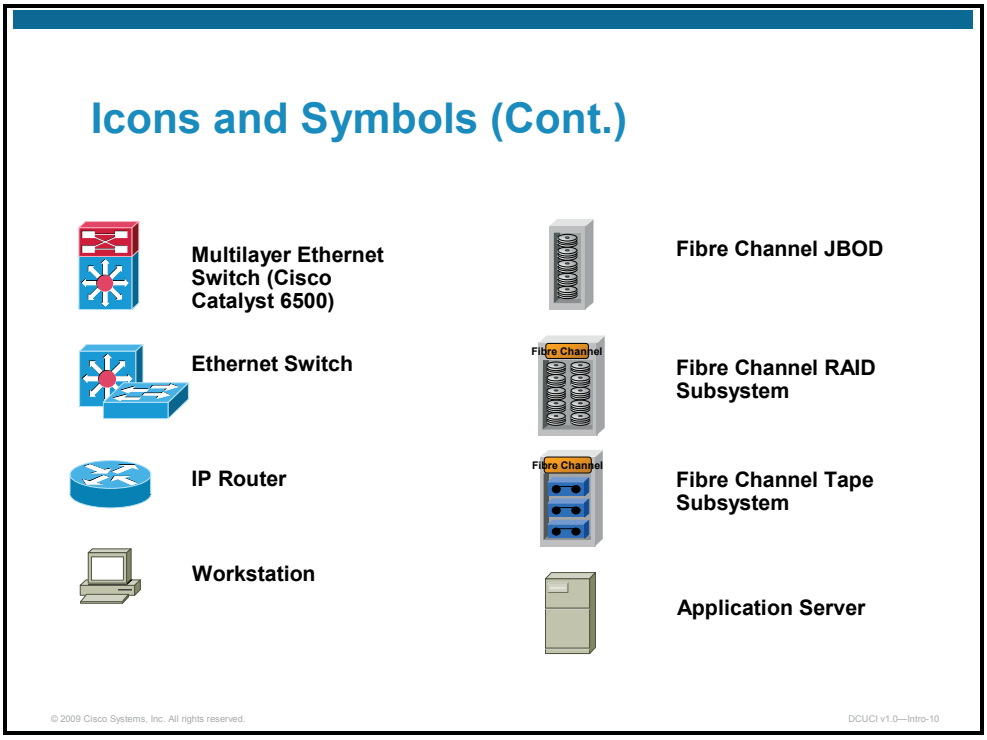

These are the icons and symbols that you will see throughout this course.

https://t.me/learningnets

Module 1

Understanding the Benefits of the Cisco Unified Computing System

Overview

This module illustrates the need for the Cisco Unified Computing System (Cisco UCS) and discusses optimum placement.

Module Objectives

Upon completing this module, you will be able to illustrate the need for Cisco UCS and discuss optimum placement. This includes being able to meet these objectives:

- Explain how to implement Cisco virtualized data centers
- Describe the benefits of the Cisco Unified Computing System

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Explaining Challenges with Data Center Server Management

Overview

This lesson explains the challenges that server, network, and storage administrators face when managing large server farms in their data centers.

Objectives

Upon completing this lesson, you will be able to explain the challenges that administrators face when managing data center servers. This includes being able to meet these objectives:

- Describe the evolution of server scalability
- Describe the management challenges that exist in modern data centers
- Describe the challenges faced by traditional blade server deployments

https://t.me/learningnets

Evolution of Server Scalability

This topic describes the evolution of server scalability.

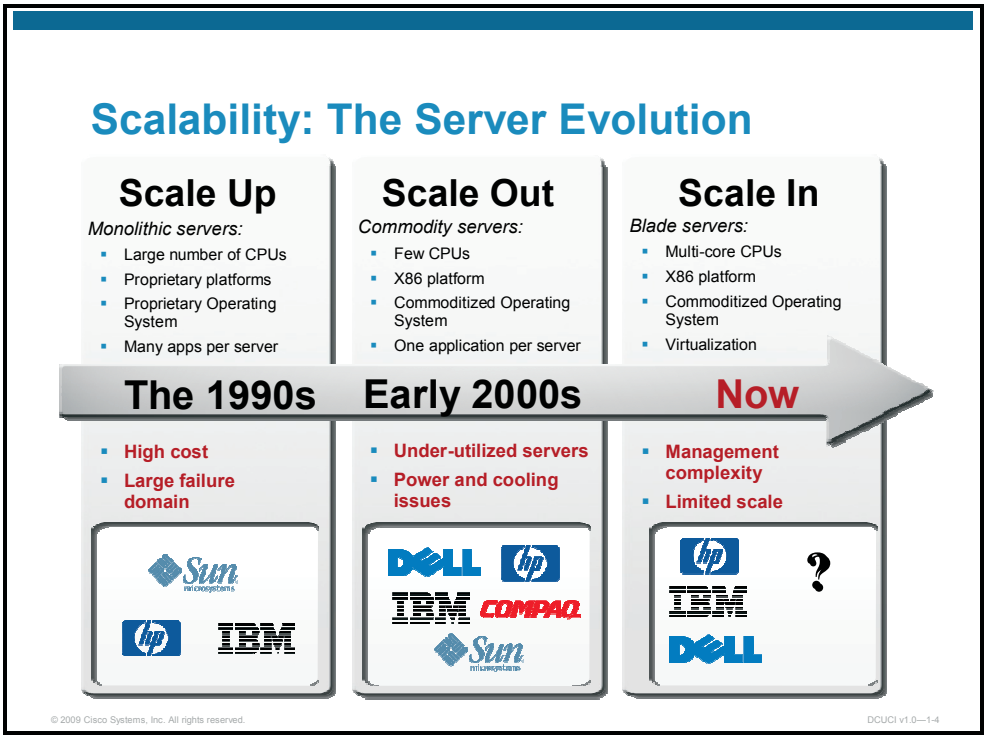

The first servers were designed as large, monolithic enclosures. At that time, the components were very expensive, so a single large system could service a large number of users more cost effectively than many smaller servers.

The next change in servers came with commoditized servers. These were smaller servers that cost less than the monolithic server, but they also had fewer processing resources, typically limiting them to a single service or application per physical server. In addition, reliability concerns further drove a one-to-one server to application architecture.

Today, the push for servers is to reduce costs by increasing efficiency. Servers are being designed into smaller form factors and utilizing shared resources. Additionally, with the introduction of virtualization software, many services are being consolidated onto single physical computing resources. While this aids in many of the challenges faced by the data center, other challenges are created.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

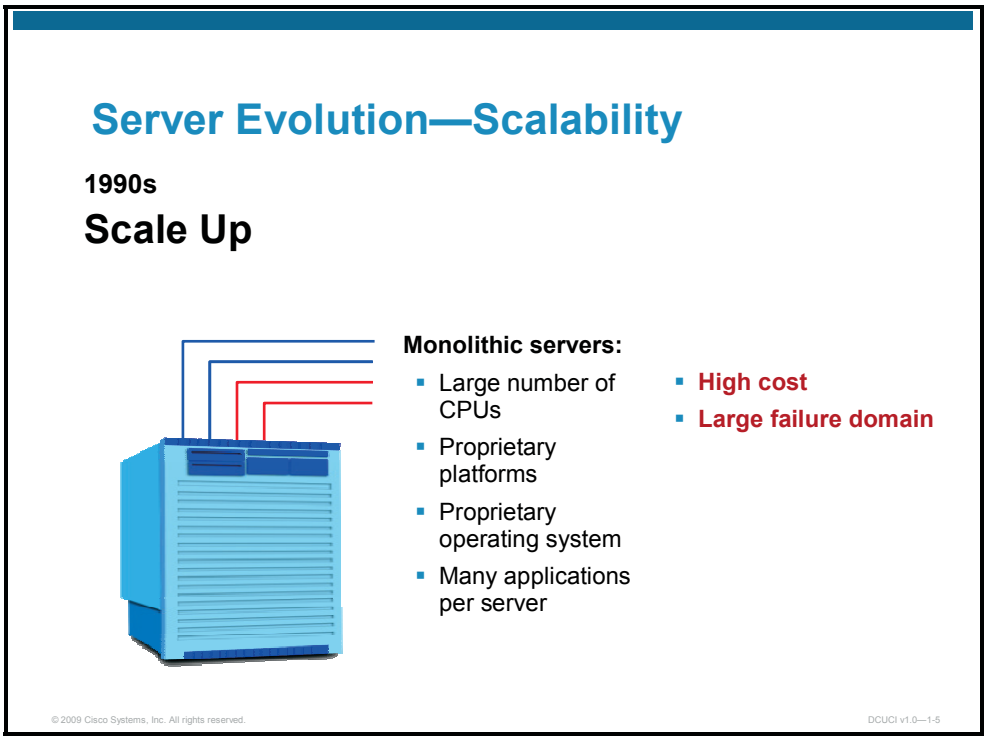

The figure describes the server scalability model through the 1990s. At that time, compute resources were typically scaled by adding more compute power to a small number of servers. The result was higher total cost of ownership (TCO) and large failure domains.

https://t.me/learningnets

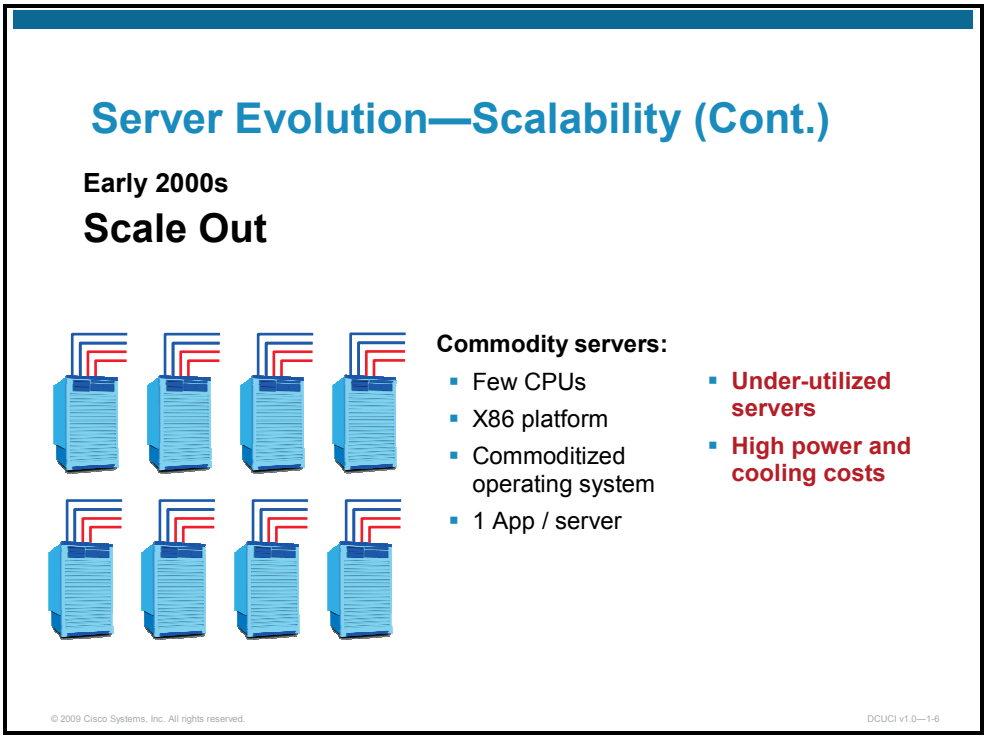

The figure describes the server scalability model during the early 2000s. During this time, the scalability model shifted away from fewer, larger servers to more numerous, smaller servers.

This transition was primarily driven by three factors:

- A shift from centralized IT control of assets to decentralized, departmental control
- The development of clustering technologies that allowed applications to scale more efficiently across multiple servers
- Increasing standardization on the Intel x86 platform to reduce cost of ownership

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

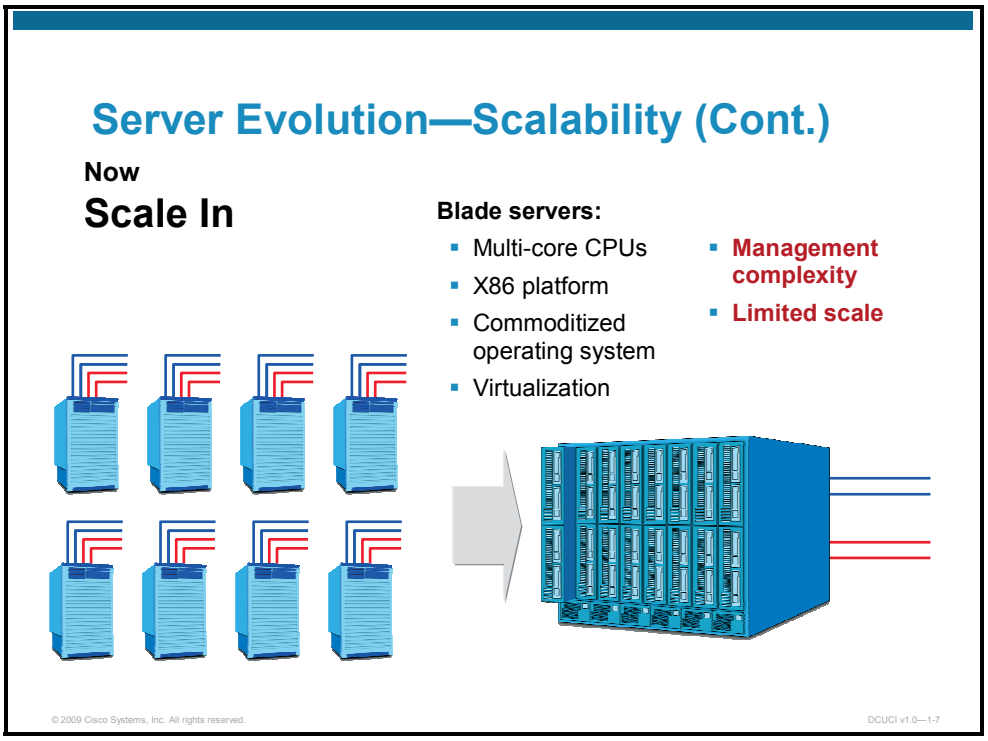

Since the early 2000s, IT has begun to move away from traditional tower or rack-mounted servers in favor of blade servers. Larger servers required more space and needed to have their own power supply, networking, and SAN cables. For smaller data centers, these infrastructure requirements are not problematic, but in a large data center these requirements result in increasing complexity and cost of ownership.

By putting the computing resources of a server in a blade form factor within an enclosure, administrators gain the following advantages:

- Maximum use of physical space
- Shared power distribution
- Shared networking and storage access
- More efficient power and cooling

https://t.me/learningnets

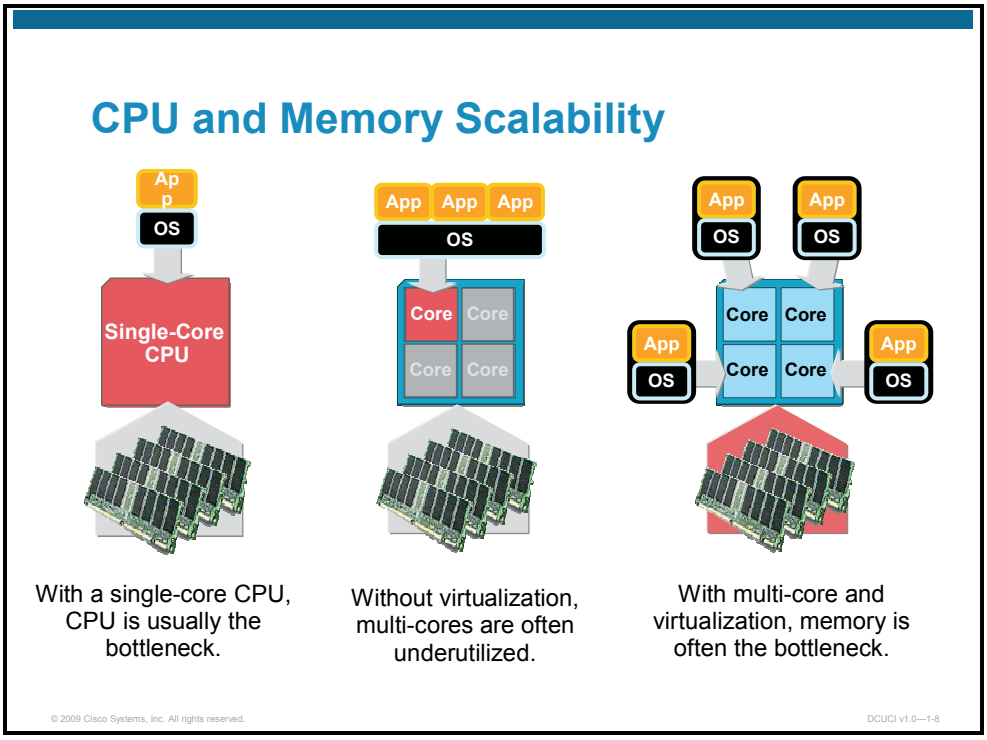

Today, server processors are multicore, offering multiple processing cores in the same space that was previously used for single-core processors.

In addition, many vendors have created memory controllers that have the ability to address large amounts of RAM into the tens of gigabytes. While blade server sizes have remained the same, the processing and memory has increased significantly, allowing blade servers to run processing and memory-intensive applications.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

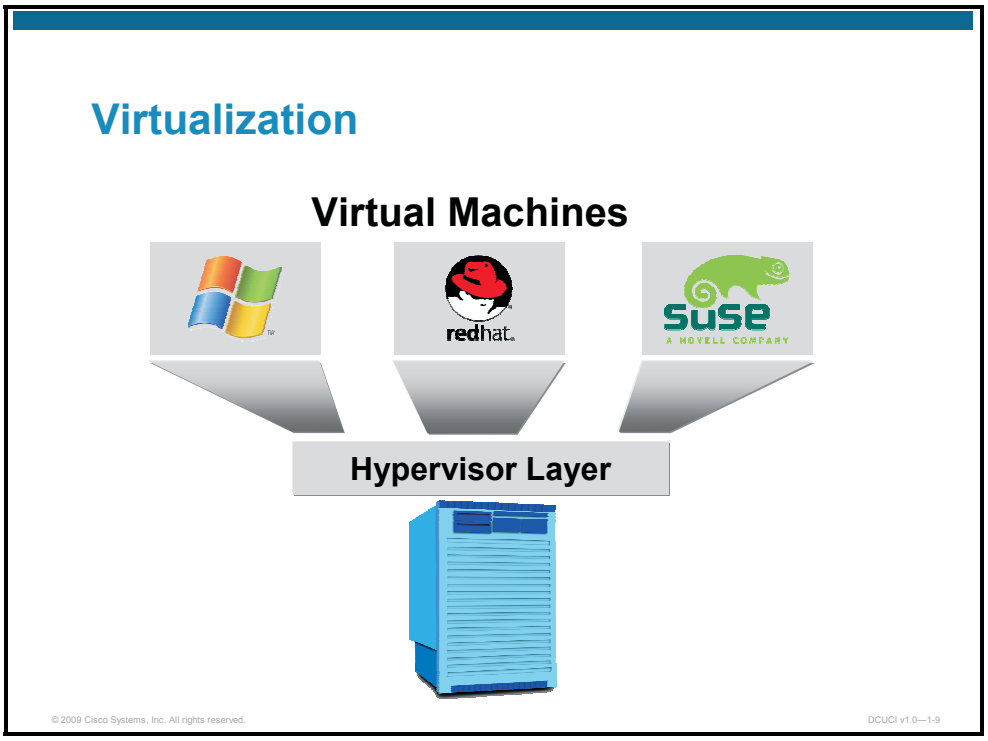

Virtualization has allowed companies to more easily consolidate servers in the data center. Instead of the standard one-to-one server model, many servers—each running in independent "virtual machines" (VMs), can run on a single physical server.

This is performed by creating individual virtual hardware that can function like a standard server. There are a number of advantages to using virtualization, including better use of computing resources, greater server densities, and seamless server migrations:

- **Virtual machine:** A virtualized set of hardware that is able to operate in a similar fashion to a physical server
- **Virtual server:** A virtual set of hardware along with the operating system, applications, and files that is able to operate comparable to a physical server
- **Hypervisor layer:** A software layer that abstracts the physical hardware and creates individual virtual hardware for each VM. For example, VMware ESX, Microsoft Hyper-V, Xen, etc.

https://t.me/learningnets

Management Challenges

This topic describes the management challenges that exist in modern data centers.

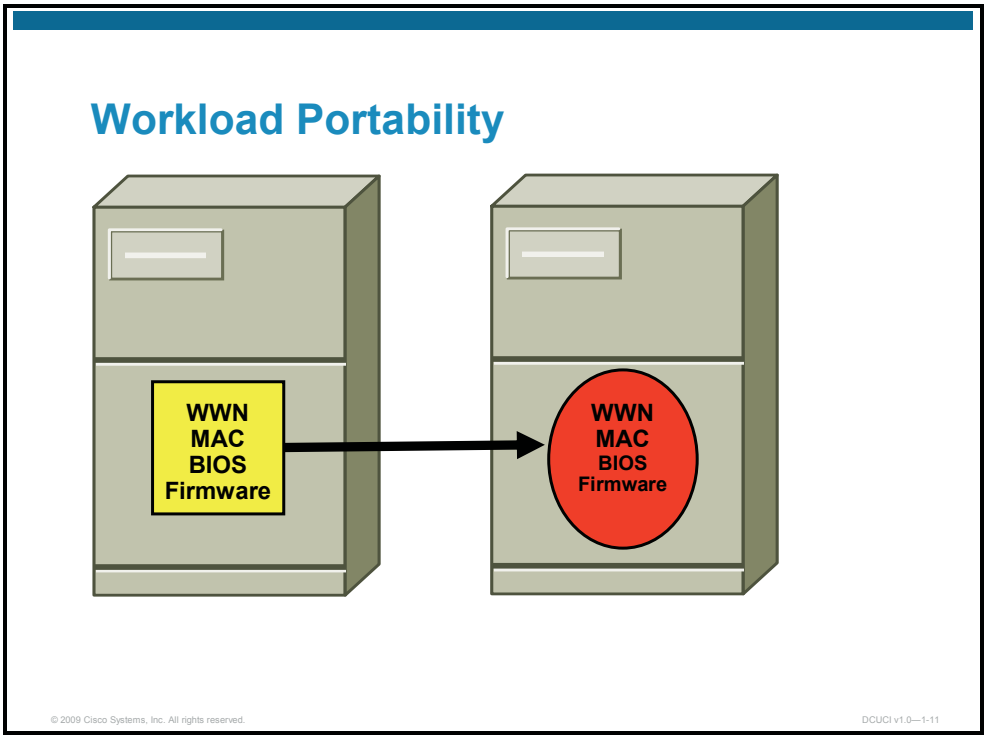

Servers have unique identifiers that identify it on various networks. These identifiers are tied to the hardware. If you change any of these items, the server will potentially lose its ability to access network resources, or even boot an operating system. These identifiers include:

- **World wide name (WWN):** Hard coded to a Host Bus Adapter (HBA), this identifier is needed for SAN access.
- **MAC address:** Hard coded to a network interface card (NIC), this identifier is needed for LAN access.
- **BIOS:** This identifier contains settings that are specific to the server hardware.
- **Firmware:** This low-level software runs on peripheral device and adapter cards to enable the operating system to interface with the device.

If a server fails and its operating system and application needs to be migrated to another physical server, the operating system/application and the network may require manual configuration changes. These manual configuration changes lead to longer recovery times and increased application downtime.

The alternative to manual configuration changes is the use of "stateless servers."

- **Stateless server:** A stateless server is a server whose operating system and application personality have no ties to the physical hardware. One way in which to accomplish this task is to use transportable, virtual, unique identifiers.
- **Server personality:** A server personality is the operating system configuration and application settings. This is a fully functional set of programs, files, and settings required to perform the given task of the server.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

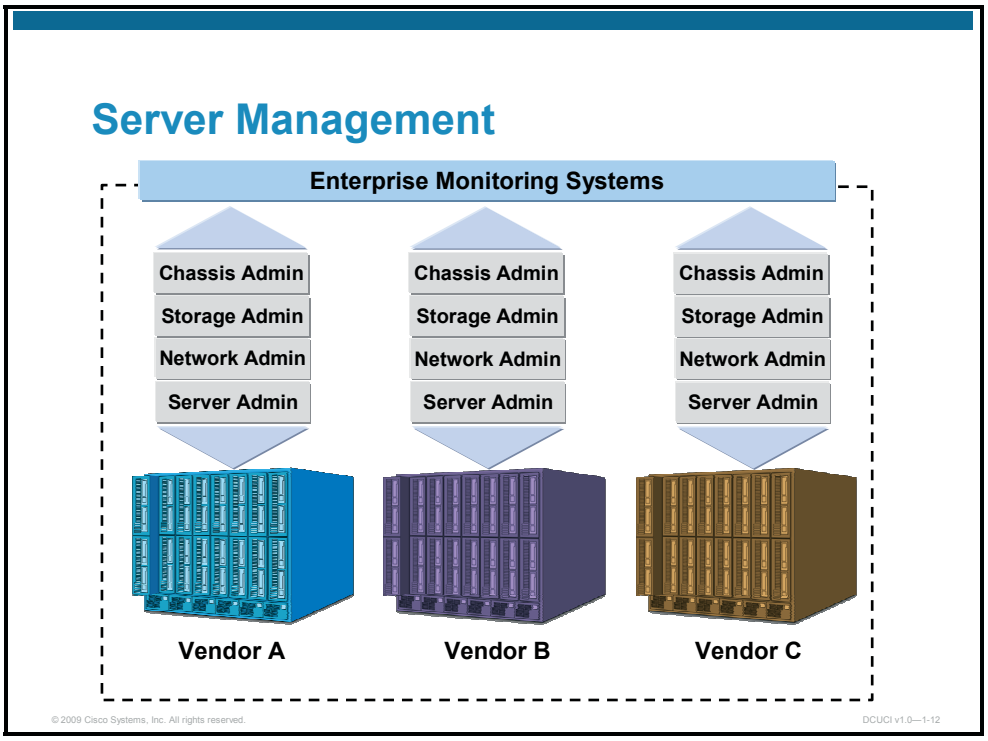

As computing capacity increases within the data center, so does complexity. Blade servers solve many issues but they also bring an additional point of management into the datacenter per chassis. Several independent systems must be managed including LAN, SAN, servers, and storage. These separate resources must be managed at each network layer—access, aggregation, and core. Typically, these resources are managed by individual teams and may be monitored using proprietary system-monitoring tools and alert aggregators. In many cases, customers must use multiple monitoring applications to cover all aspects of the data center.

https://t.me/learningnets

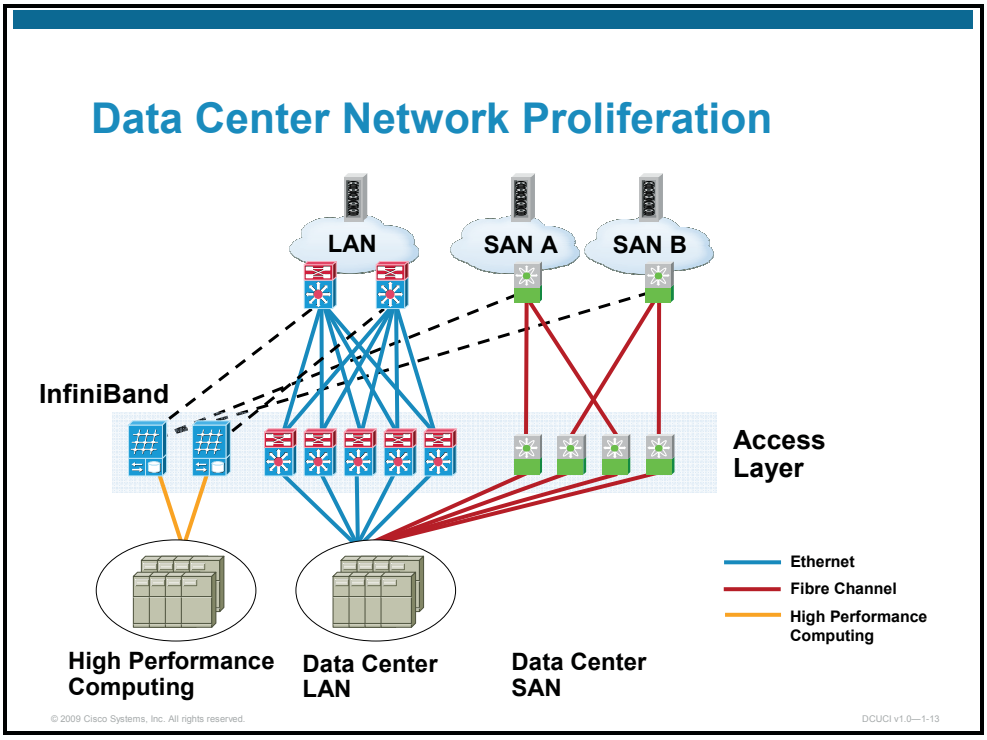

The data center consists of multiple networks. In addition to the LAN and the WAN, 20 percent to 40 percent of all servers are also connected to a SAN over Fibre Channel. The use of high performance computing is also growing; for example, with financial trading applications that demands very low latency. This leads to more complexity with regard to cabling, power, cooling, and management.

The increase in server quantity has resulted in a similar increase in network requirements. As application demand expands and new servers are implemented, the networks must grow to meet this demand. These networks include LAN, SAN, and High Performance Computing (HPC). This not only results in increased power and cooling needs but also increased cabling costs and management challenges.

Networking needs vary greatly depending on the operating system and application but can reach as high as eight 1-GE ports and four 4-G Fibre Channel ports for a virtualization platform. This is up to 12 cables per physical server and can add up to 252 cables for a rack filled to capacity with 2U servers.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Blade Challenges

This topic describes challenges faced by traditional blade server deployments.

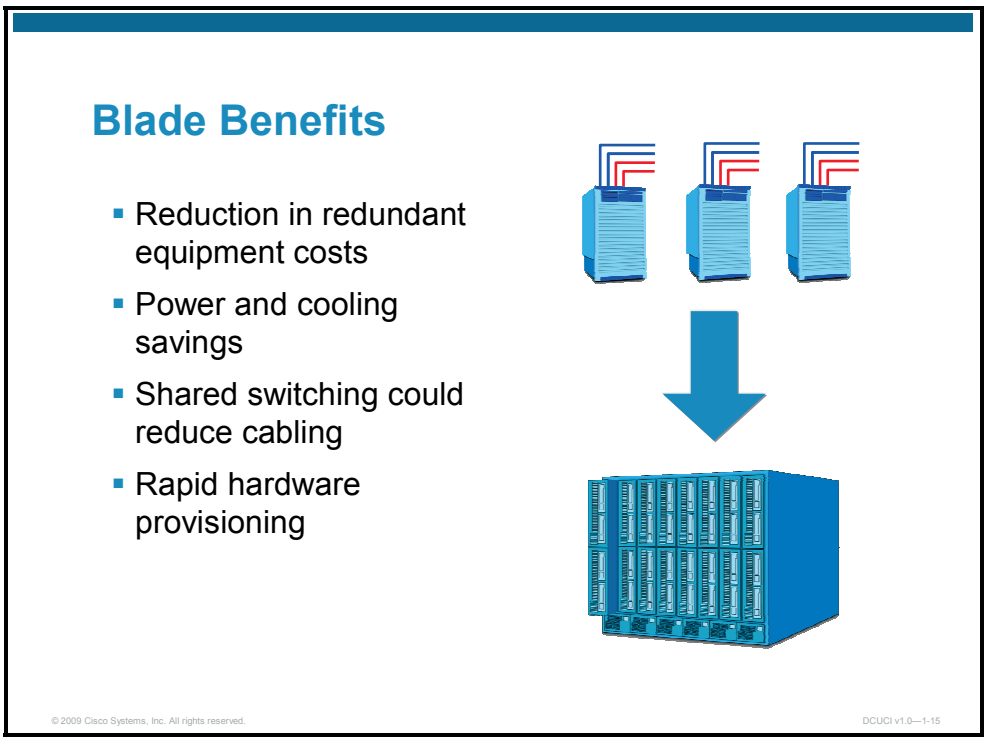

The consolidation of discreet servers into blade chassis provides a number of advantages in the data center:

- Reduced physical footprint
- Shared networking and SAN switching
- Reduced cabling
- Rapid provisioning of additional resources

https://t.me/learningnets

Blade Challenges

- Increased physical density may create power and cooling challenges
- Increased compute density could increase bandwidth and cabling requirements
- Each chassis and local switching creates additional management points

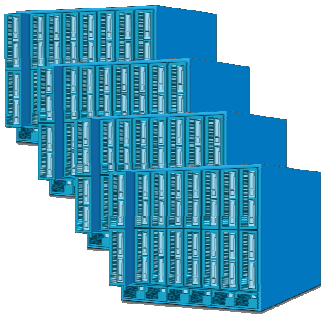

© 2009 Cisco Systems, Inc. All rights reserved.

As the physical density of a blade environment increases relative to rack mounted servers, there may be more power and cooling required in the same physical space. This can cause challenges in powering and cooling the data center.

The relatively high density of compute resources in a blade environment often requires significantly more bandwidth and cabling to fully utilize the deployment.

Each chassis in a blade environment will typically have a management IP address and require monitoring for system and blade health. Additionally, each local LAN or SAN switch (if installed) also adds management overhead and overall network design challenges.

https://t.me/learningnets

Summary

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

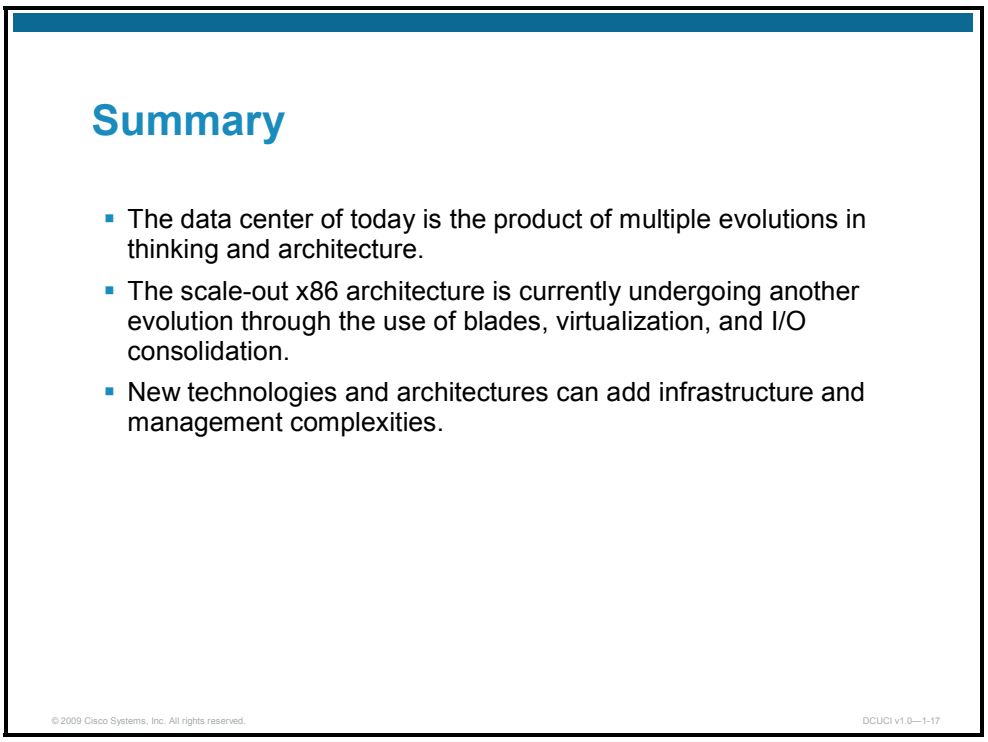

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

The Cisco Unified Computing System

Overview

This lesson describes the benefits of the Cisco Unified Computing System (Cisco UCS).

Objectives

Upon completing this lesson, you will be able to describe the benefits of the Cisco Unified Computing System. This includes being able to meet these objectives:

- Describe the logical view of the Cisco UCS
- Describe the standards and processes for I/O consolidation

https://t.me/learningnets

Logical View of UCS

This topic describes the logical view of the Cisco UCS from various perspectives.

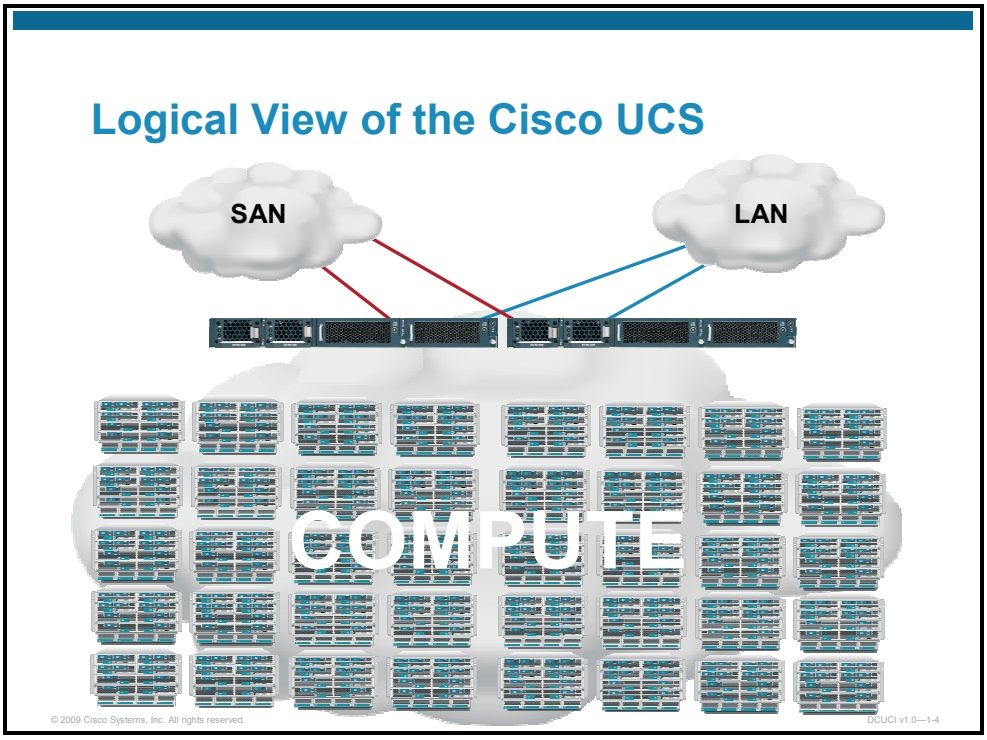

Two Cisco UCS 6100 Series Fabric Interconnects:

- Provides management for up to 320 blades and 40 chassis
- Provides consolidated I/O Fibre Channel over Ethernet (FCoE):
  - LAN and SAN on the same media infrastructure

Forty Cisco UCS 5108 Blade Server Chassis:

- Two Cisco UCS 2104XP Fabric Extenders (aka FEX) per chassis
- One to eight blade servers per chassis

Note The Cisco UCS Fabric Interconnects have uplink connectivity to both the local area network (LAN) and SAN networks. It is also important to note that support for up to 40 chassis requires use of the Cisco UCS 6140XP 40-Port Fabric Interconnect. In a high-availability configuration using two Cisco UCS Fabric Interconnects, both Fabric Interconnects can be configured to actively switch traffic simultaneously.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

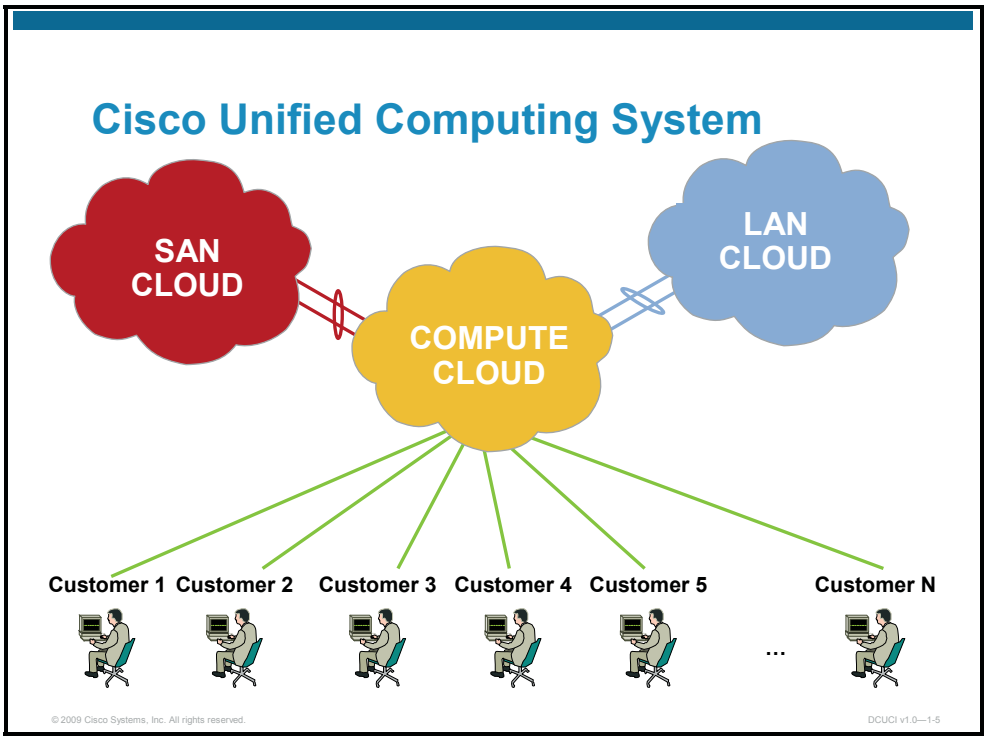

The Cisco UCS acts as a cloud of compute resources connected to existing LAN and SAN core infrastructures. This cloud is a cable once, dynamic environment that can be rapidly provisioned for new services.

https://t.me/learningnets

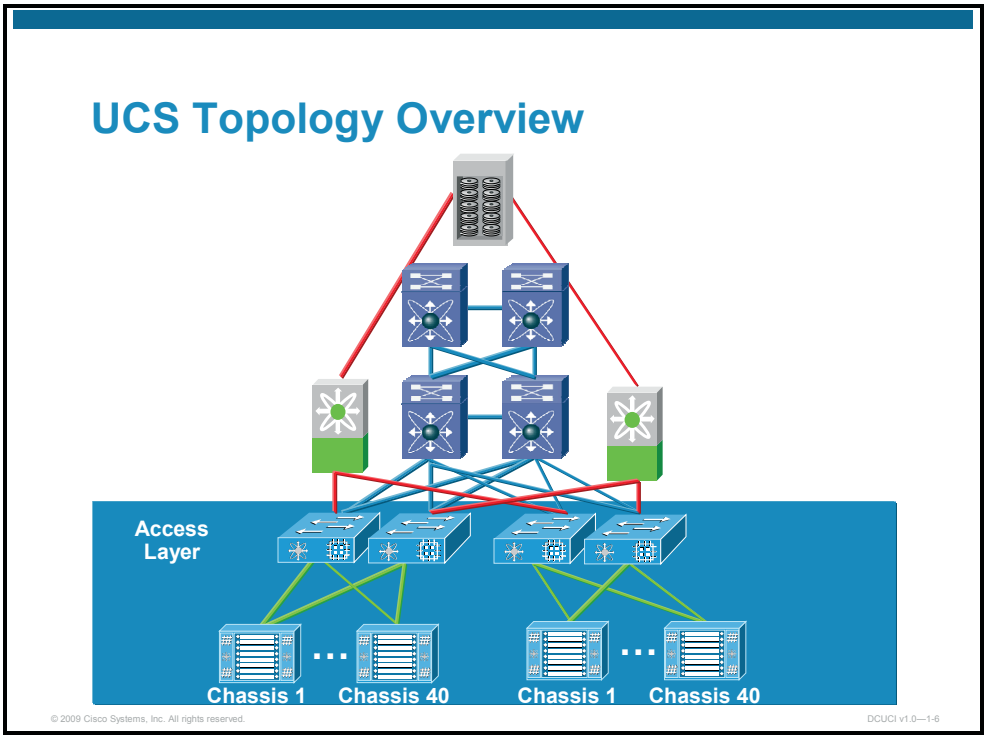

This figure shows how Cisco UCS is implemented at the access layer.

The deployment model of Cisco Unified Computing System is to replace the traditional server access layer with UCS-specific hardware. By replacing the access layer, integration with existing switching technologies is eased and management difficulties are reduced. The specific components that make up UCS are discussed later in this lesson.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

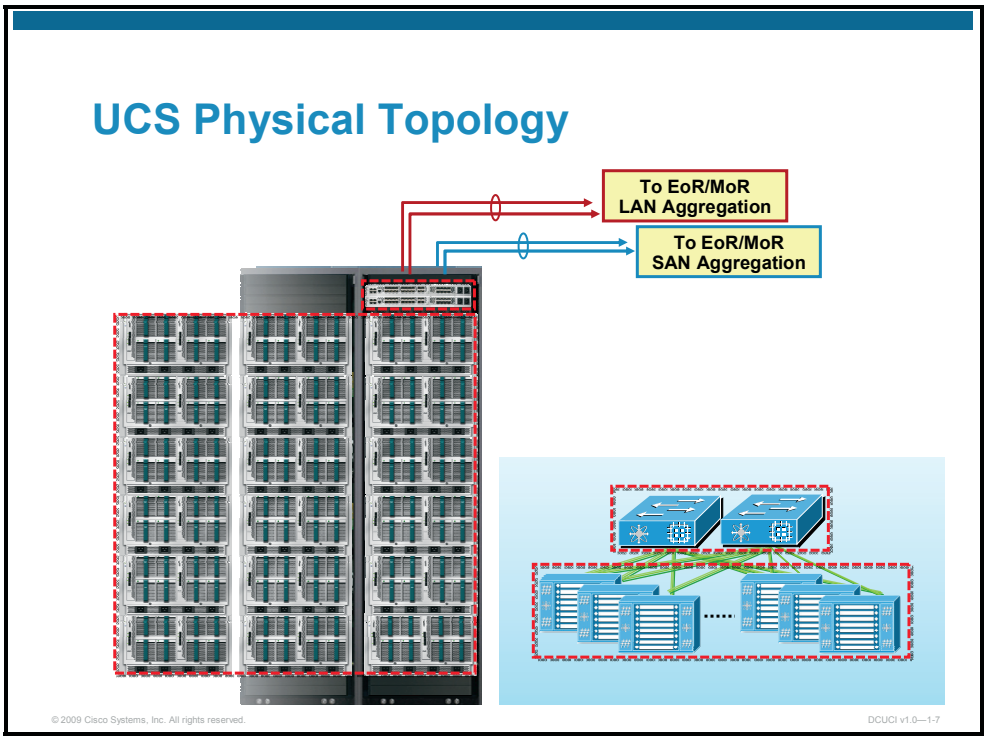

The figure shows a physical view of a UCS deployment, with a single pair of Fabric Interconnects supporting multiple racks of chassis.

https://t.me/learningnets

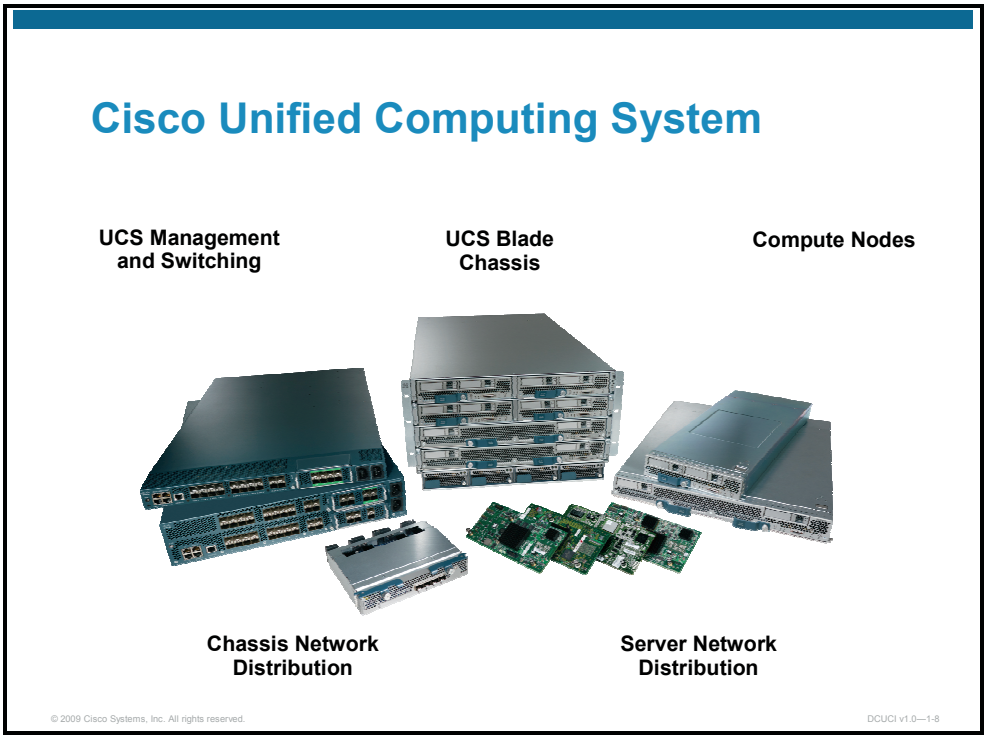

The figure shows the various components of the Cisco UCS.

UCS Management and Switching:

- UCS 6100-series Fabric Interconnects

UCS Blade Chassis:

- UCS 5100-series Blade Chassis

Compute Nodes:

- UCS B-series Blade Servers

Chassis Network Distribution:

- UCS 2100-series IO Modules

Server Network Distribution:

- M81KR Virtual Interface Card
- M71KR Converged Network Adapter
- 82598KR 10-GE Adapter

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

I/O Consolidation

This topic describes the standards and processes used for I/O consolidation.

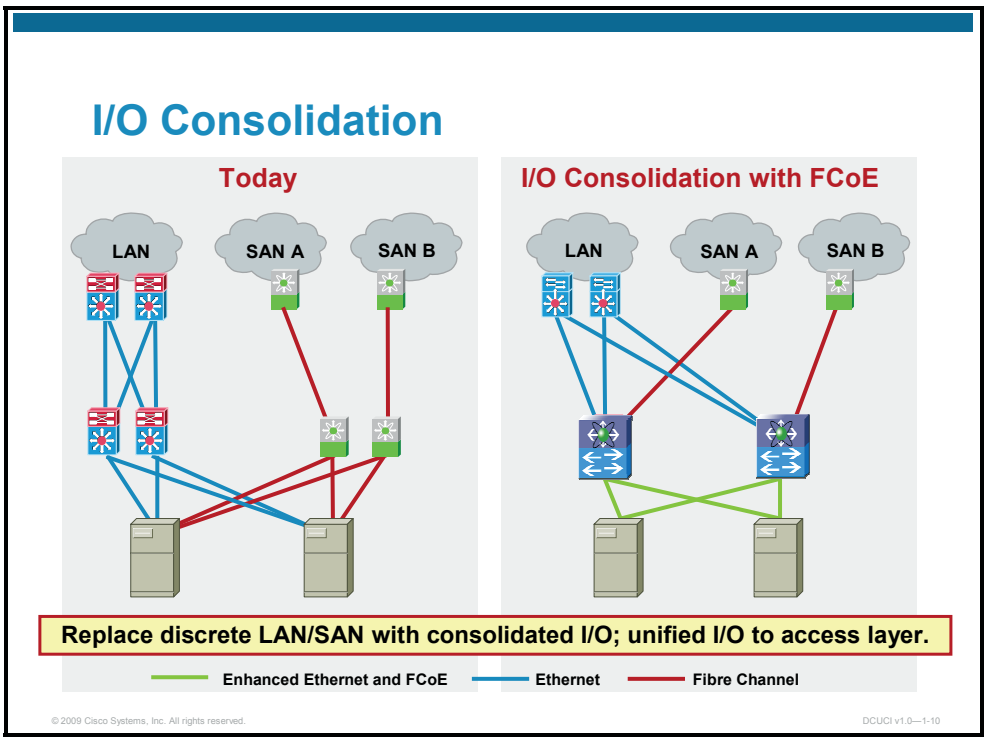

The increase in server quantity has resulted in a similar increase in network requirements. As application demand expands and new servers are implemented, the networks must grow to meet this demand. These networks include LAN, SAN, and High Performance Computing (HPC). This demand not only results in increased power and cooling needs but also increased cabling costs and management challenges.

Networking needs vary greatly depending on operating system and application but can reach as high as eight 1-GE ports and four 4-Gb Fibre Channel ports for a virtualization platform. This platform is up to 12 cables per physical server and it can add up to 252 cables per rack that is filled to capacity with 2U servers.

https://t.me/learningnets

Data Center Bridging Enhancements

- A collection of IEEE-based enhancements to classical Ethernet that provide end-to-end QoS
- Does not disrupt existing infrastructure

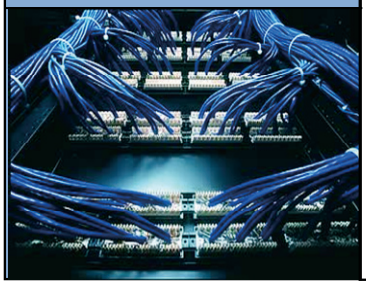

Ethernet enhancements:

- Priority groups: Virtualizes links and allocates resources per traffic classes
- Per Priority flow control
- End-to-end congestion management and notification
- Shortest-path bridging: Layer 2 multipathing

Benefits of Ethernet enhancements:

- Lossless fabric: No drop storage links
- Provides tools for handling network congestion
- Deterministic latency for HPC clusters
- Enables a converged Ethernet fabric for reduced cost and complexity

© 2009 Cisco Systems, Inc. All rights reserved.

The figure summarizes the enhancements to Ethernet that are required to achieve I/O consolidation. The traditional Ethernet specifications have some limitations when used within a data center for I/O consolidation. The main problem is a limited means for ensuring that frames are not lost. Enhanced Ethernet is sometimes called lossless when referring to its error handling and flow control, although it is not truly lossless.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

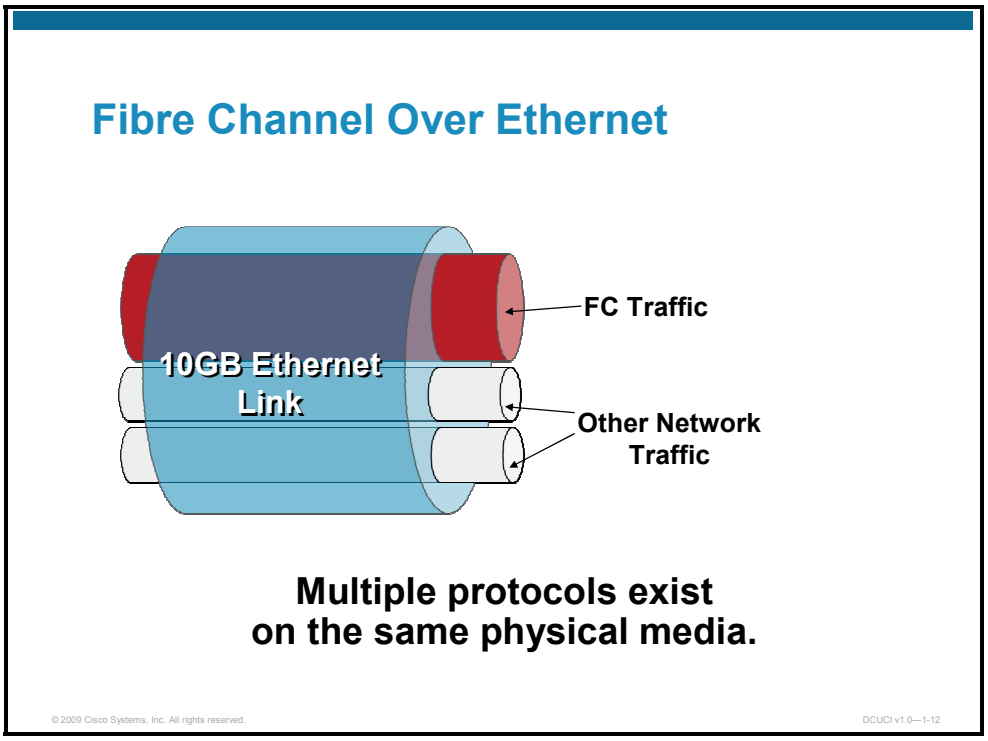

In an I/O consolidation environment separate protocols operate as tunnels on the same physical media. Fiber Channel traffic is tunneled in an Ethernet payload, while TCP/IP traffic exists separately inside Ethernet payloads.

https://t.me/learningnets

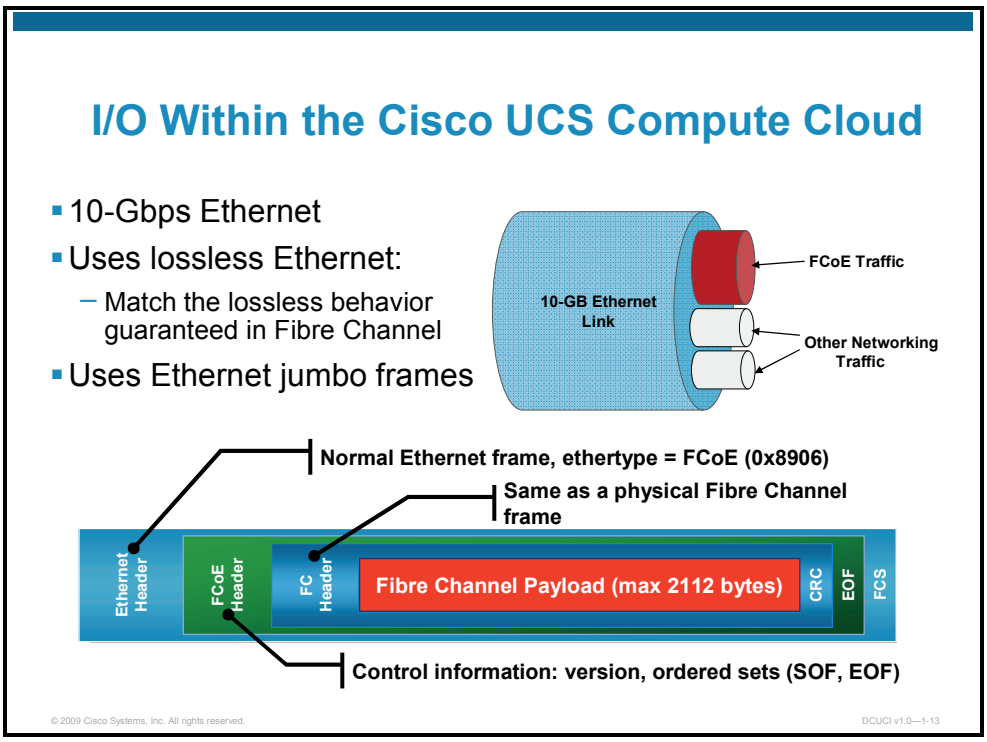

Effectively all Ethernet fabrics today utilize a form of packet-loss flow control. However, there is an IEEE 802.3X pause capability in standard Ethernet that can be configured to enable a lossless Ethernet fabric. There is also a priority-based pause capability that extends the pause functionality into logical lanes, allowing multiple lossless fabrics to exist on a single physical infrastructure. In the figure you can see that the Fibre Channel payload and Fibre Channel header are encapsulated in an FCoE header, which is further encapsulated in an Ethernet frame.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

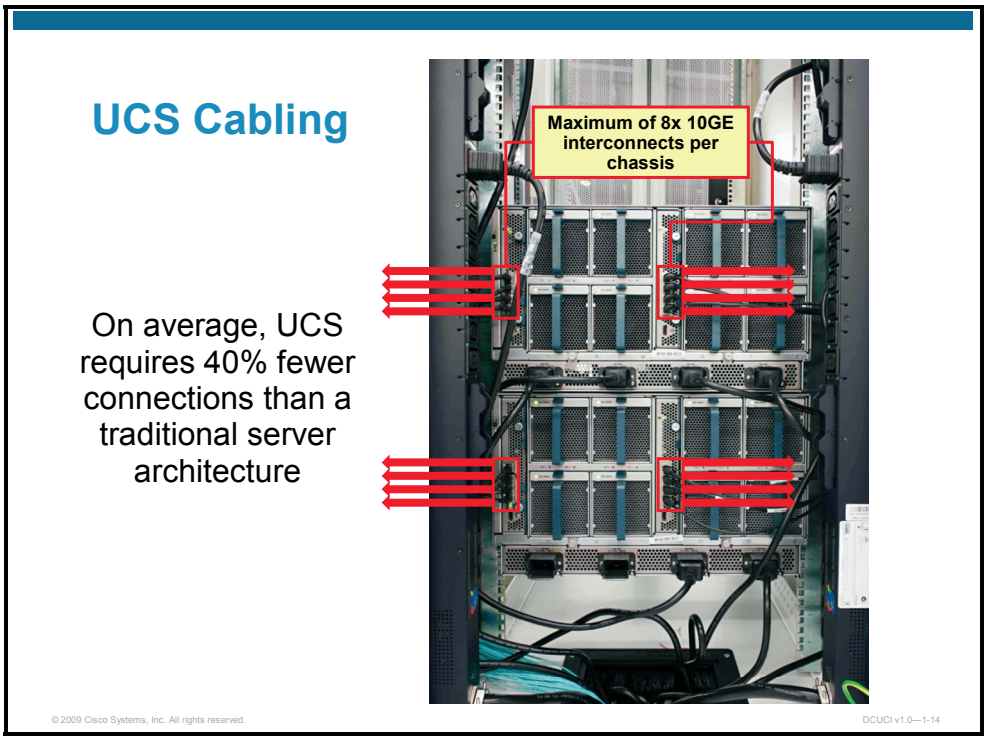

This is a picture of the UCS platform in production. UCS on average requires 40% fewer connections than a traditional server architecture and 66 percent fewer components, providing better airflow and less energy consumption. UCS also reduces the management points in a data center POD build out providing for reduce OpEx, simpler integration, automation and reduced provisioning times.

https://t.me/learningnets

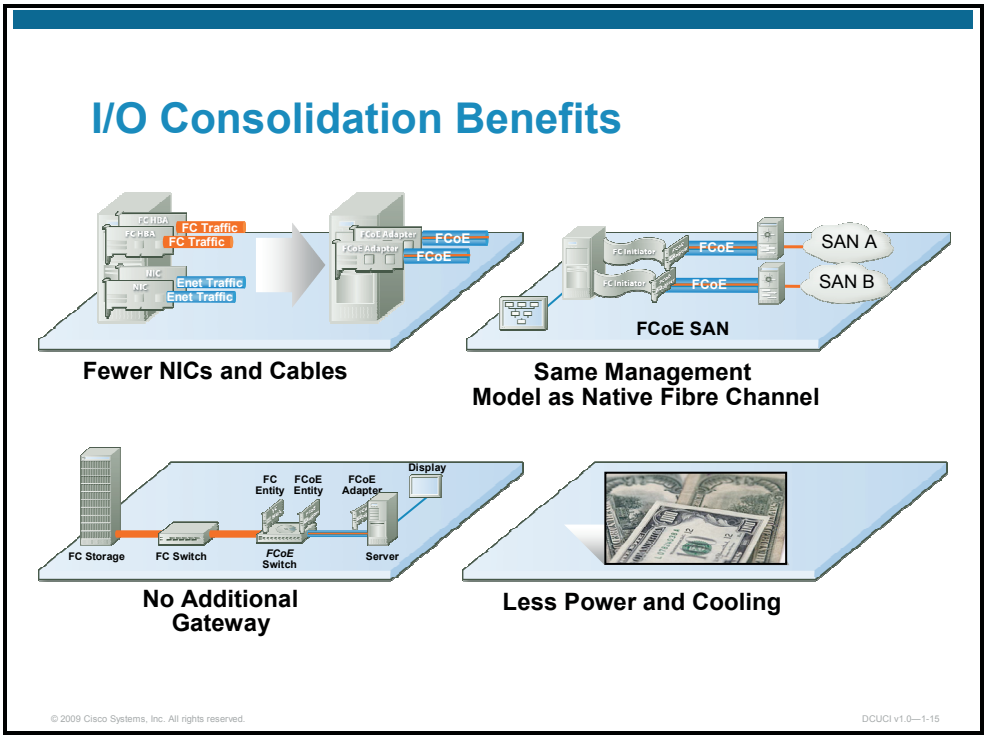

I/O consolidation using FCoE provides the following benefits:

- FCoE reduces the number of network interface cards (NICs) and cables that customers need to purchase, reducing cost and simplifying cable management.
- Because FCoE follows the same model as Fibre Channel, storage administrators will be able to manage their fabrics in the same manner they always have.
- FCoE does not have a stateful gateway but simply maps Fibre Channel traffic onto lossless Ethernet, resulting in performance benefits over technologies that require a gateway.
- Significant savings in the amount of power expended per rack results when FCoE is employed; by consolidating I/O, less power is expended on NICs and cables. Companies also save power on the cooling expended to compensate for the heat produced.
- Managed like Fibre Channel at initiator, target, and switch level.
- Extension of Fibre Channel onto lossless Ethernet fabric.
- Aligned with FC-BB-4 model, standardized in FC-BB-5.
- Easy to understand; completely based on Fibre Channel model:
  - Same operational model.
  - Same host-to-switch and switch-to-switch behavior.
  - Same techniques of traffic management.
  - Same management and security models.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

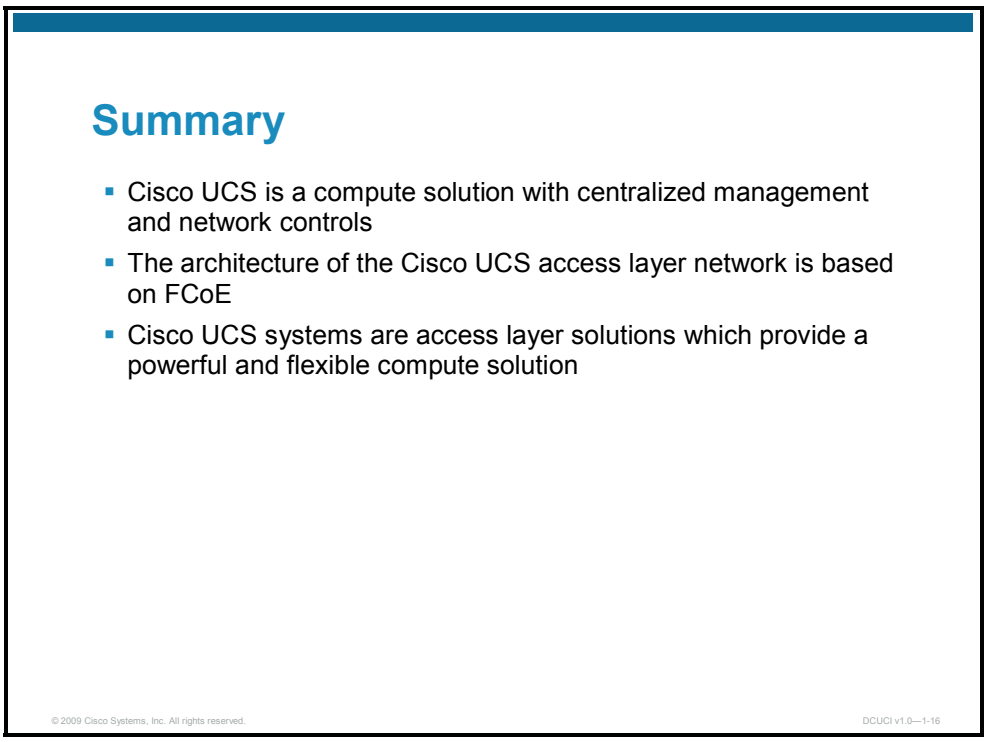

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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

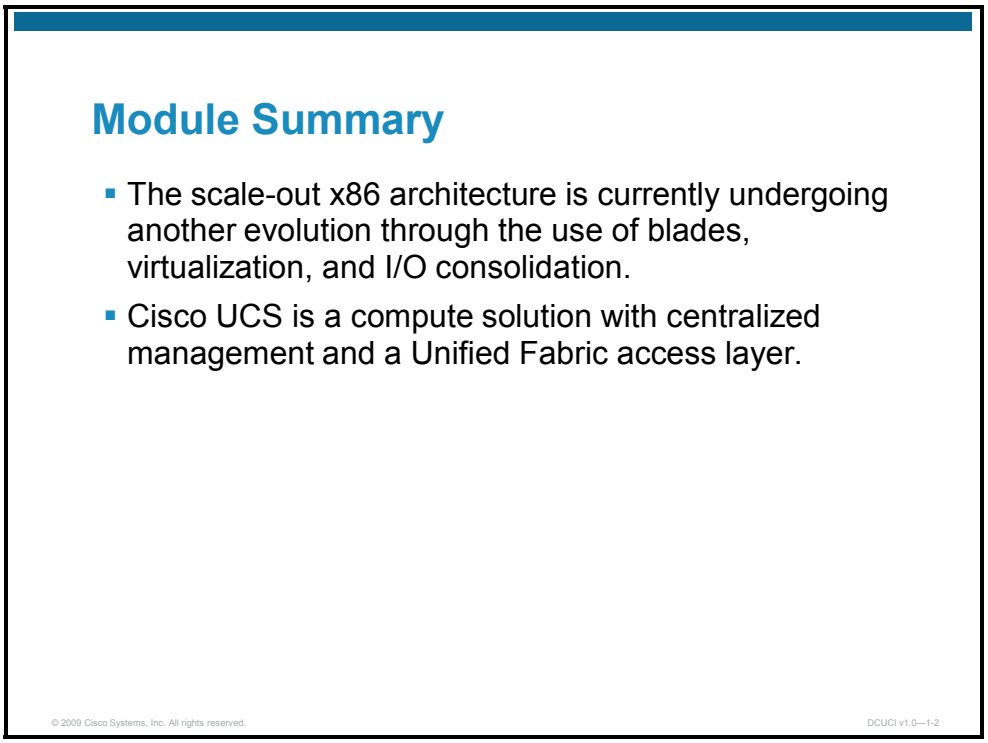

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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) What is typically the resource bottleneck in a virtualized environment? (Source:

A) Reviewing Data Center Server Trends)
B) CPU capacity
C) Storage
D) Power and cooling
E) Memory

Q2) Which of the following is a result of I/O consolidation? (Source: The Cisco Unified Computing System)

A) Reduced cabling
B) Increased host performance
C) More VLANs and VSANs
D) Increased power consumption

https://t.me/learningnets

Module Self-Check Answer Key

Q1) D

Q2) A

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 2

Understanding the Cisco Unified Computing System

Overview

This module describes the hardware, software components and management architecture of the Cisco Unified Computing System (Cisco UCS).

Module Objectives

Upon completing this module, you will be able to describe the hardware, software components, and management architecture of the Cisco Unified Computing System. This includes being able to meet these objectives:

- Define the hardware components of the Cisco Unified Computing System
- Discuss the architecture and feature set of the Cisco Unified Computing System
- Explain the process used to install and power up Cisco Unified Computing System hardware

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Defining Cisco UCS Hardware Components

Overview

In this lesson you will learn about the hardware components of the Cisco Unified Computing System (Cisco UCS) platform.

Objectives

Upon completing this lesson, you will be able to define the hardware components of the Cisco Unified Computing System. This includes being able to meet these objectives:

- Describe the major Cisco UCS hardware components
- Describe the Cisco UCS chassis hardware components
- Describe the Cisco UCS Fabric Interconnect models
- Describe the Cisco UCS Blade Servers and options
- Describe how VMware integrates into a Cisco UCS environment

https://t.me/learningnets

Cisco UCS Hardware Overview

This topic describes the major hardware components of the Cisco UCS.

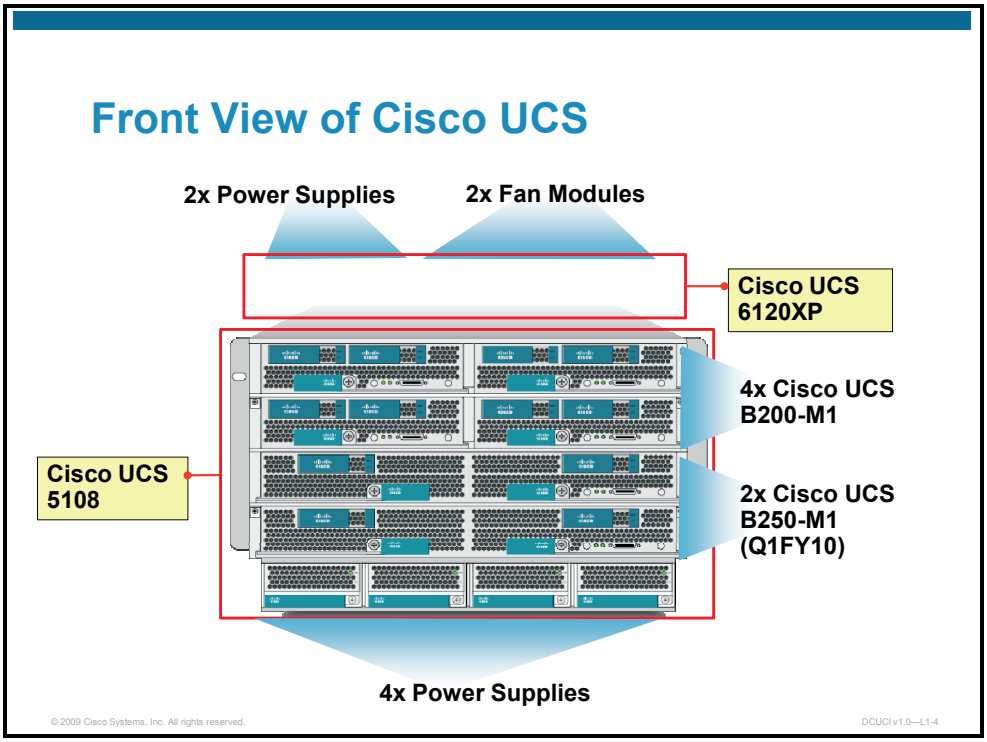

The power supplies for both the Cisco UCS 6100 Series Fabric Interconnect and Cisco UCS 5108 Blade Server Chassis are accessible from the front. Both half-width and full-width blades are supported in the same chassis. The fan modules for the Cisco UCS 6100 Fabric Interconnect are accessible from the front, while on the Cisco UCS 5100 they are accessible from the rear. Note that the recommended configuration would include two Fabric Interconnects; for clarity, only one is shown here.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

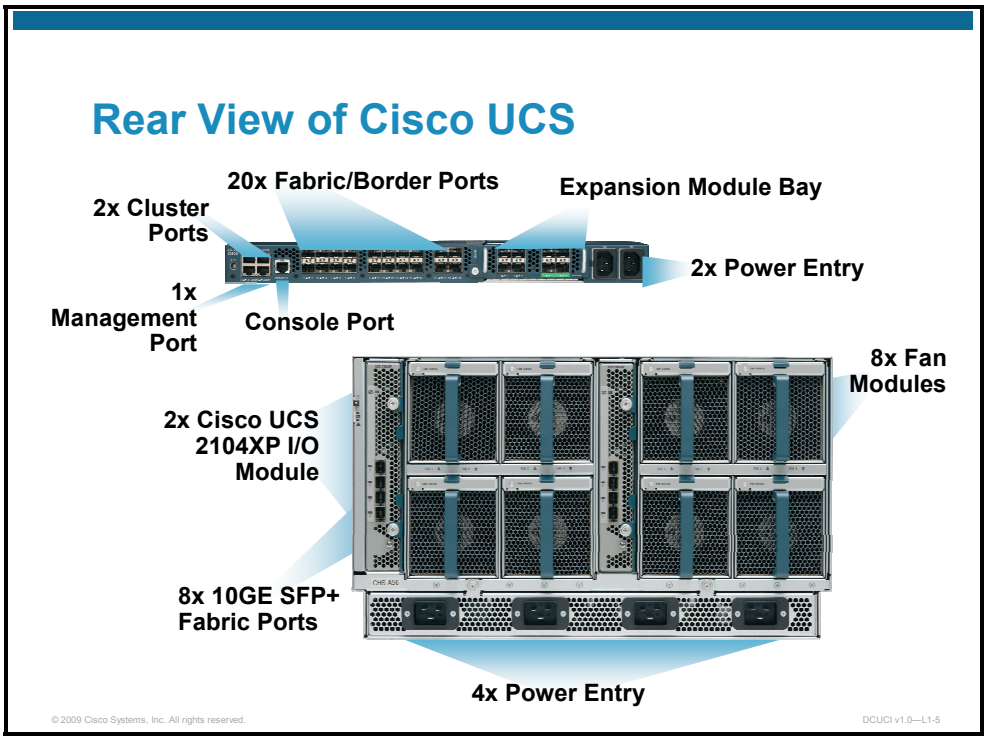

All the cabling for a Cisco UCS for both the Fabric Interconnect and the chassis (including power, Ethernet, console, and Fibre Channel) is plugged into the rear. Two Cisco UCS Fabric Interconnects are required for a fully-redundant configuration. In such a high-availability configuration, the left (as viewed from the rear of the chassis) 2104XP I/O module (IOM) connects to the left Fabric Interconnect and the right 2104XP IOM connects to the right Fabric Interconnect.

Note	The Fabric Interconnect shown in this figure is the Cisco UCS 6120XP 20-Port Fabric Interconnect.

https://t.me/learningnets

Cisco UCS Chassis Hardware

This topic describes the chassis of the Cisco UCS.

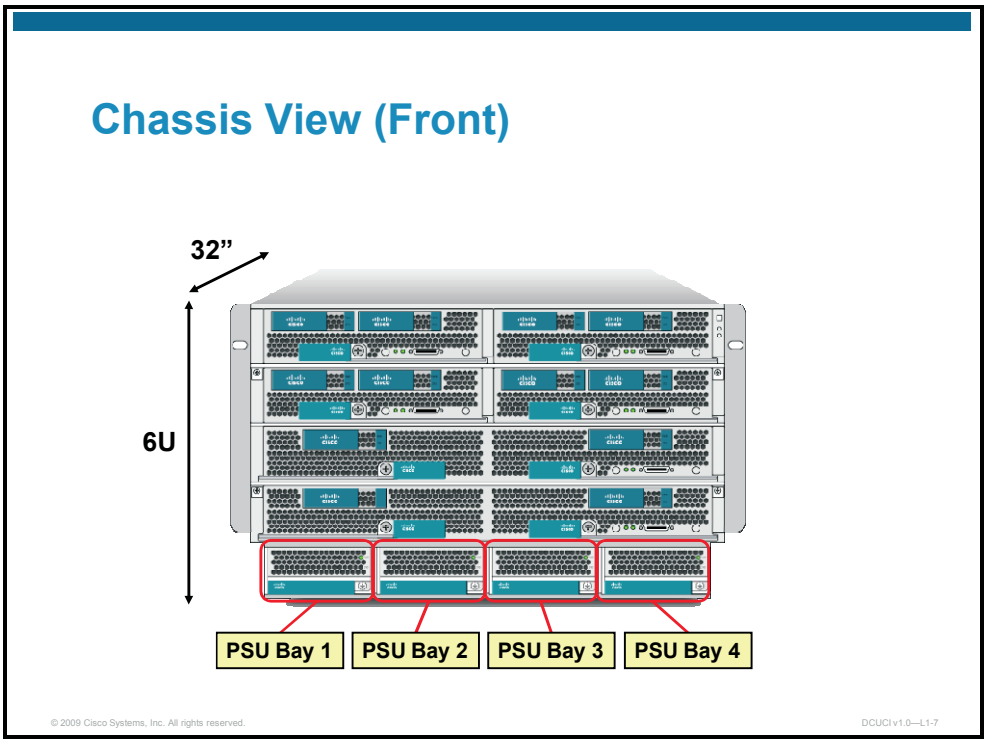

Chassis

The Cisco UCS 5108 Blade Server Chassis is 32 inches deep and 6 rack units (RU) tall. The blade servers and power supplies install in the front and lock into place using thumbscrews. The same chassis may be used for up to eight Cisco UCS B200-M1 2-Socket Blade Servers (half-width blade servers), four Cisco UCS B250-M1 2-Socket, Extended Memory Blade Servers (full-width blade servers), or any physically possible combination of half- and full-width blade servers. For the half-width blade servers, a vertical divider is used to separate horizontally adjacent blade servers. For the full-width blade servers, this divider must be removed.

Power

There are a maximum of four 2500W hot-pluggable power supplies per chassis. The actual number of power supplies required depends on the system’s hardware configuration and desired power redundancy. They can be configured as non-redundant, N+1, or N+N (grid) redundant. They are single-phase 220V, IEC320-C 19. The power supplies provide a 550W budget per slot for up to eight half-width blade servers and 1100W budget per slot for up to four full-width blade servers. This design supports future growth and power budget requirements with no service disruption.

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-plug dual fan modules per chassis with status indicators on each module.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Feature Support

Cisco UCS features supported by the chassis include the ability to scale up to as many as 20 chassis (using the 20-port Fabric Interconnect–Cisco UCS 6120XP), or 40 chassis (using the 40-port Fabric Interconnect–Cisco UCS 6140XP).

https://t.me/learningnets

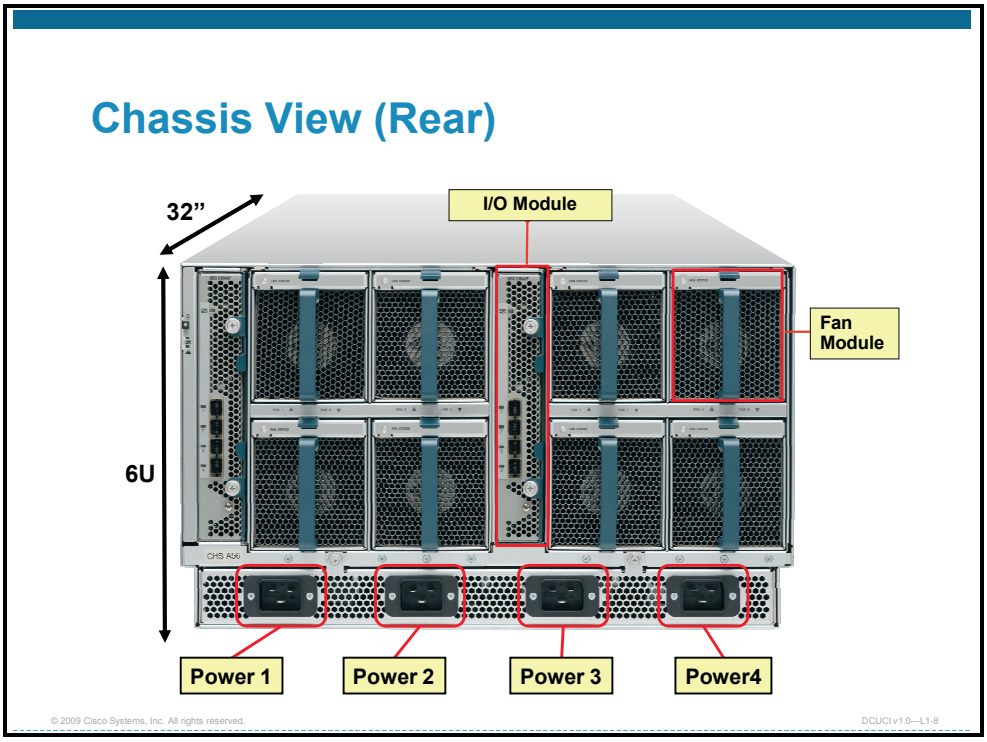

Chassis

The Cisco UCS 2100 Series IOMs are installed in the rear of the chassis. These provide connection between the blade server Ethernet mezzanine cards and the Fabric Interconnects. Eight fan modules provide cooling to the blade servers and chassis components.

Power

There are four 220V power entry ports to provide power to the front-side power supply units.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

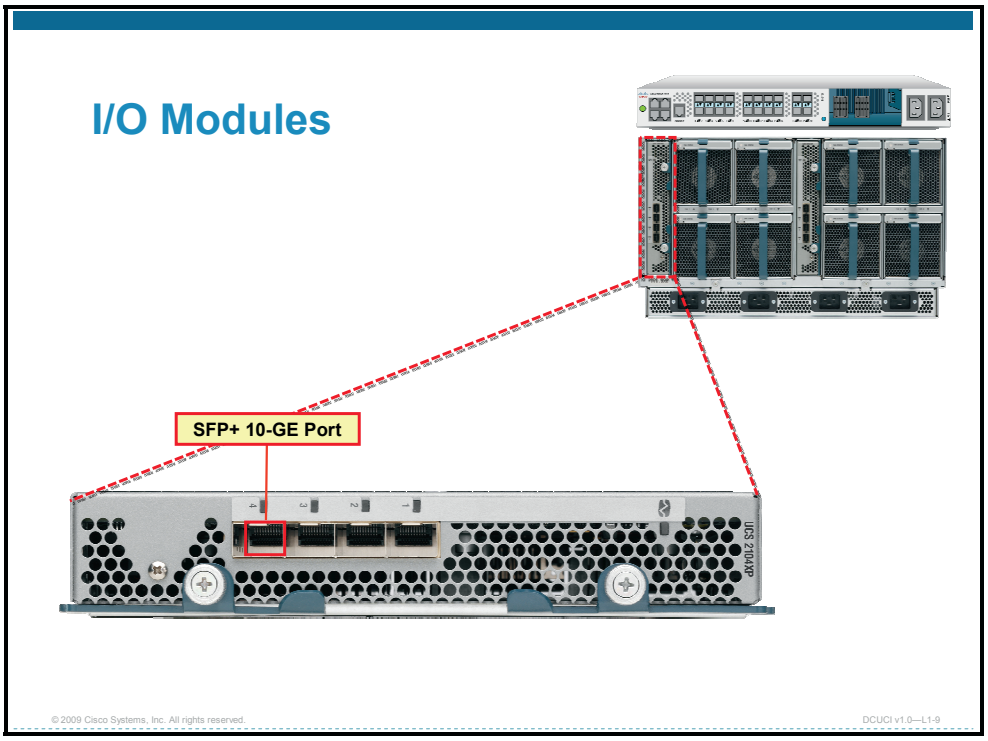

I/O Module

The IOM logically extends the fabric from the Fabric Interconnect to the blade server. Its primary function is to connect the blade servers to the Fabric Interconnect, which handles all switching for blade-to-blade and blade-to-fabric communications. No local switching occurs in the IOM. There are four 10-Gbps ports that can be used to connect the chassis to the switch. The chassis contains two IOM slots for increased redundancy and bandwidth, though only one IOM is required for a minimal configuration. If only one IOM is installed, it must be placed into the left bay as viewed from the rear of the chassis. If both are used, they are hot-swappable and fully redundant from an Ethernet perspective. Fibre Channel redundancy will be discussed later in this course. The IOM provides 10-Gbps per blade server port (there are two ports per blade server). Each blade server connects to both IOMs in the chassis via the mezzanine card.

The IOM supports increased scalability by multiplexing up to an eight-to-one oversubscribed blade servers to ports ratio. This would be the case for a fully populated chassis with half-width blades, using a single 10-GB uplink from the IOM to the Fabric Interconnect.

CMC

The IOM has a chassis management controller (CMC) which is a key part of management infrastructure. The CMC collects status data from the IOM using the Intelligent Platform Management Interface (IPMI) protocol over the interintegrated circuit (I2C) serial bus. The CMC then communicates that information to the management node using the Ethernet server link. The CMC controls power supply and fan speeds, and serves as a proxy for the Cisco UCS Manager to the blade servers for certain functions. We will also see that the CMC plays a part in the high-availability protocols.

https://t.me/learningnets

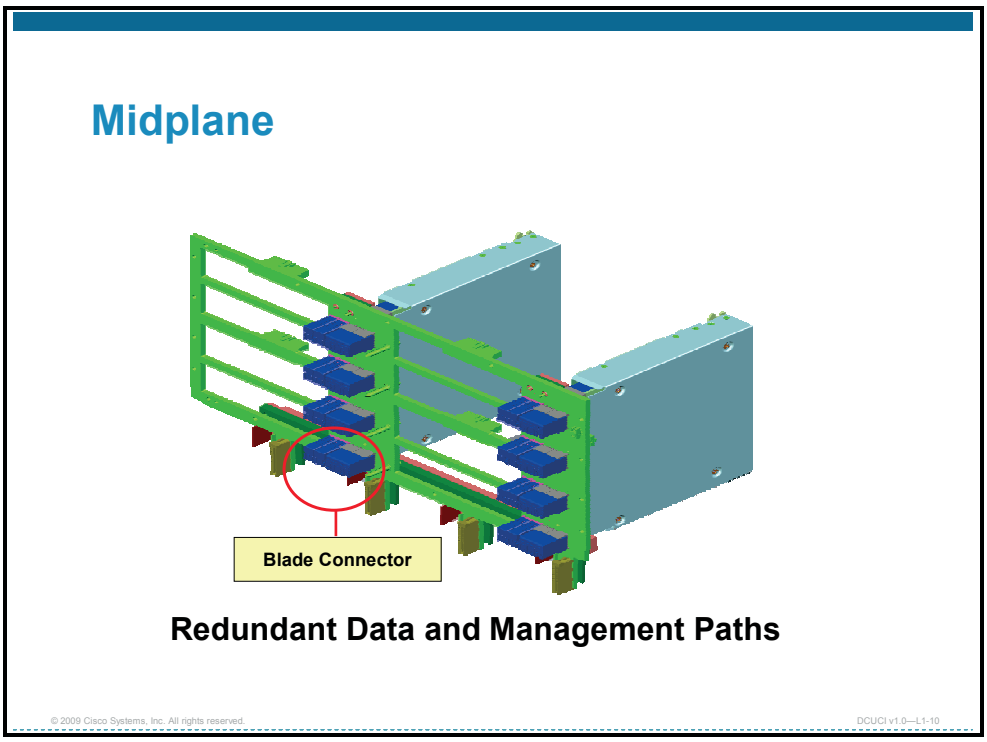

The Midplane

The midplane provides redundant power to the IOM (and in turn, the CMC) and blade servers that plug into it. The midplane also provides redundant data network (Ethernet) connectivity between the IOM and the blade servers. Lastly, the midplane has redundant management I2C paths and dedicated management network (Ethernet) connectivity, and supports auto-discovery of all components plugging into it.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Fabric Interconnect

This topic describes the Cisco UCS Fabric Interconnect models and expansion modules.

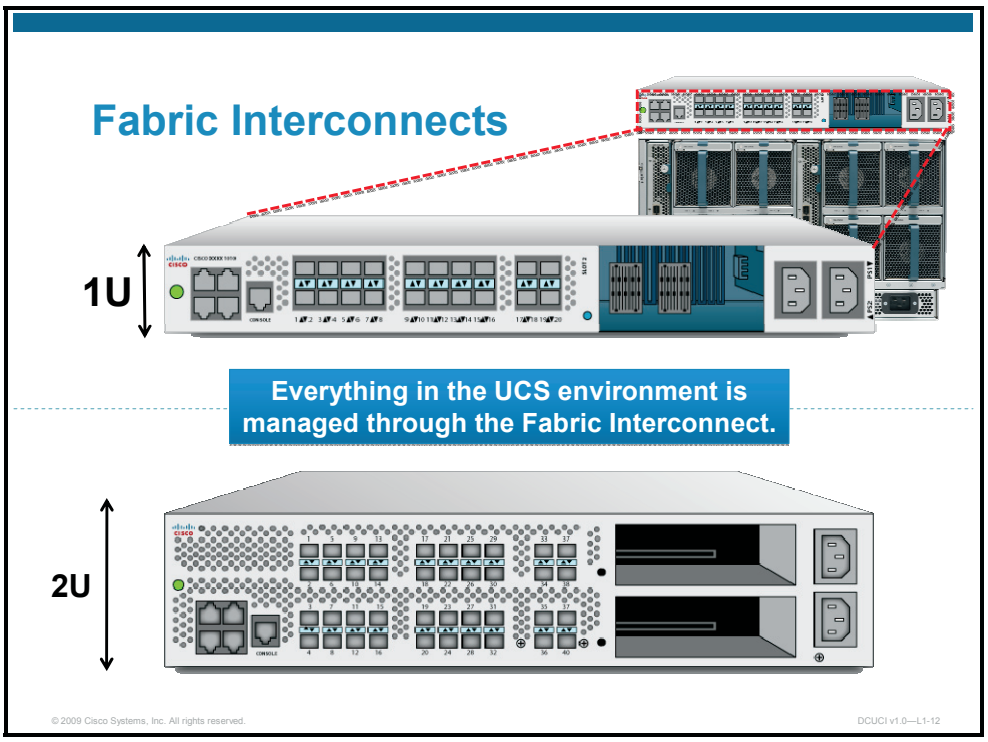

Fabric Interconnect

The Cisco UCS Fabric Interconnect is offered in 20-port or 40-port configurations.

The 20-port Fabric Interconnect is 1U high. It has twenty small form-factor pluggable (SFP+) ports for ten GE/Fibre Channel over Ethernet (FCoE). There is one expansion module for connecting to the LAN and SAN networks. The 20-port model supports up to a 560-Gbps (full-duplex) fabric. None of these ports are enabled or configured by default.

The 40-port Fabric Interconnect is 2U high. It has 40 SFP+ ports for 10-GE/FCoE. There are two expansion module slots for connecting to the LAN and SAN networks. The 40-port model supports up to a 1.12-Tbps (full-duplex) fabric. None of these ports are enabled or configured by default.

https://t.me/learningnets

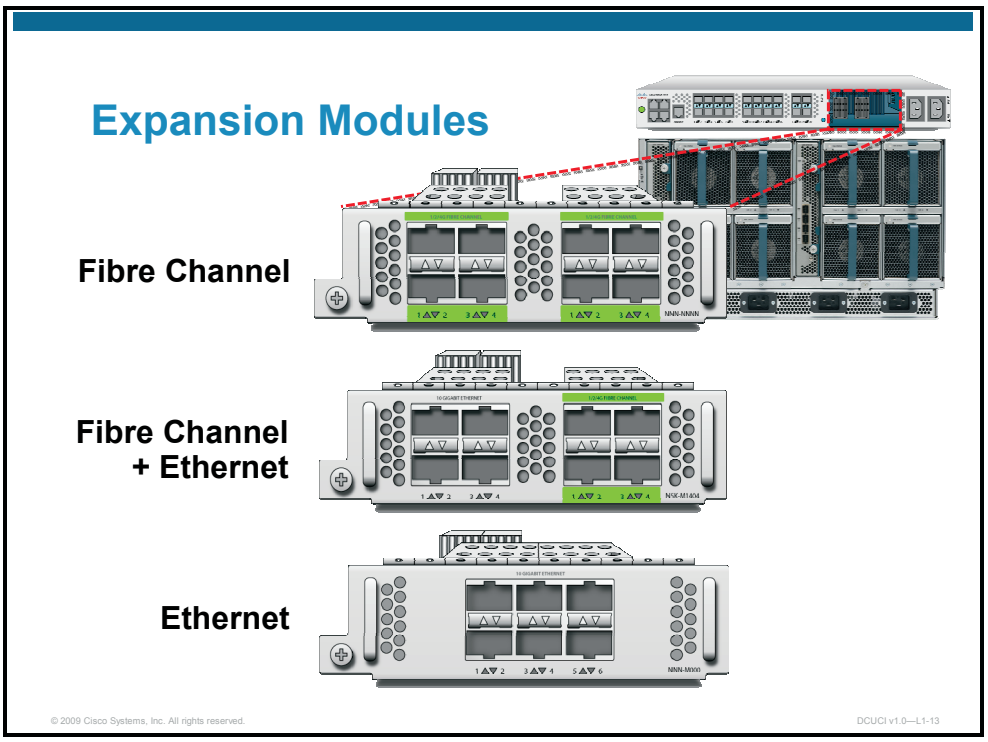

Expansion Modules

The expansion modules provide connectivity into your enterprise Ethernet LAN and Fibre Channel SAN networks. There are three expansion modules from which to choose: Fibre Channel-only, combination Fibre Channel and Ethernet, and Ethernet-only:

- The Fibre Channel-Only expansion module contains eight SFP ports that support 1-, 2-, and 4-Gbps Fibre Channel.
- The combination expansion module contains four SFP+ ports that support 10-GE and four SFP ports that support 1-, 2-, and 4-Gbps Fibre Channel.
- The Ethernet-Only expansion module contains six SFP+ ports that support 10-GE.

Note that Expansion Module ports cannot be connected to a chassis—these ports are for uplinks only.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Blades

This topic describes the Cisco UCS Blade Server hardware.

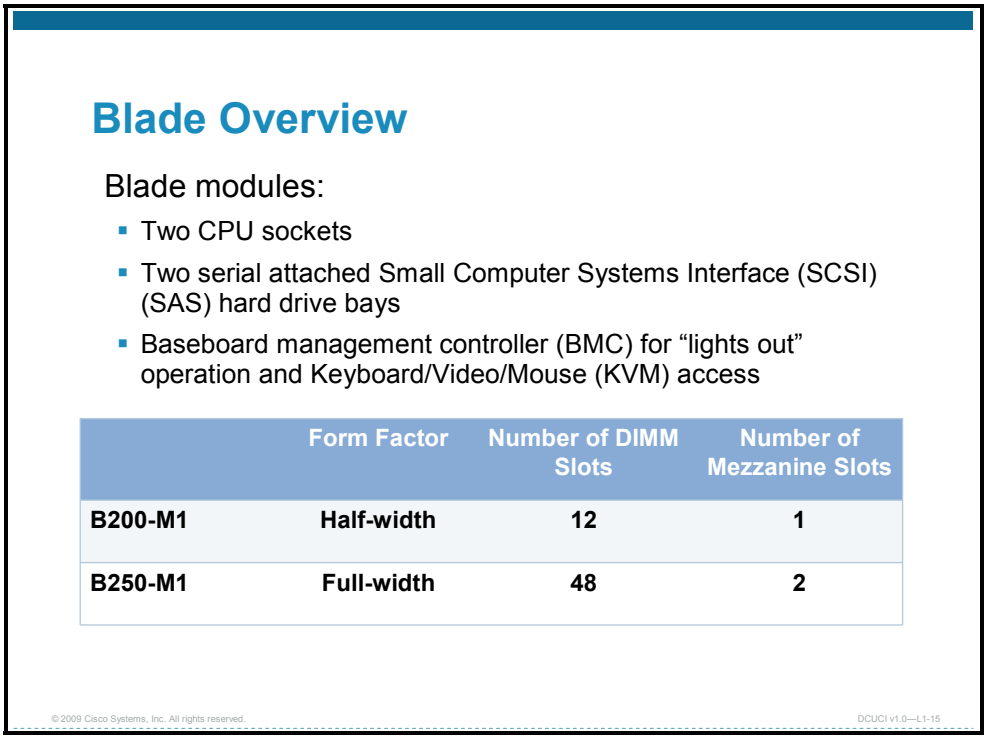

There are currently two blade modules for the Cisco UCS platform:

- B200-M1:
  - Half-width
  - Two processor sockets
  - Twelve DIMM slots
  - One mezzanine slot
- B250-M1:
  - Full-width
  - Two processor slots
  - Forty-eight DIMM slots
  - Two mezzanine slots

The two blade models share a common x64 Intel Xeon 5500 architecture and other features:

- Hot-swappable
- Two CPU sockets:
  - Customer option
- Two internal SAS hard disk drive (HDD) bays
- Customer option
- Baseboard management controller (BMC) for "lights out" operation

https://t.me/learningnets

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

The BMC is a microcontroller on the motherboard that provides “lights out” hardware status and configurability to the system management software.

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 KVM access to the end user through the switch over the dedicated Ethernet links.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

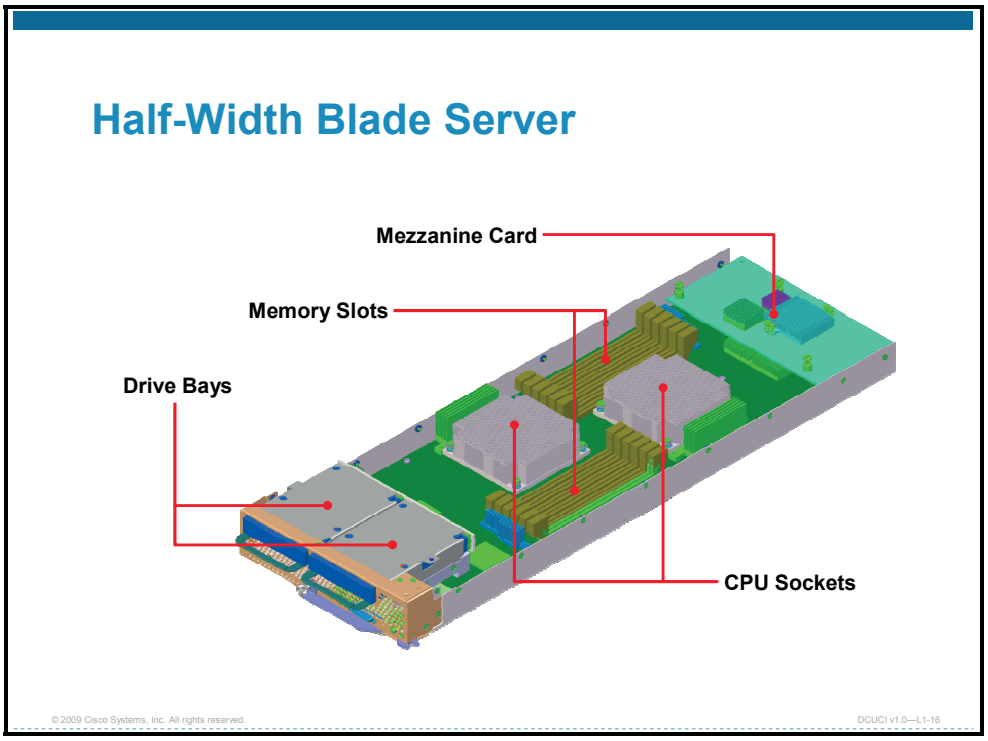

The Cisco UCS B200-M1 2-Socket Blade Server (half-width) contains two small form-factor (SFF) drive bays that support SAS disks in a hardware RAID-0 or RAID-1 configuration. The blade server has a 550W power and cooling budget. The chassis can hold up to eight half-width blade servers. The motherboard contains two CPU sockets and twelve DDR3 DIMM memory sockets. There is a single mezzanine adapter slot.

https://t.me/learningnets

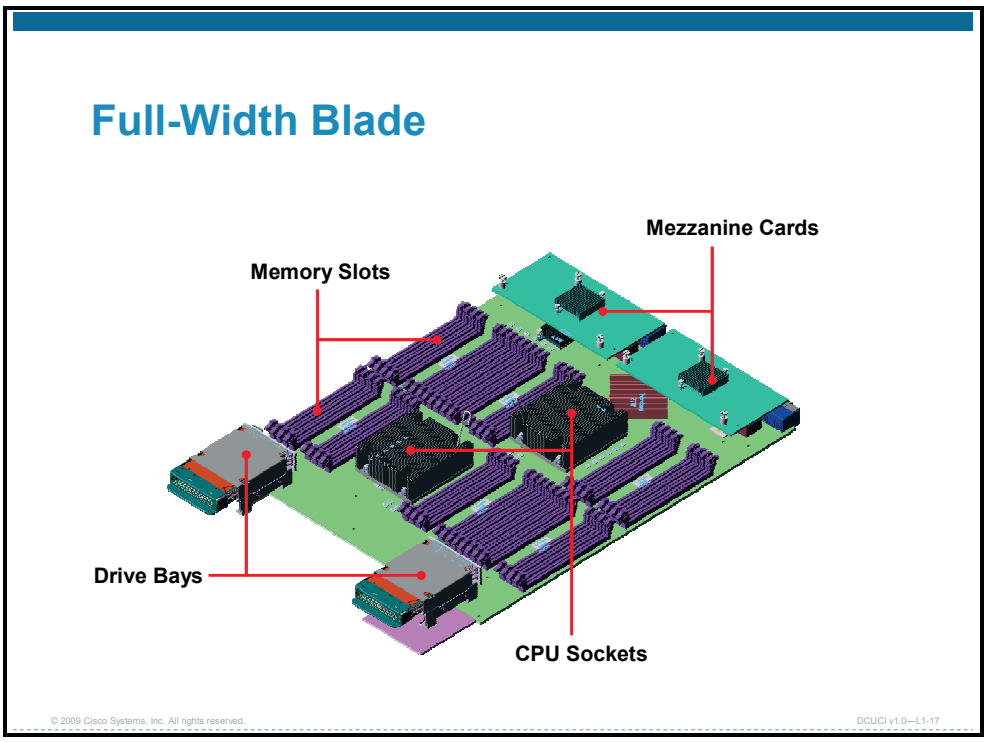

The Cisco UCS B250-M1 2-Socket Extended Memory Blade Server (full-width) has two SFF drive bays. The blade has an 1100W power and cooling budget. The chassis can hold up to four full-width blades. The motherboard contains two CPU sockets and forty eight DDR3 DIMM memory sockets, providing four times the memory capacity of the half-width blade.

There are two mezzanine adapter slots, doubling the I/O bandwidth of the Cisco UCS B200-M1.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Nehalem Memory Scalability

- Intel Nehalem Xeon processors support three memory channels of 3 DIMMS each
- Traditional servers limited to nine DIMMS per populated CPU socket

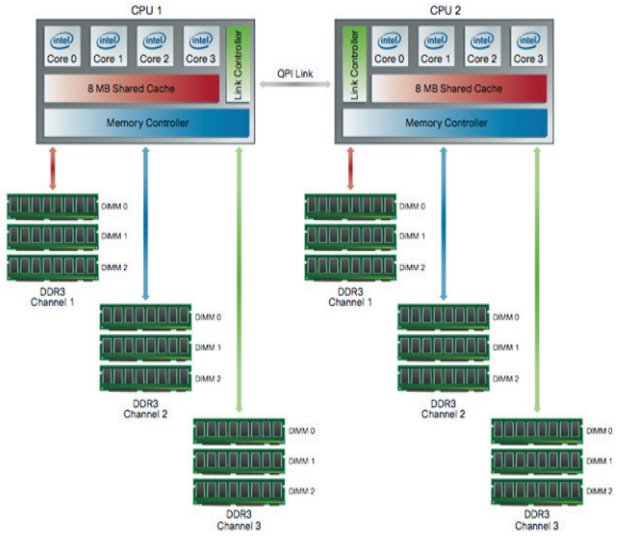

© 2009 Cisco Systems, Inc. All rights reserved.

The Intel Nehalem line of processors support three memory channels of three DIMMs each. This limitation is largely due to the electrical characteristics of the memory busses. A side effect of these characteristics is that as the number of DIMMs installed on a bus increase, the speed at which that bus runs may decrease. As such, many models of servers using the Nehalem processor use only two DIMMs per channel, for a total of six DIMMs per processor socket. In a dual-socket system, twelve DIMMs of 8-GB would limit the system to 96 GB, a commonly seen limit in today's server platforms, and in the B200-M1 Cisco UCS blade as well. To increase the amount of supported memory, additional sockets and processors must be added.

https://t.me/learningnets

Cisco Extended Memory Technology

- B250-M1 blades use Cisco Extended Memory Technology
- Custom ASICs virtualize 8 DIMMS per memory channel
- 192 GB supported per processor, 384 GB per blade
- Can also use lower cost DIMMs to reduce overall memory cost

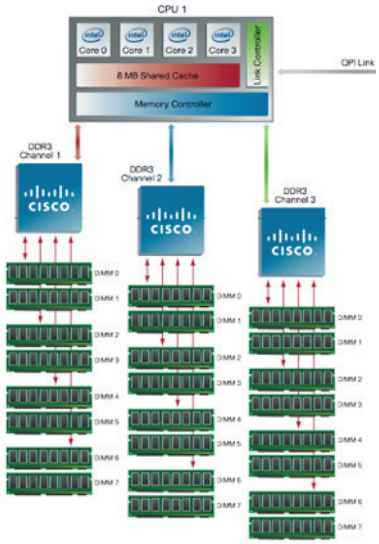

© 2009 Cisco Systems, Inc. All rights reserved.

Cisco has developed a custom memory controller chipset which virtualizes DIMMs, as seen by the CPU memory controller. Each memory channel can now support eight DIMMs instead of three, without a reduction in memory bus speed. The Cisco ASIC presents a virtual DIMM to the CPU memory controller, which is accessed normally. By allowing eight DIMMS per memory channel, or twenty-four per CPU socket, a great deal more memory may be installed on a server using this method. Alternately, smaller and lower cost DIMMs can be used to reduce the overall memory cost of a compute solution. Memory prices do not scale linearly with capacity, thus using more of the smaller capacity DIMMs results in a lower cost per gigabyte of memory.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Maximizing DIMMs in Intel Architecture

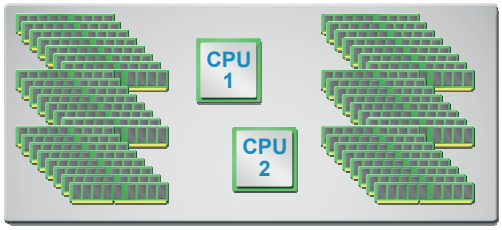

Capacity	Other Systems		Cisco Extended Memory	
	Configuration	Speed	Configuration	Speed
96 GB	12x 8GB	1066Mhz	48x 2GB	1066Mhz
144 GB	18x 8GB	800Mhz	32x 4GB + 16x 2GB	1066Mhz
192 GB	n/a		48x 4GB	1066Mhz
384 GB			48x 8GB	1066Mhz

© 2009 Cisco Systems, Inc. All rights reserved.

By allowing eight DIMMS per memory channel, or twenty-four per CPU socket, a lot more memory may be installed in a server—up to 384GB on the full-width blade. This is the largest amount of DIMM memory available in any existing two-socket system.

Alternately, smaller and lower-cost DIMMs can be used to reduce the overall memory cost of a compute solution.

Note that on other manufacturers' boards, scaling beyond 6 DIMM slots per CPU requires reduction of bus speed to 800Mhz.

https://t.me/learningnets

Maximizing DIMMs in Intel Architecture

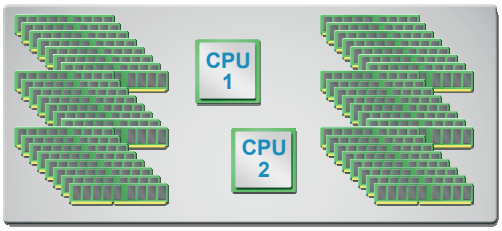

Capacity	Other Systems		Cisco Extended Memory	
	Configuration	Cost	Configuration	Cost
96 GB	12x 8GB	$11,640	48x 2GB	$2,880
144 GB	18x 8GB	$17,460	32x 4GB + 16x 2GB	$7,232
192 GB	n/a		48x 4GB	$9,407
384 GB			48x 8GB	$46,560

© 2009 Cisco Systems, Inc. All rights reserved.

Because memory prices do not scale linearly with capacity, using more smaller capacity DIMMs results in a lower cost per gigabyte of memory.

The costs shown in this table are based on the list price of equivalent form-factor Kingston server-class memory (240-pin DDR3 1066MHz registered ECC) as found on the CDW website (www.cdw.com) on September 28, 2009:

- 2GB DIMM: price = 59.99, price per GB = 29.99
- 4GB DIMM: price = 195.99, price per GB = 48.99
- 8GB DIMM: price = 969.99, price per GB = 121.25

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

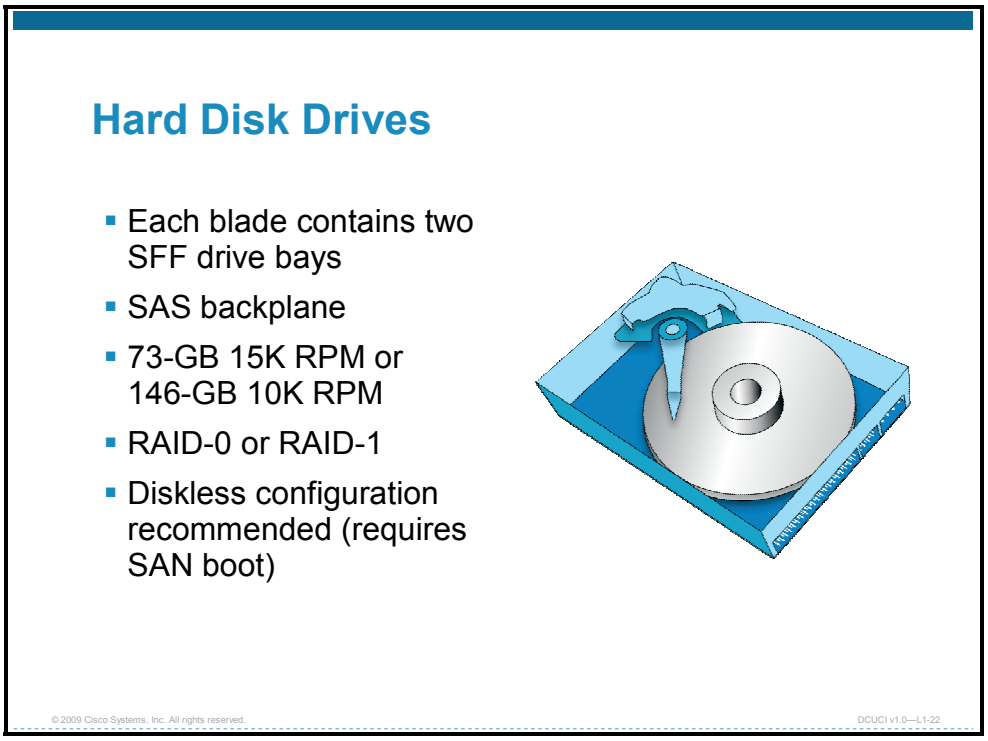

Hard Disk Drives

Both half-width and full-width blade servers have two drive bays and an SAS backplane. There are two supported drives; a 73-GB 15K RPM SAS SFF and a 146-GB 10K RPM SAS SFF drive.

While no local hard drive is required, if only one is installed, 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 serial ATA (SATA) controller supports RAID-0 and RAID-1.

UCS Feature Support

Cisco UCS features supported at the disk drive include statelessness, which is accomplished by scrubbing any data on a disk drive during a Service Profile migration.

Also, local disk policies such as RAID-0 or RAID-1 can be configured, and can migrate with logical Service Profiles between blade servers.

https://t.me/learningnets

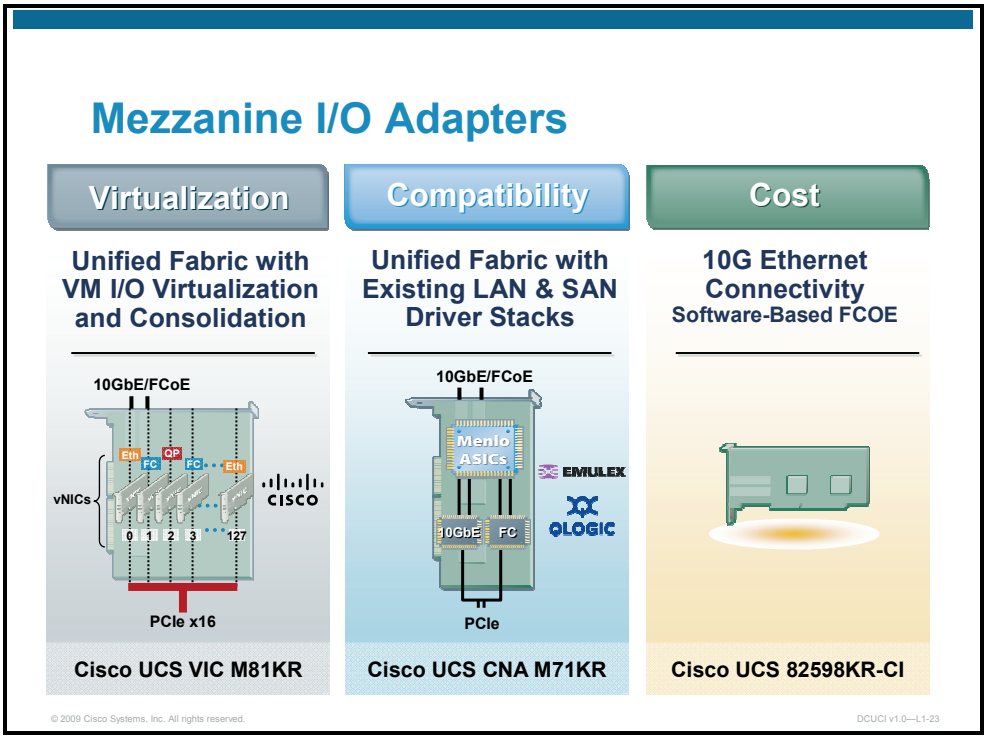

Mezzanine Cards

The types of adapters, or mezzanine cards, available in a Cisco UCS blade server include:

- Cisco UCS VIC M81KR: A virtual interface card
- Cisco UCS CNA M71KR-E: An Emulex converged network adapter
- Cisco UCS CNA M71KR-Q: A QLogic converged network adapter
- Cisco UCS 82598KR-CI: A 10-GE network adapter

Cisco UCS VIC M81KR

The Cisco UCS VIC M81KR supports multiple Ethernet virtual network interface cards (vNICs) or Fibre Channel Host Bus Adapters (vHBAs), running at 500 KIOPs in both initiator and target mode. This card supports kernel/hypervisor bypass. The Cisco UCS VIC M81KR is a converged network adapter (CNA) with dual 10-GE ports and dual Fibre Channel ports to the backplane. The card provides failover between the redundant Ethernet links. No multipathing software is required on the host operating system for this functionality.

Cisco UCS CNA M71KR

The Cisco UCS CNA M71KR is a converged network adapter (CNA) with two host-side 10-GE ports and two Fibre Channel ports to the backplane. The two network ports can run either native Ethernet or FCoE protocols and can be configured for failover. This failover is performed by the FCoE ASIC and does not require multipathing software on the host operating system:

The Cisco UCS CNA M71KR ASIC is a Cisco-designed multiplexor and FCoE protocol offload engine with a 350 MHz 24K MIPS processor. There are two versions of this card (Cisco UCS CNA M71KR-E Emulex Converged Network Adapter, and Cisco UCS CNA M71KR-Q QLogic Converged Network Adapter), thereby supporting existing proven driver stacks to the customer. Both cards use an Intel 82598 10GE Ethernet chipset.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 82598KR-CI

The Cisco UCS 82598KR-CI Network Adapter is made by Intel. It has two 10-GE ports to the backplane which run native Ethernet. operating system-implemented FCoE may be run over the Cisco UCS 82598KR-CI to provide “free SAN access” to the host. The card does not provide failover functionality.

Comparison

- The Cisco UCS VIC M81KR provides enhanced virtualization. This adapter supports mulitple vNICs presented to the ESX server which can then be used as physical NICs for virtual machines (VMs).
- The Cisco UCS CNA M71KR adapter provides increased scalability. The FCoE protocol encapsulation is offloaded from the host and performed in hardware on the Cisco UCS CNA M71KR.
- The Cisco UCS 82598KR-CI provides standard 10-GE connectivity to servers not needing Fibre Channel protocol support, or servers providing Fibre Channel protocol and FCoE encapsulation in software.

https://t.me/learningnets

Describe VMware Integration

This topic describes the VMware integration strategies in a Cisco UCS environment.

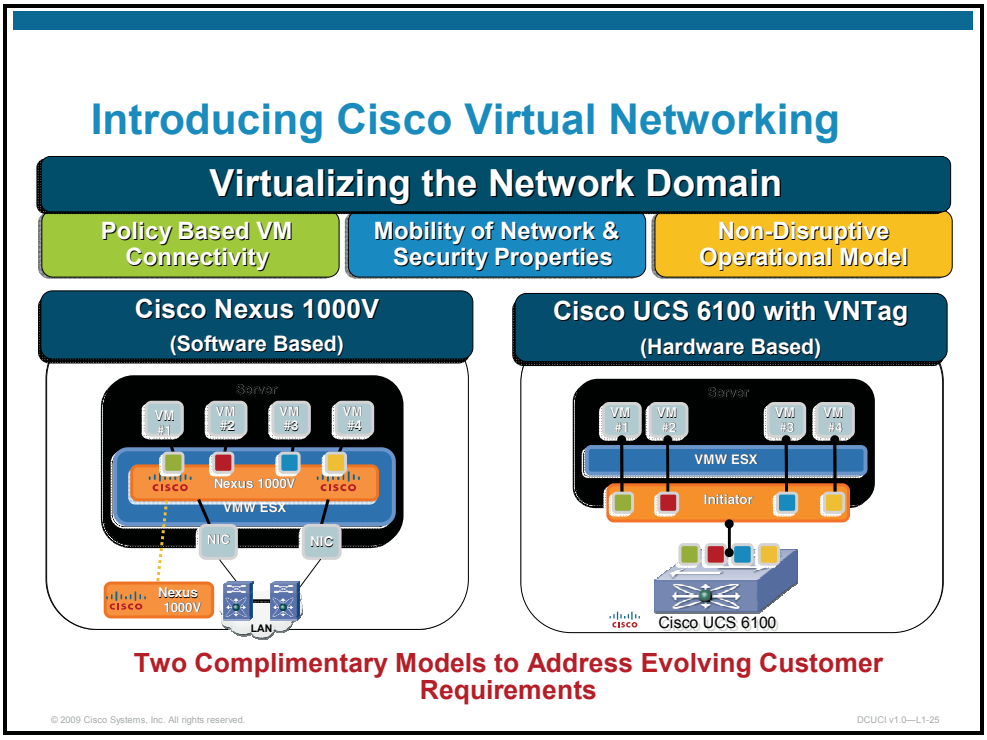

Using technology jointly developed by Cisco and VMware, the network access layer is moved into the virtual environment in order to provide enhanced network functionality at the VM level. This can be deployed as a hardware- or software-based solution, depending on the data center design and demands. Both deployment scenarios offer VM visibility, policy-based VM connectivity, policy mobility, and a non-disruptive operational model.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 1000V

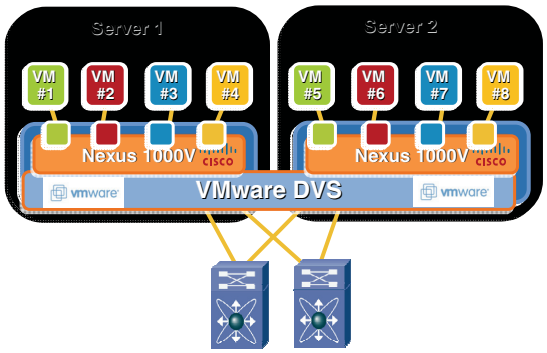

- Cisco Nexus 1000V provides enhanced VM switching for VMware ESX
- Features:
  - Policy-based VM connectivity
  - Mobility of network and Security Properties
  - Non-disruptive operational model
- Ensures proper visibility and connectivity during vMotion

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco Nexus 1000v bypasses the VMware vSwitch with a Cisco software switch, leveraging VMware's vNetwork Distributed Switch (DVS). This model provides a single point of configuration for the networking environment of multiple ESX hosts. Additional functionality includes policy-based connectivity for the VMs, network security mobility, and a non-disruptive software model.

https://t.me/learningnets

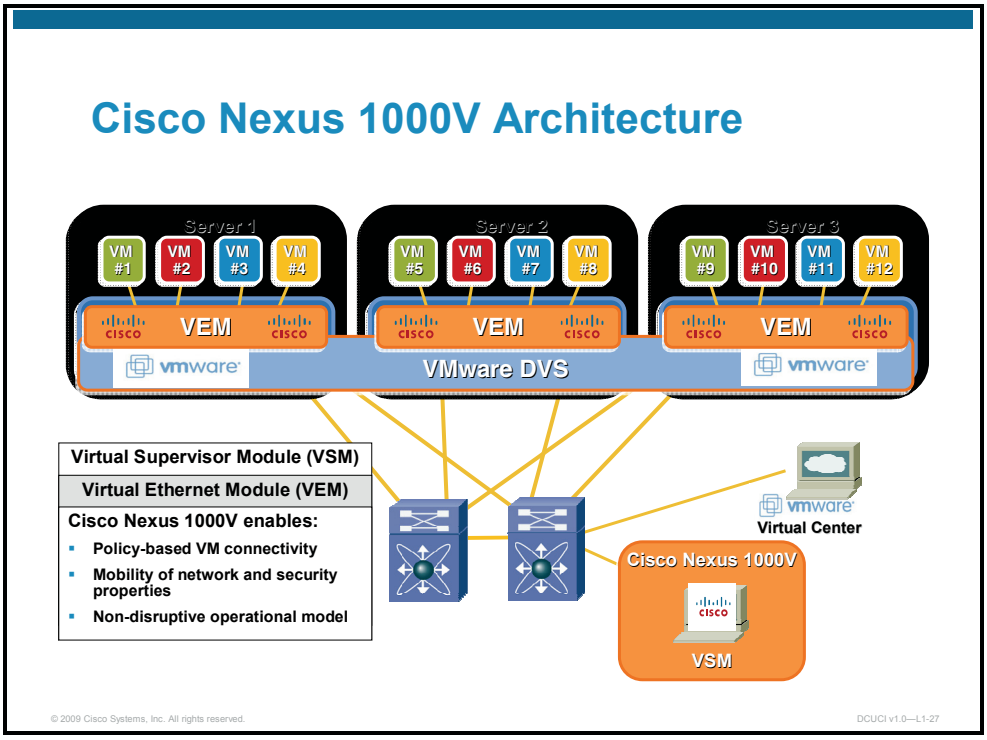

Virtual Supervisor Module (VSM):

- Virtual or physical appliance running Cisco operating system (supports high availability)
- Performs management, monitoring, and configuration
- Tight integration with VMware Virtual Center (vCenter)

Virtual Ethernet Module (VEM):

- Enables advanced networking capability on the hypervisor
- Provides each VM with a virtual dedicated “switch port”

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

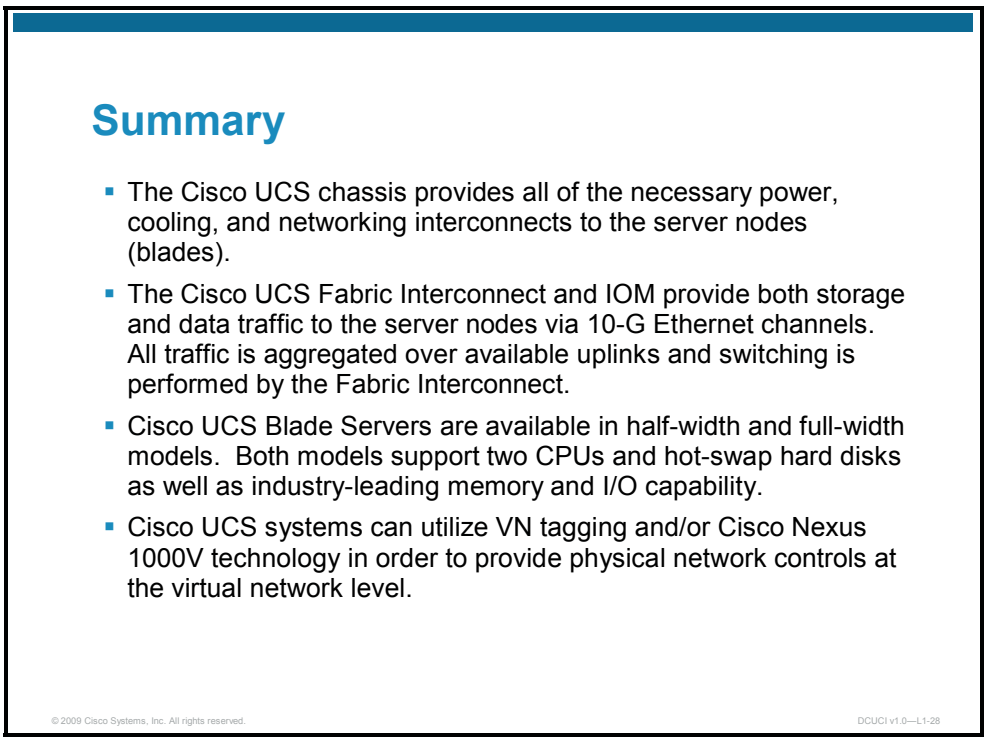

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Architectural and Feature Overview

Overview

In this lesson you will learn about the architecture and feature set of the Cisco Unified Computing System (Cisco UCS) platform.

Objectives

Upon completing this lesson, you will be able to discuss the architecture and feature set of the Cisco Unified Computing System. This includes being able to meet these objectives:

- Describe the Cisco UCS management architecture
- Describe the various responsibility assignments within the Cisco UCS environment
- Describe the diagnostic information available in Cisco UCS
- Describe external management solution integration options
- Describe the high availability and disaster recovery capabilities of the Cisco UCS Manager

https://t.me/learningnets

Cisco UCS Management Architecture

This section describes the management architecture of the Cisco UCS environment.

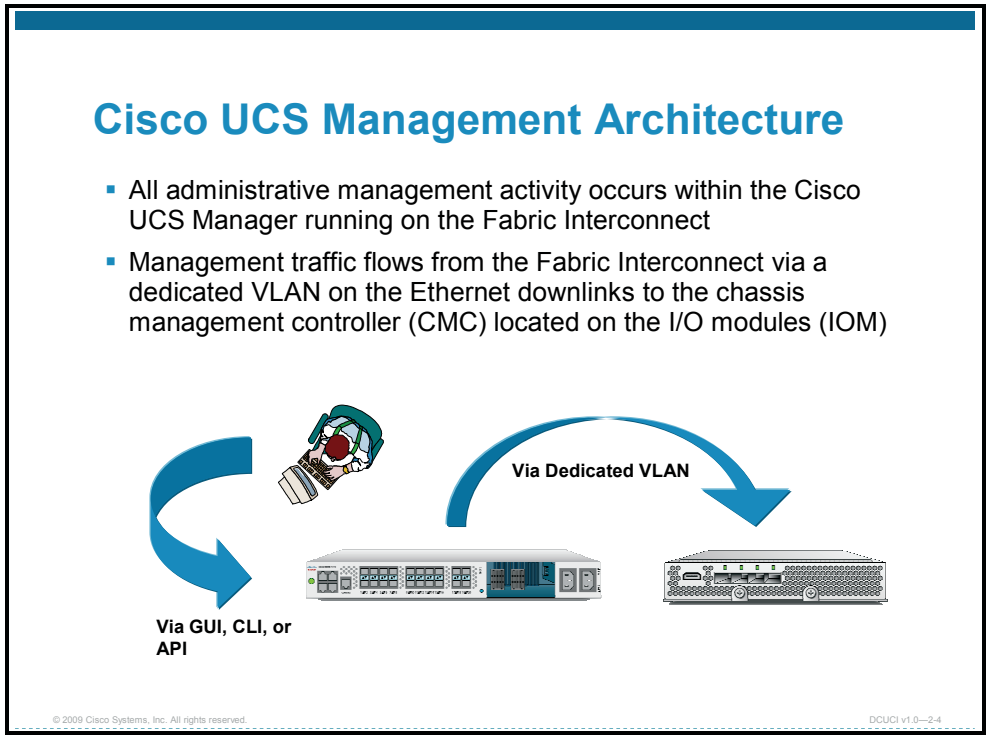

All Cisco UCS device management is performed via in-band VLAN channels. Using VLAN 4044, the IOM and chassis are managed by the Cisco UCS 6100.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Management Architecture (Cont.)

- Chassis-related management tasks are passed via redundant serial links to the appropriate device
- IOM-related management tasks are passed via interintegrated circuit (I2C) links to the appropriate IOM chipset
- Blade-related management tasks are passed via dedicated, redundant Ethernet links to the blade's baseboard management controller (BMC)

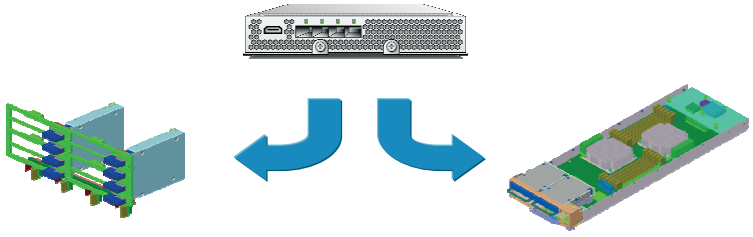

© 2009 Cisco Systems, Inc. All rights reserved.

Chassis, IOM, and processing node management is passed via various internal channels, allowing the Cisco UCS 6100 to be the single point of management for the entire Cisco UCS.

https://t.me/learningnets

Cisco UCS Fabric Interconnect Management

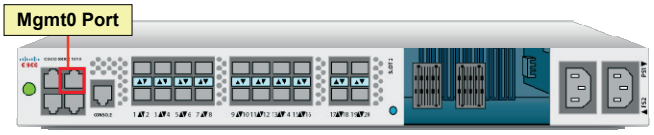

- Management of the Cisco UCS environment is performed through the Cisco UCS Manager, installed on the Fabric Interconnect
- The Cisco UCS Manager accepts connections via the GUI, command-line interface (CLI), and Extensible Markup Language (XML) application programming interface (API)
- Some physical blade management and reporting services may be optionally exposed via the Intelligent Platform Management Interface (IPMI) open standard
- Blade Keyboard/Video/Mouse (KVM) functionality is also provided through this port

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS 6100 series Fabric Interconnect devices are the single management point in Cisco UCS deployments. The Cisco UCS management software runs within the Cisco Nexus Operating System (Cisco NX-OS) on the Fabric Interconnect. The management system can be accessed via the Fabric Interconnect management port. The management port provides CLI, GUI and XML access, and presents remote KVM functionality for processing nodes (blades).

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

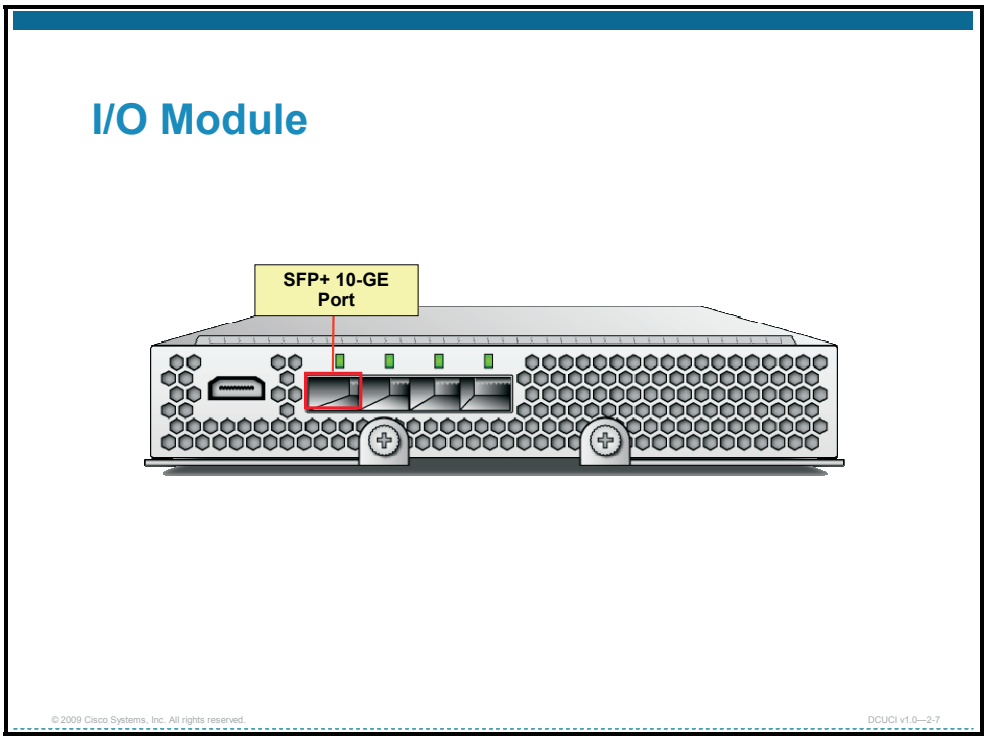

Chassis Management Controller

The IOM has a CMC that is a key part of management infrastructure. The CMC collects status data from the IOM using the IPMI protocol over the I2C serial bus. The CMC then communicates that information to the management node using the Ethernet server link. The CMC controls power supply and fan speeds, and serves as a proxy for the Cisco UCS Manager to the blade servers for certain functionality. The CMC also supports some of the Cisco UCS Manager high availability functions, which will be described in a later lesson.

https://t.me/learningnets

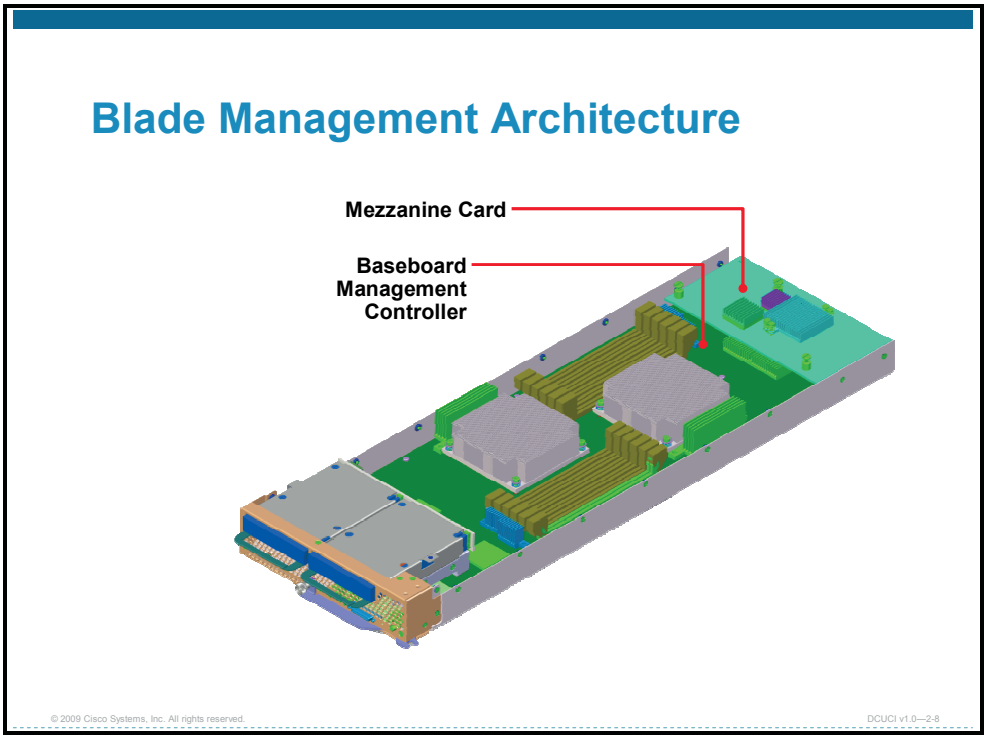

The figure depicts a one-half width Cisco UCS blade. The concept is similar for the full width blade.

Each blade server contains a BMC which communicates “lights-out” status information, KVM, and other status information to the CMC.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Management Assignments

This topic describes the various responsibility assignments within the Cisco UCS environment.

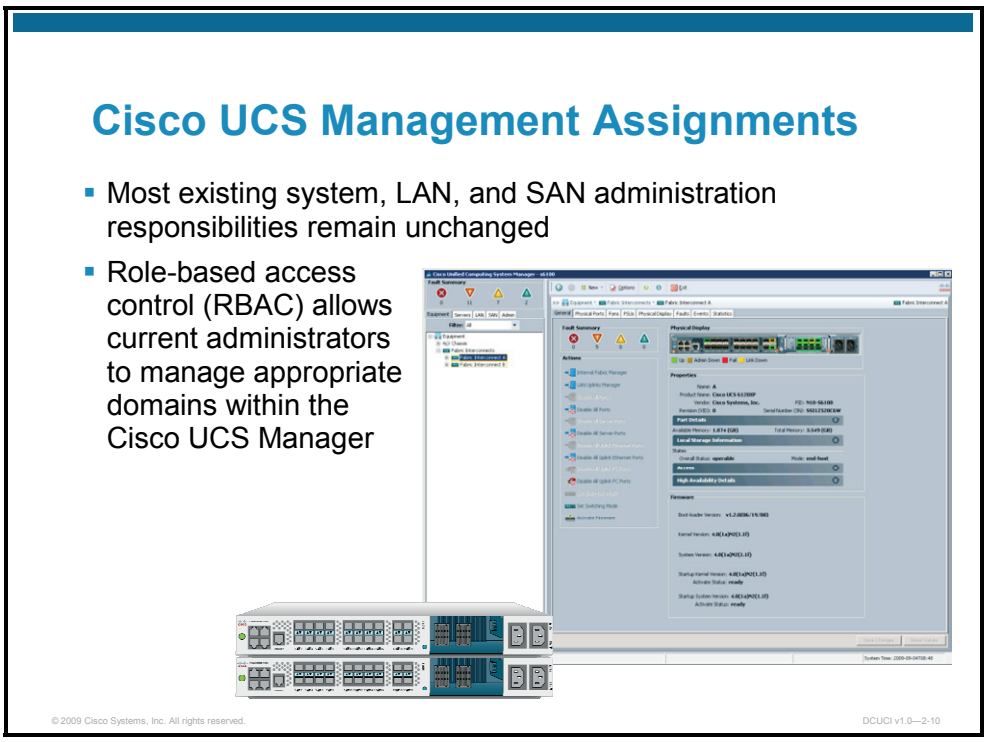

Cisco UCS is designed to be deployed with maximum flexibility and minimum disruption within existing environments. The unified switching infrastructure of the Cisco UCS 6100 Fabric Interconnect can operate in End-Host Virtualizer (EHV) mode or as a traditional switch for Ethernet. The Fabric Interconnect operates in N Port Virtualizer (NPV) mode for SAN uplinks. This allows full compatibility with no disruption when implementing into networks that contain switching infrastructures from other vendors.

https://t.me/learningnets

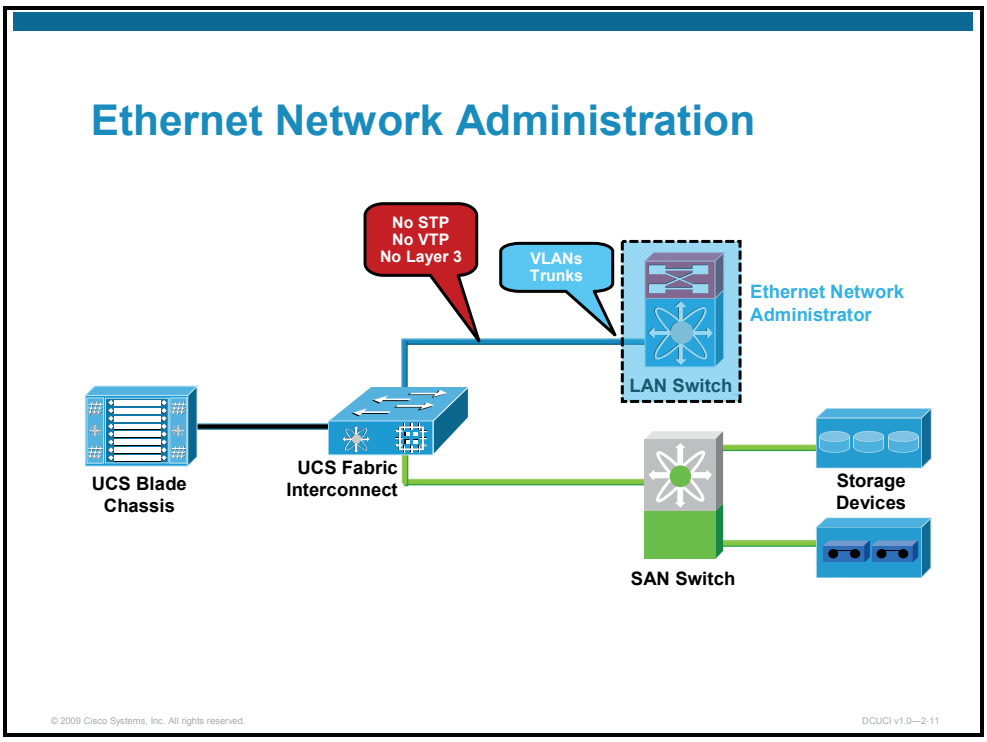

When running in switch mode, the Fabric Interconnect is managed in the same manner as any other Layer 2 access switch in the data center. The Ethernet Network Administrator is still responsible for LAN configuration, including configuration of 802.1Q VLAN trunks and EtherChannel aggregates on uplink switches.

When the Fabric Interconnect runs in EHV mode it does not participate in VLAN Trunking Protocol (VTP) or Spanning Tree Protocol (STP) schemes. The Cisco UCS 6100 is a Layer 2 device and does not participate in routing protocols. Uplink network configuration consists of link aggregation and trunking settings.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

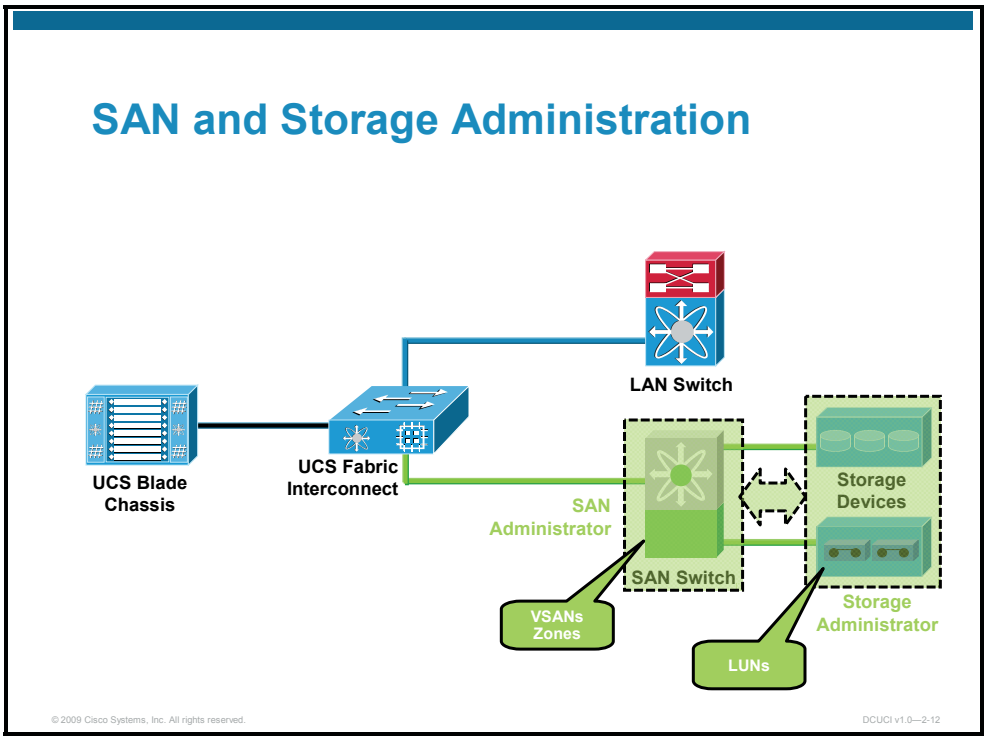

In a UCS environment, the SAN administrator is still responsible for placing uplink ports in appropriate VSANs (Cisco) or physical fabrics, zoning, quality of service (QoS), and other network services with traditional tools. The storage administrator remains responsible for LUN creation and masking in the same fashion as in any standard server or blade environment.

The Fabric Interconnect operates in NPV mode which prevents it from participating in Fabric Shortest Path First (FSPF) routing and domain name schemes. This mode is based on the N Port Identifier Virtualization (NPIV) industry standard and will allow the Cisco UCS to seamlessly integrate into multi-vendor SAN environments as long as the connected SAN switch supports NPIV. The SAN administrator must ensure that all uplink switch ports are placed in the appropriate Virtual SAN (VSAN) on Cisco switches or physical fabric when connecting to other vendor's devices.

https://t.me/learningnets

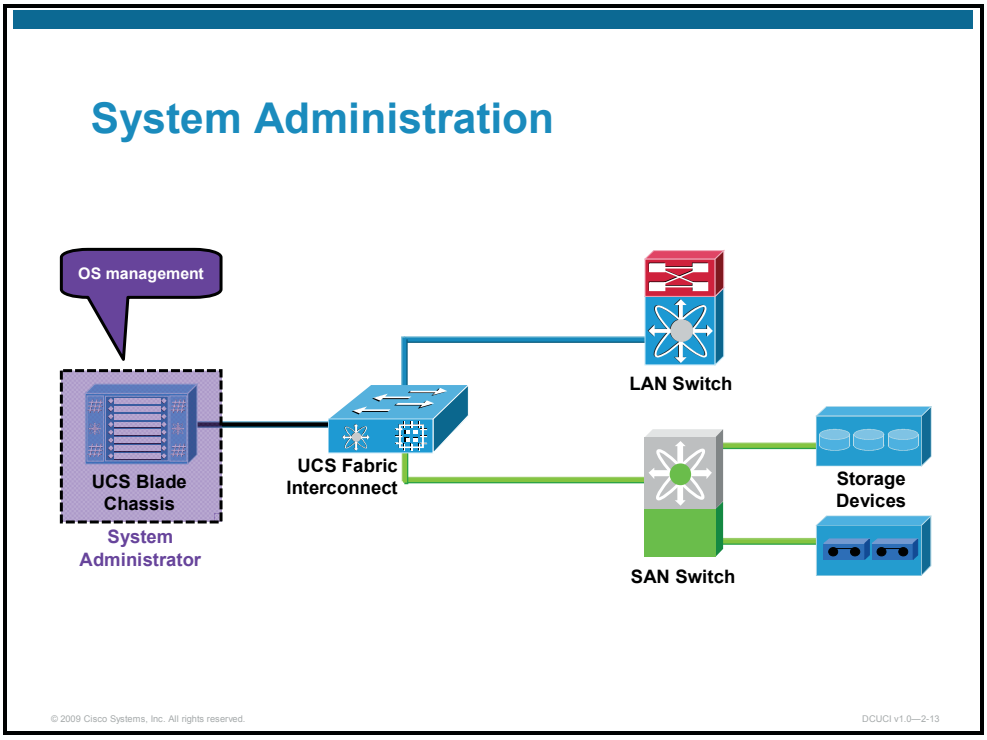

In Cisco UCS environments, system administrators remain responsible for all operating system tasks, following the same model as traditional server environments, including installing, tuning, and maintenance of operating systems. Cisco UCS performs no operating system-level configuration or management.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

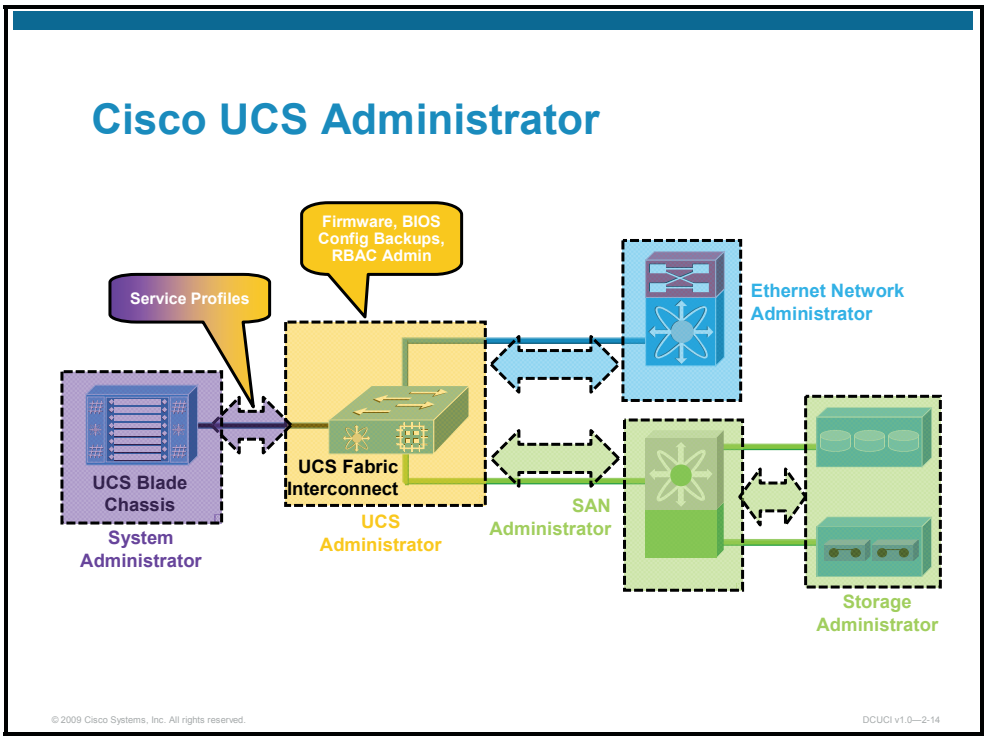

The Cisco UCS administrator is a new role within the datacenter. The Cisco UCS administrator is typically responsible for coordination with each data center team, physical resource assignment, server profile creation, creation and management of internal VSAN and VLAN objects, BIOS and firmware management, and Cisco UCS configuration backup. This can be a new role for a current team, a role for a new team, or these tasks can be divided among existing teams.

https://t.me/learningnets

Diagnostic Information

This section overviews the diagnostic information available in a Cisco UCS environment.

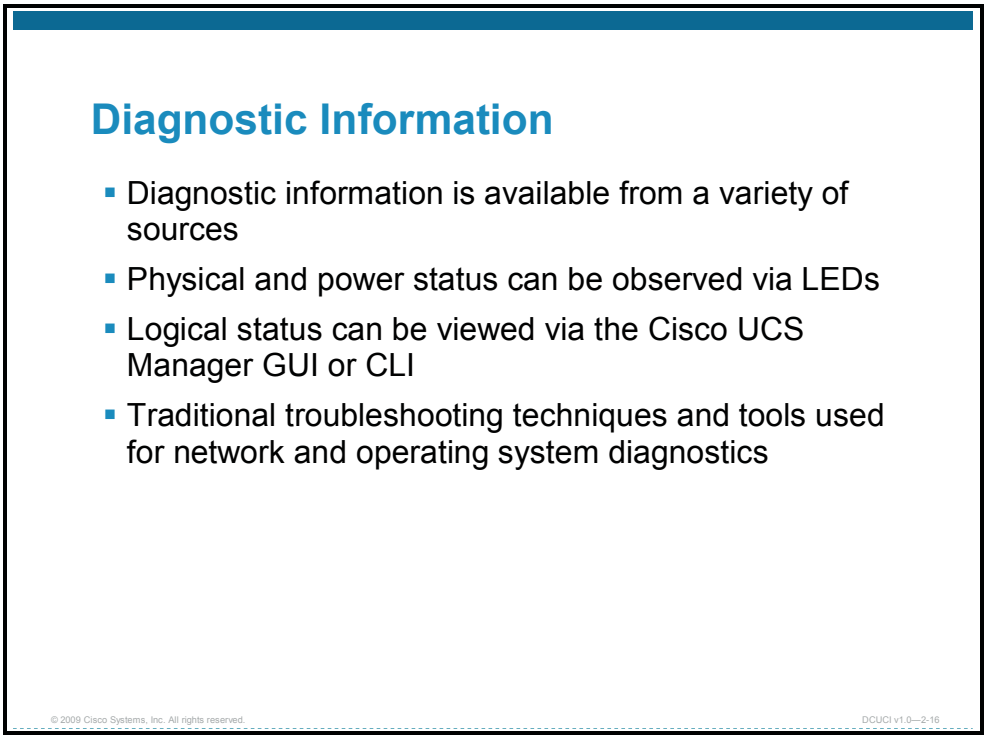

The Cisco UCS offers various diagnostic tools including LED, GUI, and CLI troubleshooting indicators and information. All traditional troubleshooting tools and techniques apply within Cisco UCS environments.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

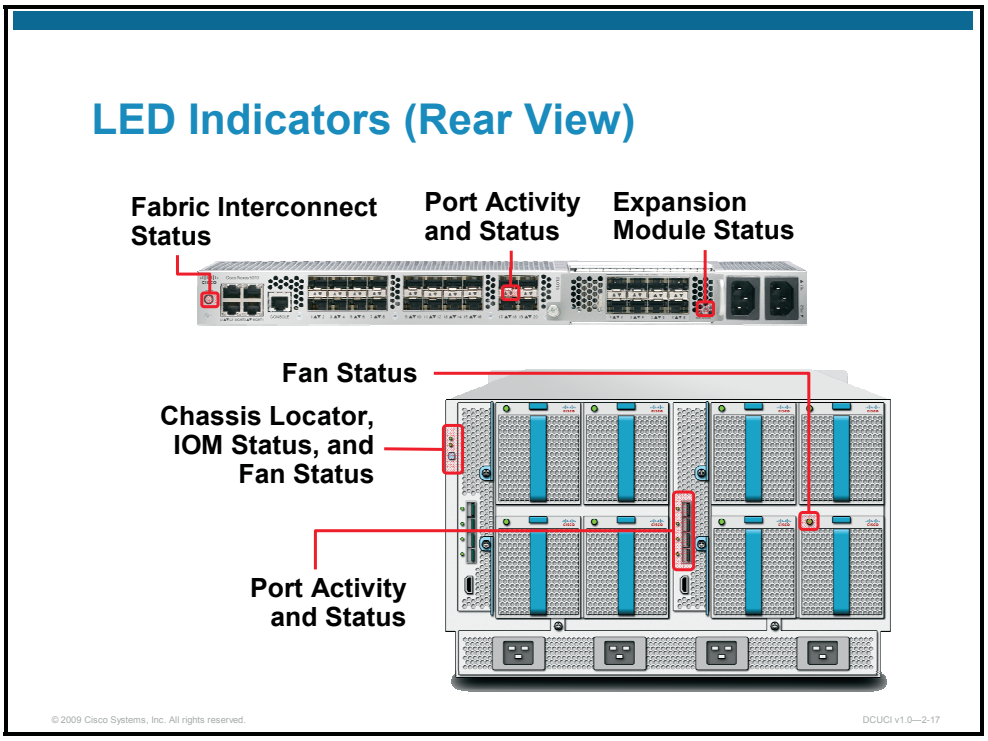

This rear view of a Cisco UCS system shows the various LED indicators that the operator can use to identify components and their states. Different LEDs have different colors and flashing patterns to indicate their status. Some of the components support beaconing (for example, I/O modules) to help the support staff to locate them.

The L1, L2, and Mgmt0 ports also have status and activity lights in the rear.

https://t.me/learningnets

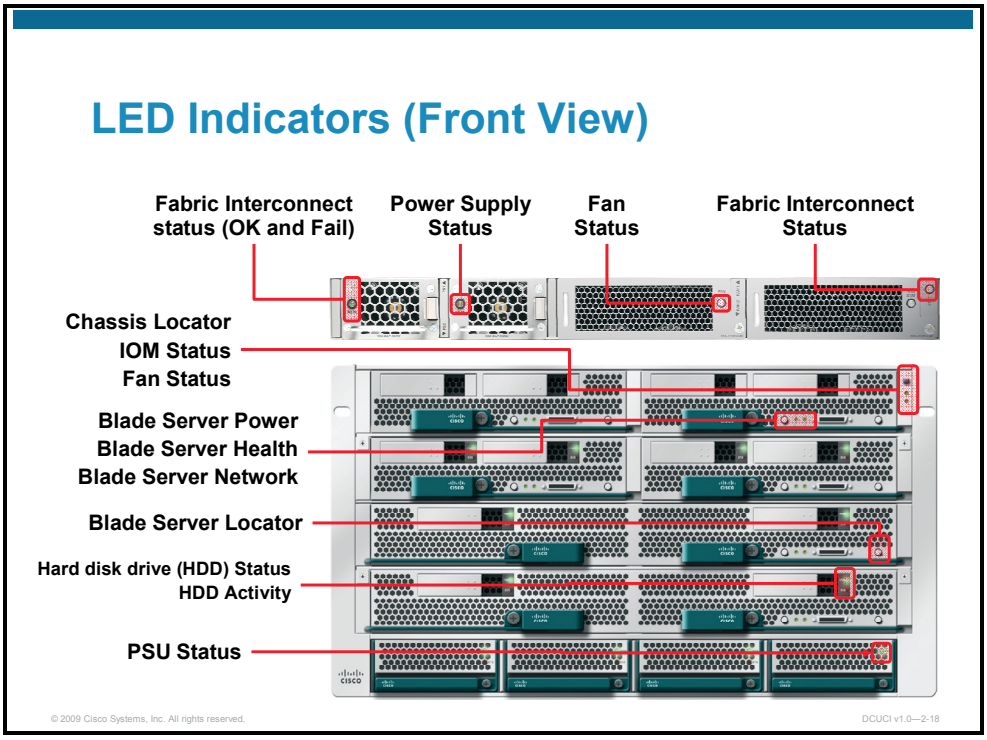

The front-view of the Cisco UCS system also has diagnostic lights that show status and locators for certain components. Some of these diagnostic LED indicators are duplicated on the front and back (for example, chassis locator, IOM status, fan status, and Fabric Interconnect status).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Management Solution Integration

This section describes the protocols and options for integration with external management solutions.

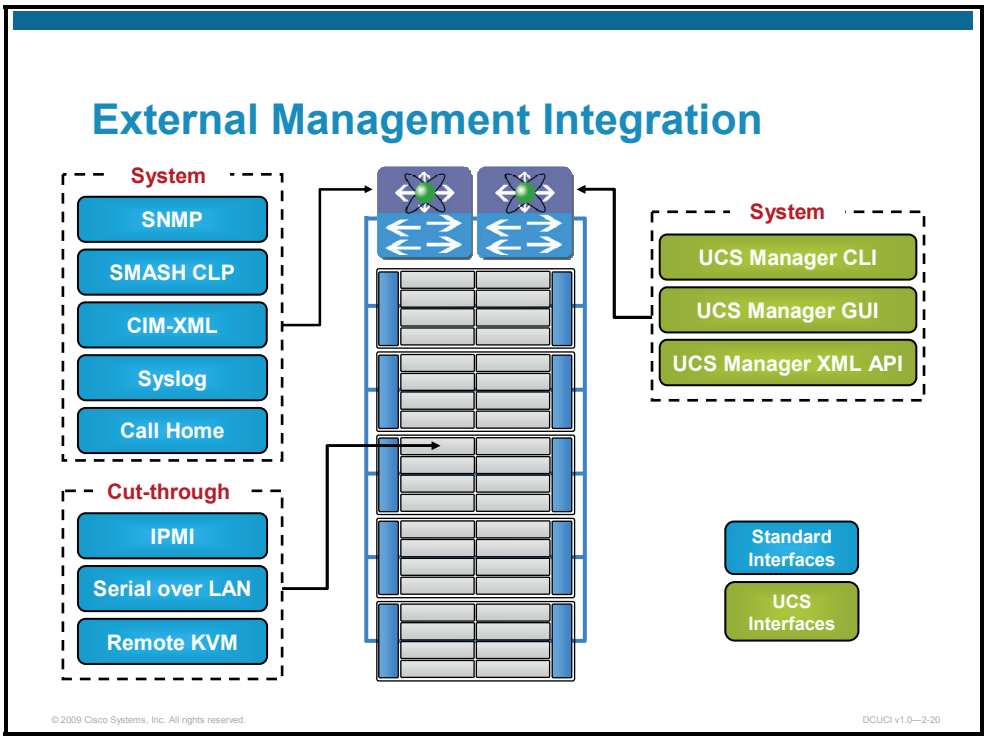

External management tools can be easily integrated into Cisco UCS environments using the provided tools. Cisco UCS offers industry standard tools including: XML-API, web services management (WBEM), SNMP, and IPMI for third-party management integration. Software development kits (SDKs) are available for full control of the Cisco UCS environment.

Cisco UCS management supports a variety of industry-standard protocols for monitoring, configuration, and Call Home. The Cisco UCS Manager XML API and corresponding SDKs allow you to create custom wrappers for the Cisco UCS Manager configuration and monitoring, to accomplish specific functions:

- Creating a multi-tenant portal with your own presentation/authorization/authentication mechanisms
- Integrating with custom or industry-standard orchestration tools
- Populating external configuration management databases (CMDBs) to include management changes in a Cisco UCS environment into your enterprise management system
- Scripting and integrating into custom management solutions

https://t.me/learningnets

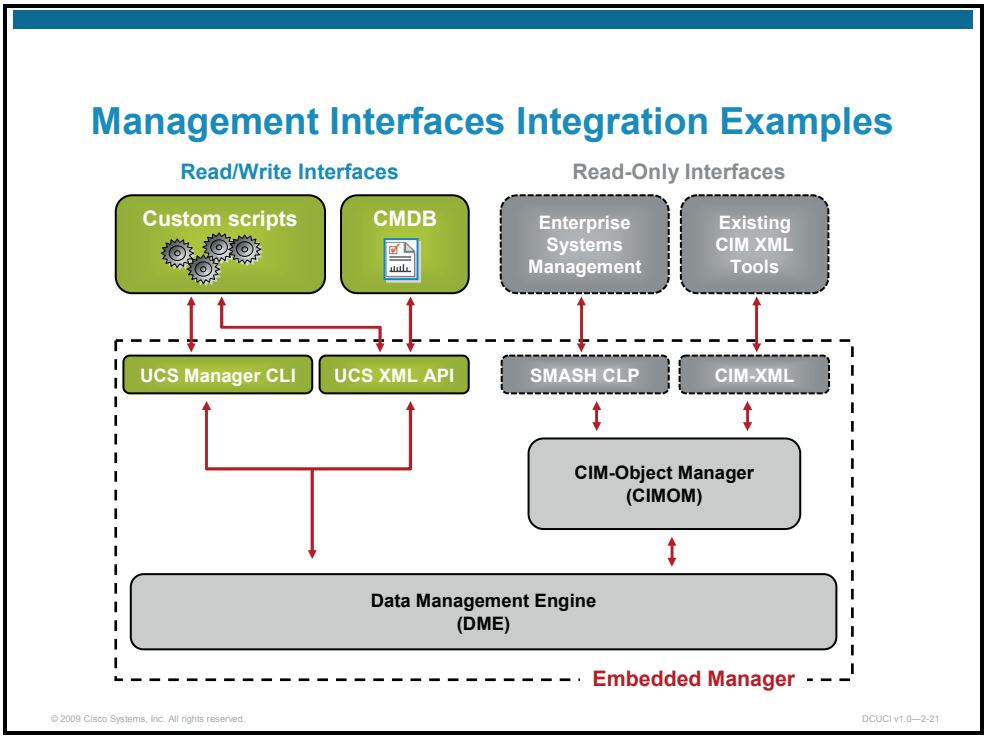

The Data Management Engine (DME) is a UCS Manager service that enables interaction between the UCS hardware and software components and external applications. The DME provides API services—all interfaces are ultimately translated into the XML API—and authentication servers.

Custom scripting can be done using the CLI or using the XML API directly. Configuration Management Databases (CMDBs) can be integrated using the XML API. Everything that is available in the UI in the product is available through the XML API.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

BMC BladeLogic

- End-to-end provisioning tool
- Addresses multiple management aspects
  - Integrated UCS management
  - Operating system provisioning
    - Installation
    - Boot
  - Application Installation
  - Patch installation
  - Configuration and compliance management
  - Business service deployment orchestration

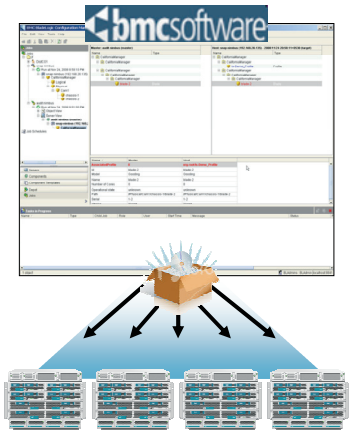

© 2009 Cisco Systems, Inc. All rights reserved.

BMC BladeLogic integrates and supports UCS management by leveraging unique UCS features like policies and service profiles. The integration allows administrators to manage and to deploy servers, and to perform compliance operations on UCS. The BMC BladeLogic enables administrators to perform day-to-day UCS system management without using native UCS Manager CLI or GUI management interfaces.

BMC BladeLogic is a purpose-built automation solution for the Unified Computing System. It exploits the UCS XML API to provide the automation intelligence. UCSManager does the work BladeLogic tells it to do when BladeLogic tells it to do it.

https://t.me/learningnets

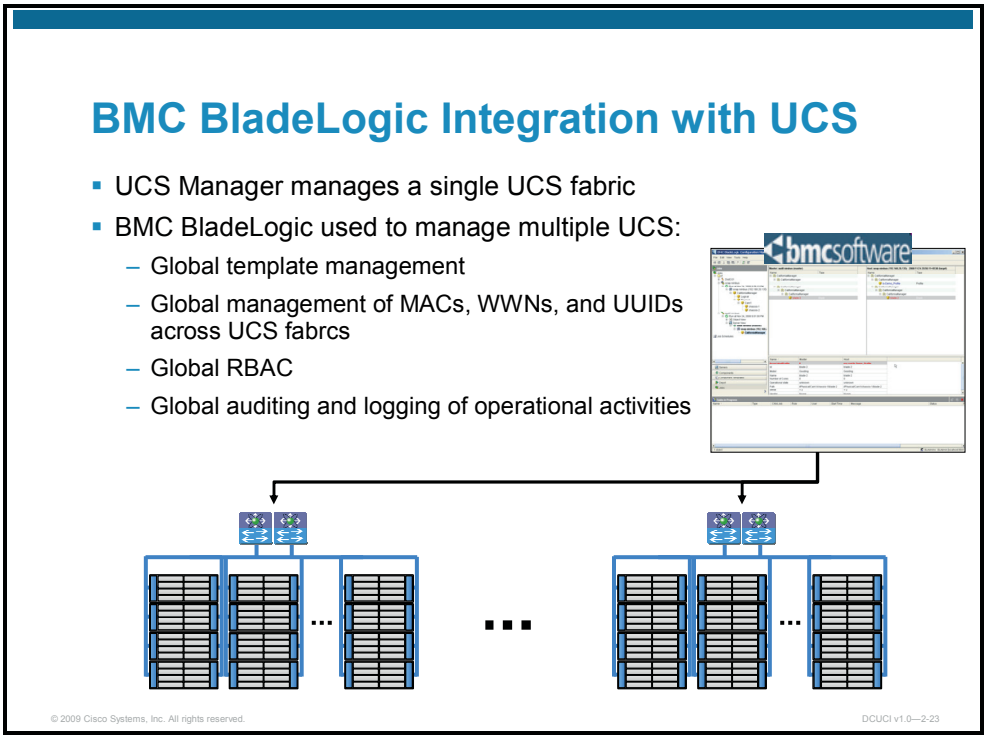

BMC BladeLogic enables administrators in a multi-UCS deployment scenario to manage multiple UCS fabrics from a single application. BMC BladeLogic manages across physical UCS systems, enabling system administrators to define a template once (for an Oracle Database server, for example) and then replicate that template across the farm of UCS. It also manages unique IDs, such as MAC addresses and WWN addresses, for the entire server farm, ensuring that an administrator never assigns an address that is already in use.

BMC BladeLogic reaches beyond the UCS into the server software including the hypervisor layer, guest operating system and the applications.

BMC BladeLogic uses custom objects to manage UCS environment. The discovery and import of UCS objects is done using XML API. The objects imported include general inventory, profiles, templates, relationships, topology, connectivity, and policy definitions. BMC BladeLogic stores and edits the hardware templates initially created by the 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.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS Manager High Availability and Disaster Recovery

This section describes high availability capabilities of the Cisco UCS Manager.

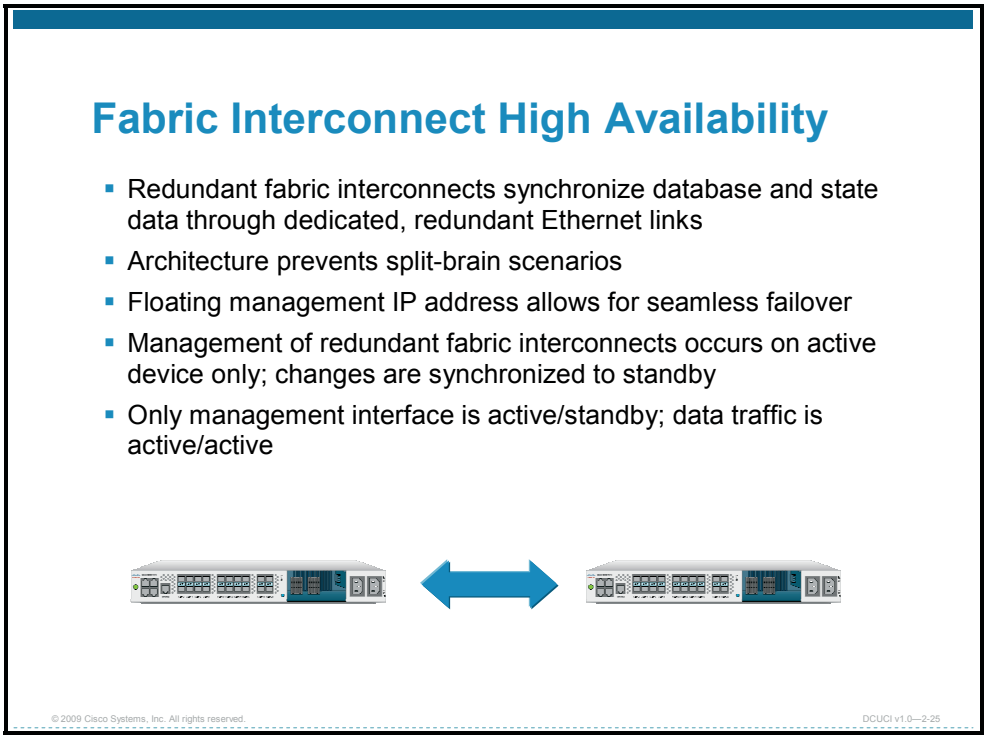

Cisco UCS 6100 Fabric Interconnects can be configured for high availability, providing physical redundancy. These devices can maintain a constant state sync using the high availability connection ports. The high-availability architecture is designed for fast failover and maximum protection against split-brain scenarios.

https://t.me/learningnets

Cisco UCS Manager Disaster Recovery

- Recovering a full Cisco UCS Manager implementation after a disaster can use multiple backup types
- Backups can be scheduled and automated to an external TFTP, Secure Copy Protocol (SCP), or FTP server
- Full database backups include physical hardware states such as service profile associations
- Configuration data can be backed up as XML:
  - config-all
  - config-system
  - config-logical
- XML-format backups do not include associations and can be edited before restoration

© 2009 Cisco Systems, Inc. All rights reserved.

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.

https://t.me/learningnets

Summary

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

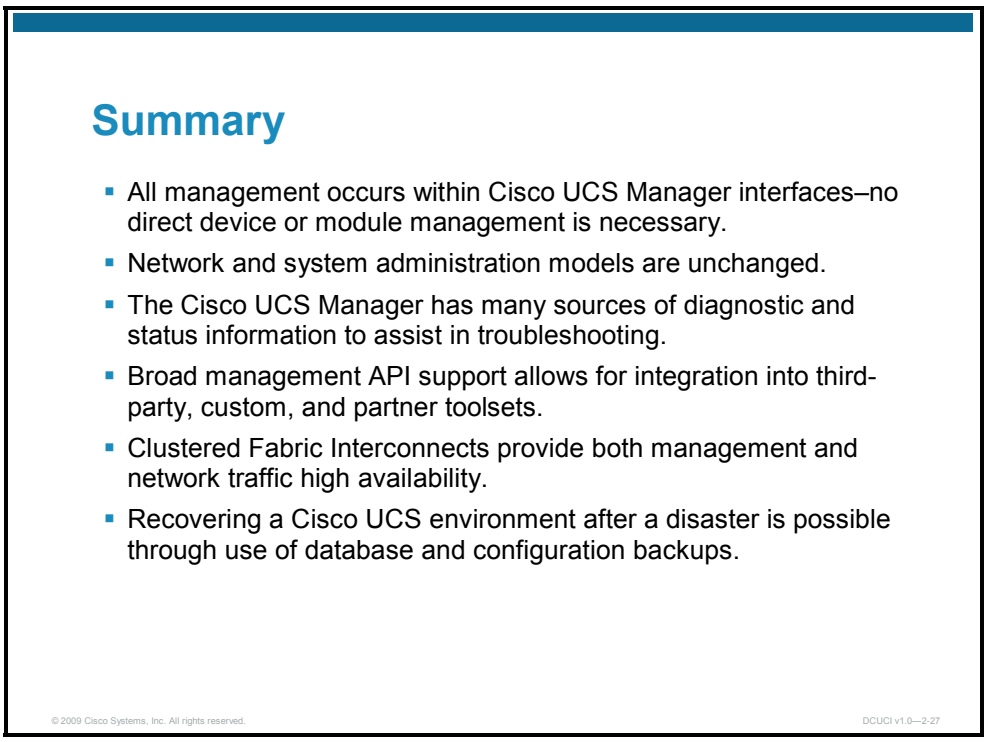

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Installing Cisco UCS Hardware

Overview

This lesson describes the interfaces, functions, and installation practices of Cisco Unified Computing System (Cisco UCS) hardware to provide knowledgeable answers to questions related to these system component topics. You must be familiar with the system components before installing these components in a Cisco UCS environment.

Objectives

Upon completing this lesson, you will be able to explain the process used to install and power up Cisco Unified Computing System hardware. This includes being able to meet these objectives:

- Describe the installation guidelines for the Cisco UCS platform
- Describe the power supply configuration options for the Cisco UCS 5108 chassis
- Describe how to install the various modules of a Cisco 5108 Chassis

https://t.me/learningnets

Installation Guidelines

This topic describes the installation guidelines for the Cisco UCS platform.

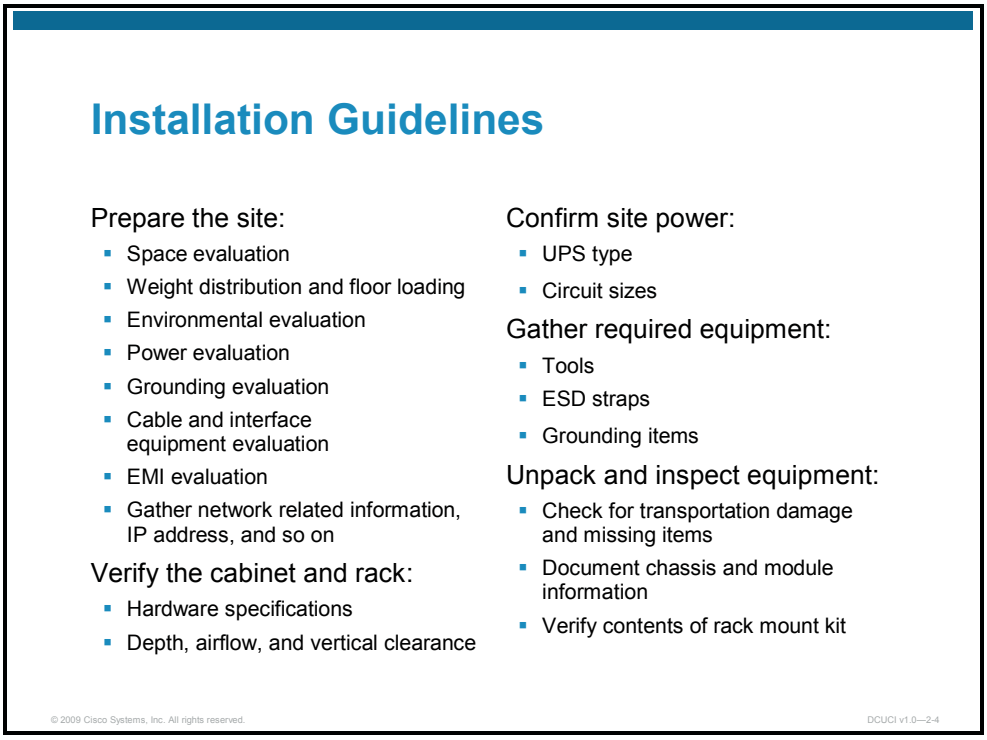

Follow these guidelines when installing the Cisco UCS:

- Plan your site configuration and prepare the site before installing the hardware. It is recommended that you use the site planning tasks listed in the Cisco UCS Hardware Installation Guides.
- Ensure there is adequate space around the switch to allow for servicing the switch and for adequate airflow.
- Ensure the air conditioning meets the heat dissipation requirements listed in the Cisco UCS Hardware Installation Guides.
- Ensure the cabinet, or rack, meets the requirements listed in the Cisco UCS Hardware Installation Guides.

Note	Jumper power cords are available for use in a cabinet.

- Ensure the site power meets the power requirements listed in the Cisco UCS Hardware Installation Guides. If available, you can use an uninterruptible power supply (UPS) to protect against power failures.
- Ensure circuits are sized according to local and national codes.

Note	To prevent loss of input power, ensure that the total maximum loads on the circuits supplying power are within the current ratings of the wiring and breakers.

- Record your installation and configuration information as you work. See *Site Planning and Maintenance Records* in the *Cisco UCS Hardware Installation Guides*.

https://t.me/learningnets

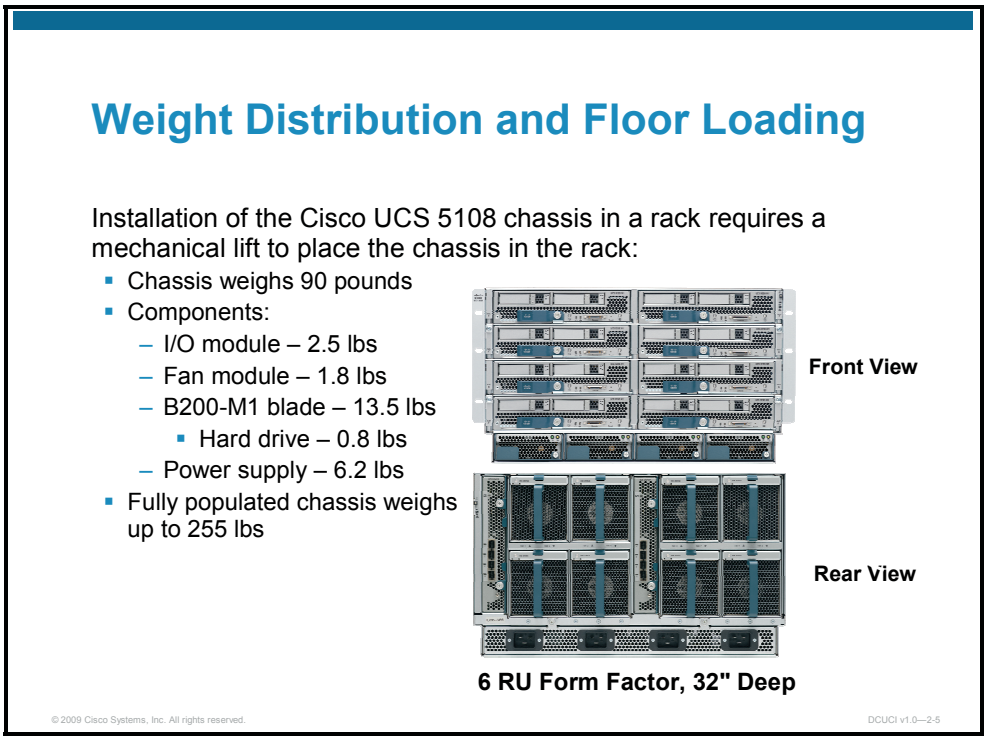

Installation of the Cisco UCS 5108 Chassis in a rack requires a mechanical lift to place the chassis in the rack. Make sure you have access to the lift during the installation process. A fully loaded Cisco UCS 5108 can weigh about 255 pounds. Installations with multiple Cisco UCS 5108 chassis per rack should include a floor-loading assessment to evaluate the static-load rating of the flooring as part of the site evaluation process. For additional information about floor loading requirements, consult the UL document *GR-63-CORE, Network Equipment-Building System (NEBS) Requirements: Physical Protection.*

Chassis Weight

An unpopulated Cisco UCS 5108 chassis weighs 90 pounds (40.83 kg).

Module Weights

The following components are listed with their weights:

- I/O module: 2.5 pounds (1.13 kg)
- Fan module: 1.8 pounds (.82 kg)
- B200-M1 blade: 13.5 pounds (6.12 kg)
- Hard drive: .8 pounds (.36 kg)
- Power supply: 6.2 pounds (2.81 kg)

https://t.me/learningnets

Fabric Interconnect

- **UCS 6120XP– 35 lbs**
  - 2 power supplies
  - 2 fan modules
  - 1 expansion module
  - 1 RU form factor
  - 30” depth

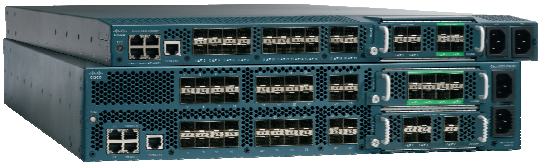

- **UCS 6140XP– 50 lbs**
  - 2 power supplies
  - 5 fan modules
  - 2 expansion modules
  - 2 RU form factor
  - 30” depth

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS 6120XP and Cisco UCS 6140XP can weigh up to 35 or 50 pounds (15.9 kg or 22.7 kg) respectively.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cabinet and Rack Considerations

- Supported racks and cabinets:
  - Open EIA rack
  - Perforated or solid-walled EIA cabinet
  - Requires square-hole posts
- Airflow:
  - All Cisco UCS components use front-to-back airflow
  - Minimum of 6 inches (15.2 cm) of clearance between chassis is recommended in open racks
  - Minimum of 2.5 inches (6.4 cm) of distance between the chassis air vents and any walls is required

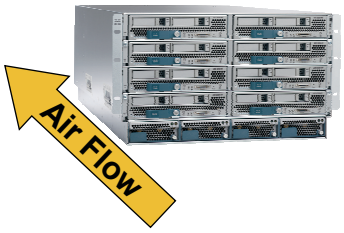

© 2009 Cisco Systems, Inc. All rights reserved.

Cabinet and rack considerations include weight, size, and airflow.

https://t.me/learningnets

Configuring Chassis Power Supplies

This topic describes the power supply configuration options for the Cisco UCS 5108 chassis.

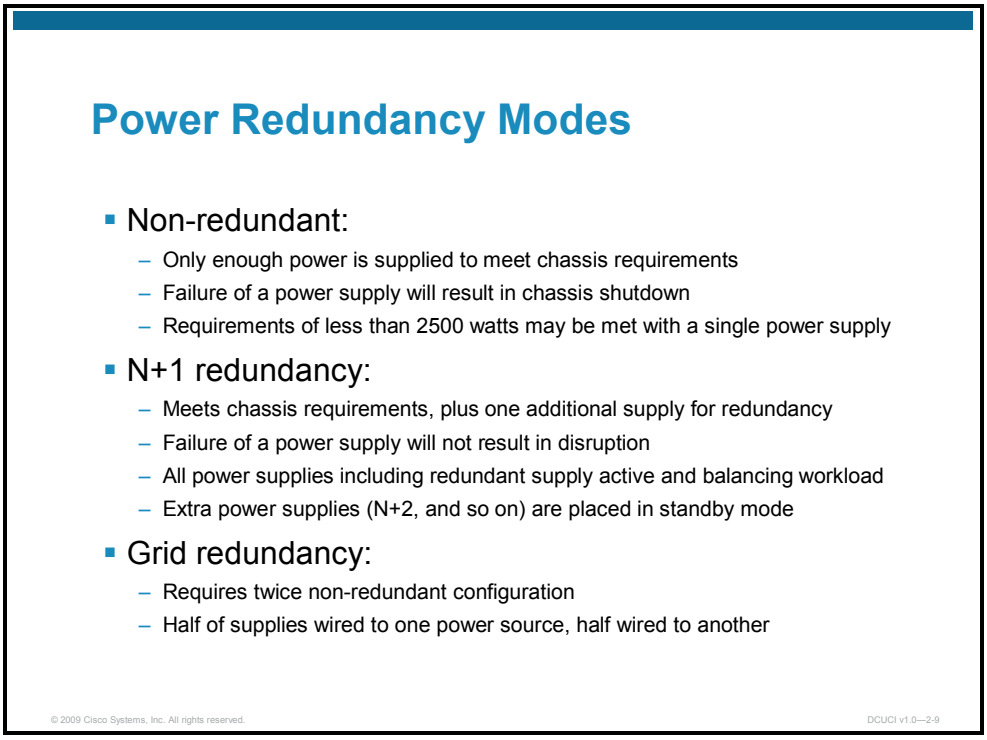

Cisco UCS supports multiple modes of power supply redundancy.

A non-redundant configuration has only enough capacity to support the running system with all installed power supplies active. If any power supply fails, some or all components may be affected.

In an N+1 redundancy configuration, Cisco UCS can tolerate the failure of any single power supply. This configuration requires one more power supply than the non-redundant configuration.

Grid redundancy requires twice the non-redundant configuration, but when properly cabled can survive the loss of an entire power grid within the data center.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Power Supply Connectivity

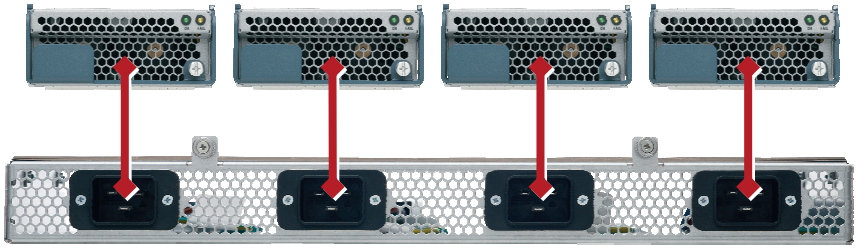

- Inbound power connections are not matrixed:
  - Failure of input power lead will cause loss of power on associated power supply

© 2009 Cisco Systems, Inc. All rights reserved.

The power inlet ports on the Cisco UCS 5108 chassis are not matrixed – that is, there are no crossover connections between the inlet ports and the chassis power supplies. If any component in the path to the power supply fails, that power supply will not receive power.

https://t.me/learningnets

Power Supply Connectivity (Cont.)

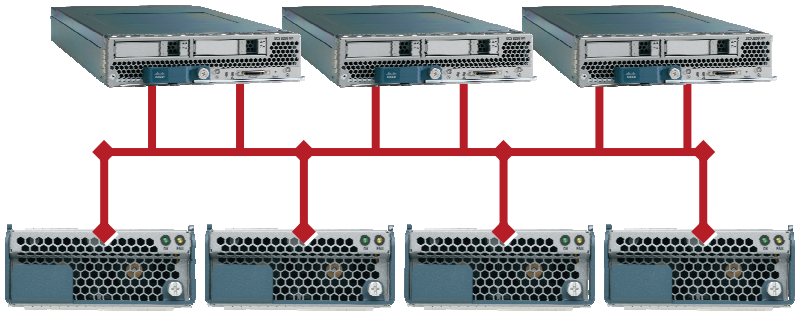

- Outbound power connections are matrixed:
  - All active power supplies pool available power for all components

© 2009 Cisco Systems, Inc. All rights reserved.

The power supply outputs are matrixed within the Cisco UCS 5108 chassis. This means that a failure of any one power supply does not affect any other component, assuming a properly redundant configuration.

https://t.me/learningnets

Chassis Module Installation

This topic describes how to install the various modules of a Cisco 5108 Chassis.

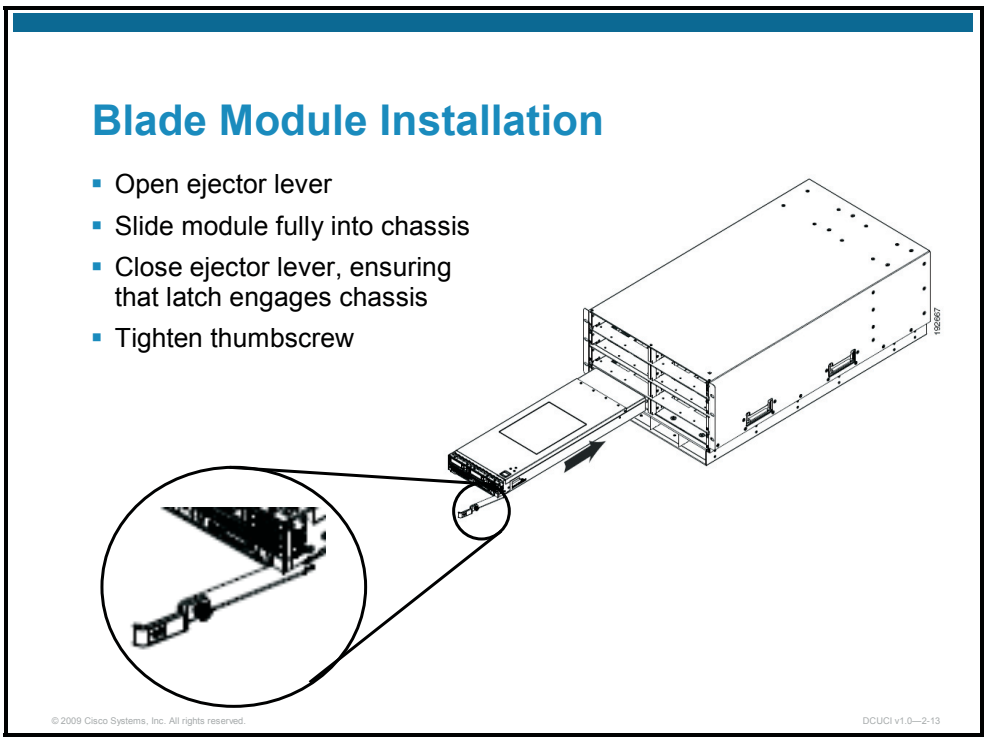

To install a blade module, slide the blade fully into the chassis, close the ejector lever, and tighten the thumbscrew. Removal is the reverse of installation.

https://t.me/learningnets

Blade Hard Drive Installation

- Press release button
- Extend lever
- Slide hard drive into bay
- Close lever

1 – Release Button
2 – Lever

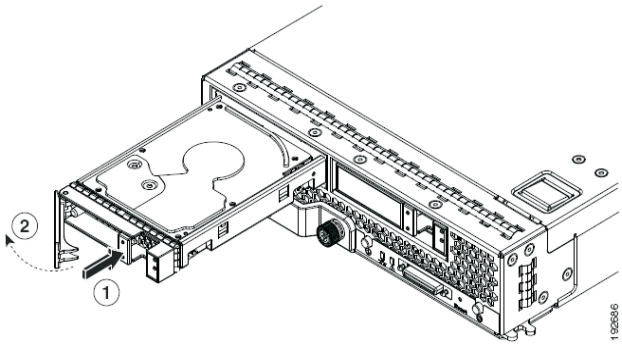

© 2009 Cisco Systems, Inc. All rights reserved.

To install a blade hard drive, press the release button and extend the lever. Slide the hard drive into the bay and close the lever. Removal is the reverse of installation.

https://t.me/learningnets

I/O Module Installation

- Loosen thumbscrews
- Lift both levers
- Slide module into chassis
- Close both levers to seat modules
- Tighten thumbscrews

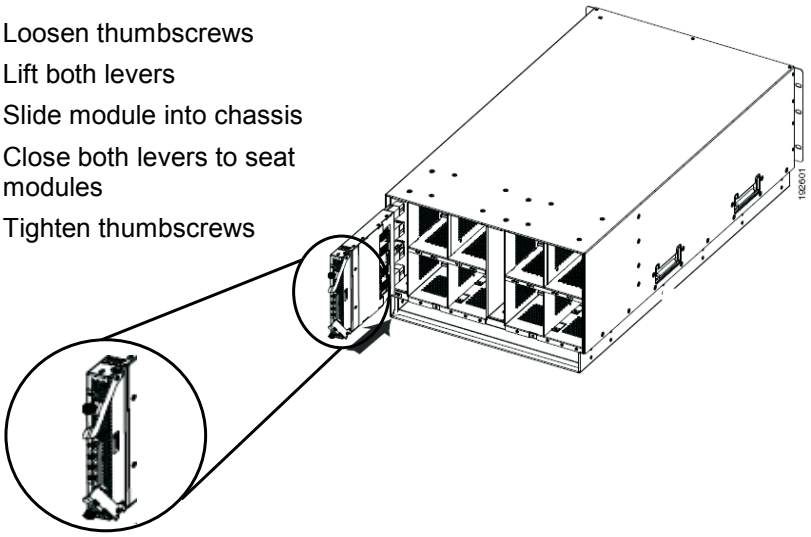

© 2009 Cisco Systems, Inc. All rights reserved.

To install an I/O module, loosen the thumbscrews, left both ejector levers, and slide the module into the chassis. Close both levers and tighten the thumbscrews. Removal is the reverse of installation.

https://t.me/learningnets

Fan Module Installation

- Slide module into chassis until spring latch snaps into place

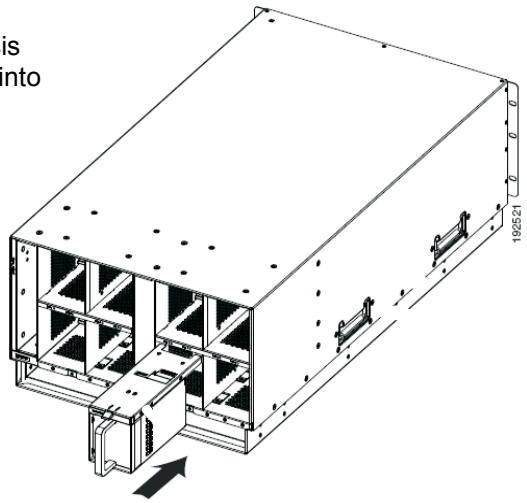

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows how to install a fan module.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

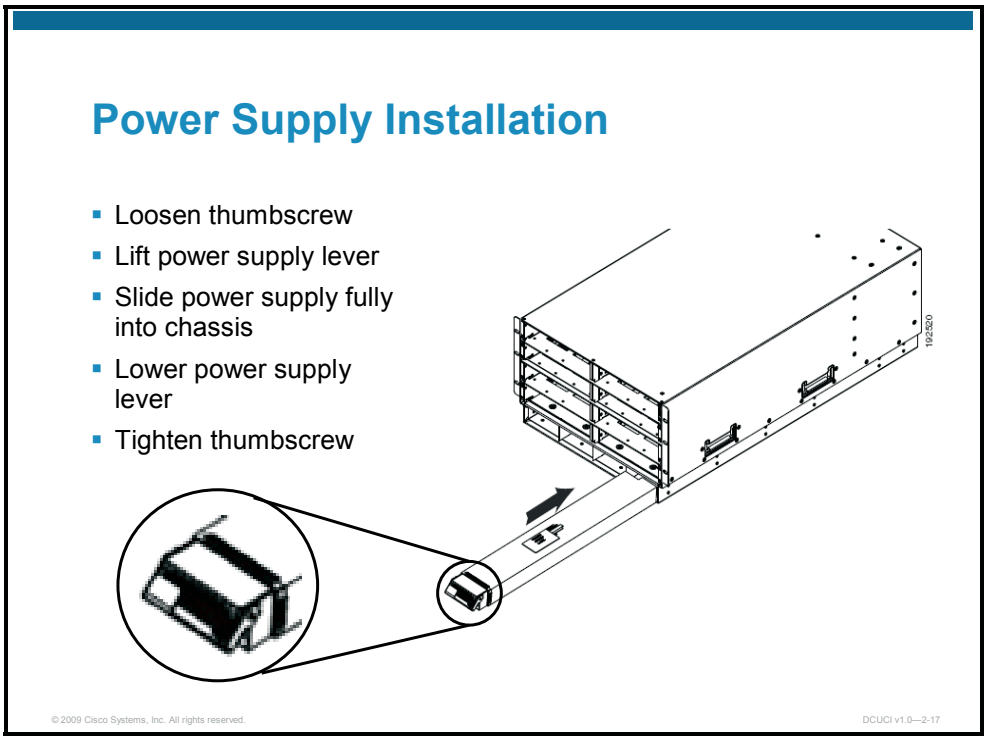

The figure shows how to install a power supply.

https://t.me/learningnets

Summary

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

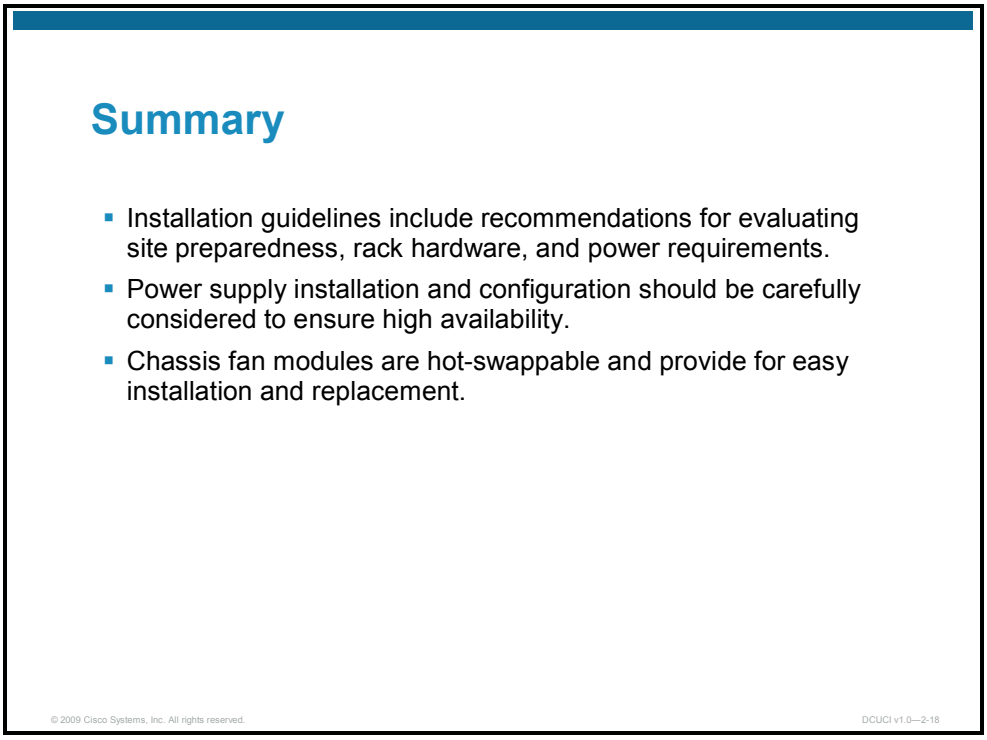

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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

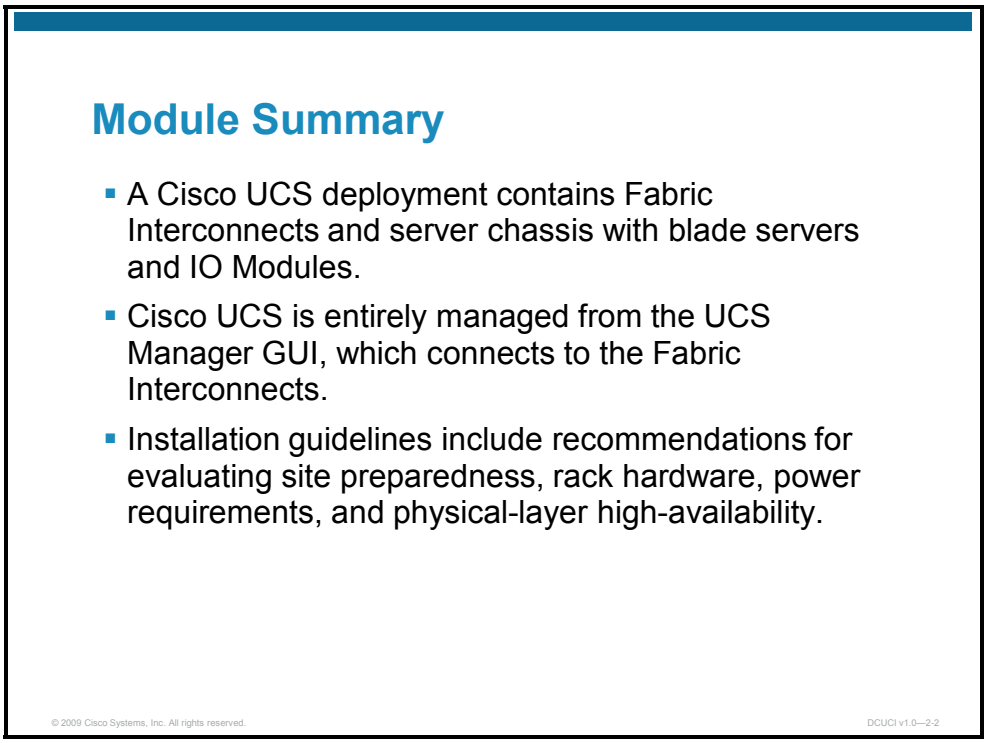

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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 of the following is a valid combination of compute nodes in a Cisco UCS 5108 Chassis? (Source: Defining Hardware Components)

A) 2 UCS B250-M1, 3 UCS B200-M1
B) 3 UCS B250-M1
C) 4 UCS B250-M1 2 UCS B200-M1
D) 10 UCS B200-M1

Q2) Which of these tasks is the responsibility of the UCS Administrator job role? (Source: UCS Architectural Overview)

A) Layer 3 configuration
B) Fibre Channel zoning
C) Service Profile creation
D) Operating System driver updates

Q3) If two power supplies are required to fully power a Cisco UCS 5108 chassis with a given configuration, how many power supplies are required to provide N+1 redundancy? (Source: Installing UCS Hardware)

A) Two
B) Three
C) Four
D) Five

https://t.me/learningnets

Module Self-Check Answer Key

Q1) B

Q2) C

Q3) B

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 3

Understanding Cisco UCS Connectivity

Overview

This module describes the connectivity of the Cisco Unified Computing System (Cisco UCS) platform.

Module Objectives

Upon completing this module, you will be able to describe the connectivity of the Cisco Unified Computing System platform. This includes being able to meet these objectives:

- Describe unified connectivity
- Explore the unified computing user interfaces
- Describe Cisco Unified Computing System startup and shutdown procedures
- Upgrading software in the Cisco Unified Computing System

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Describing Unified Connectivity

Overview

In this lesson you will learn how to describe unified connectivity.

Objectives

Upon completing this lesson, you will be able to describe unified connectivity. This includes being able to meet these objectives:

- Describe hardware component connectivity
- Describe the connectivity between Fabric Interconnects for high availability
- Describe the Fabric Interconnect network configuration

https://t.me/learningnets

Hardware Component Connectivity

This section describes connectivity to the various hardware components of the Cisco Unified Computing System (Cisco UCS).

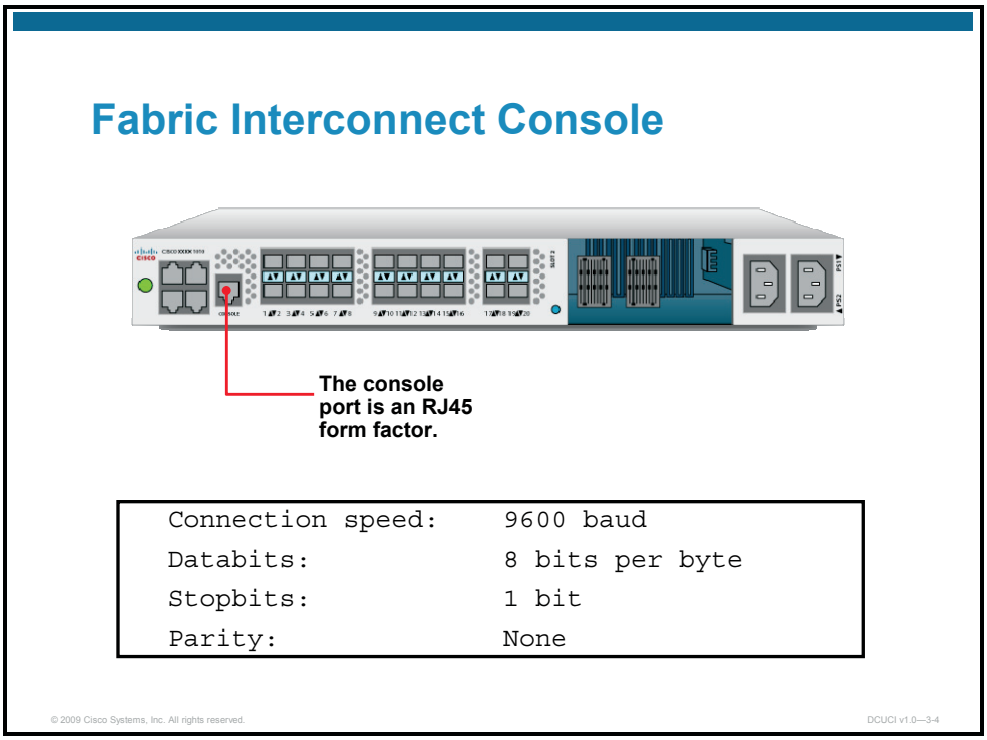

The console terminal of the Cisco UCS 6100 Fabric Interconnect is accessed using the RJ45-based serial console port on the rear of the switch. The figure illustrates the port location.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Fabric Interconnect Management

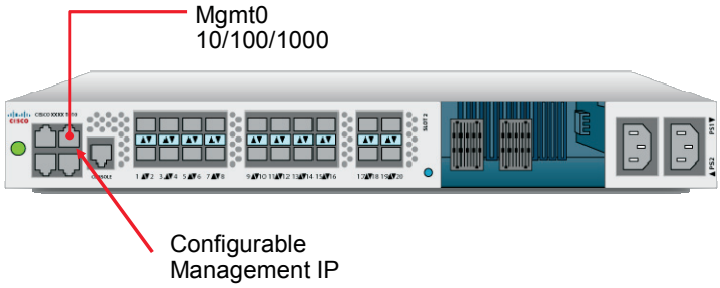

Management port is RJ45 form factor.

© 2009 Cisco Systems, Inc. All rights reserved.

Out of band management is performed using the management port located on the rear of the Cisco UCS 6100 chassis. Once this port has been assigned an IP address, the console can be accessed from a remote workstation. The speed of the interfaces is 10/100/1000.

https://t.me/learningnets

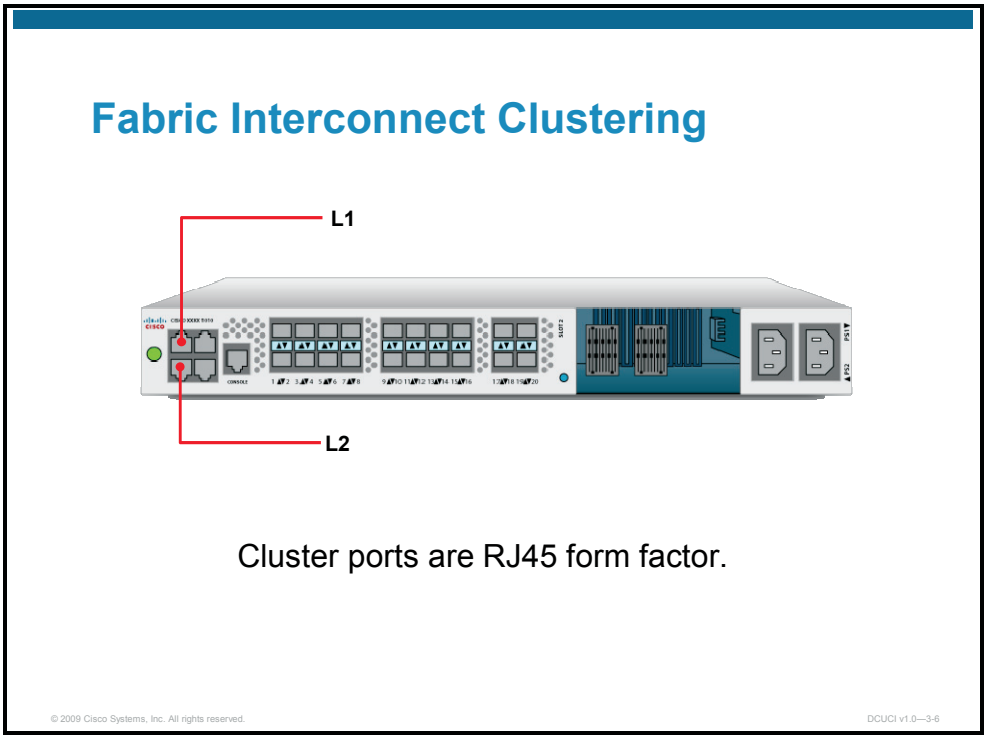

Cluster Interfaces

These interfaces provide a cluster link between two Cisco UCS 6100 Fabric Interconnects. They carry the cluster heartbeat messages between the two Fabric Interconnects as well as high level messages between Cisco UCS Manager elements. These links are part of an 802.3ad bond and is managed by the underlying operating system. The bond is configured to run Link Aggregation Control Protocol (LACP) which brings the bond link up 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 ports are labeled L1 and L2 on the Fabric Interconnect and operate at 10/100/1000 speeds. As discussed in a later lesson, clustered Fabric Interconnects must be cabled from L1-to-L1 and L2-to-L2.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

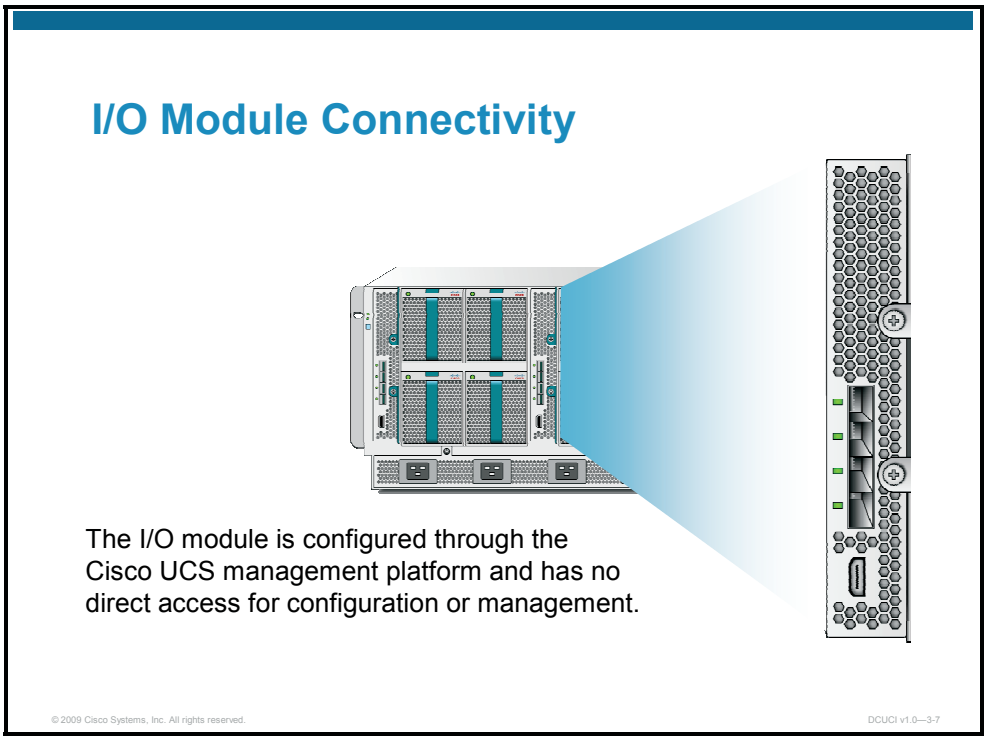

IOM Management

The I/O module (IOM) itself is not an independent point of management (described in a later lesson).

https://t.me/learningnets

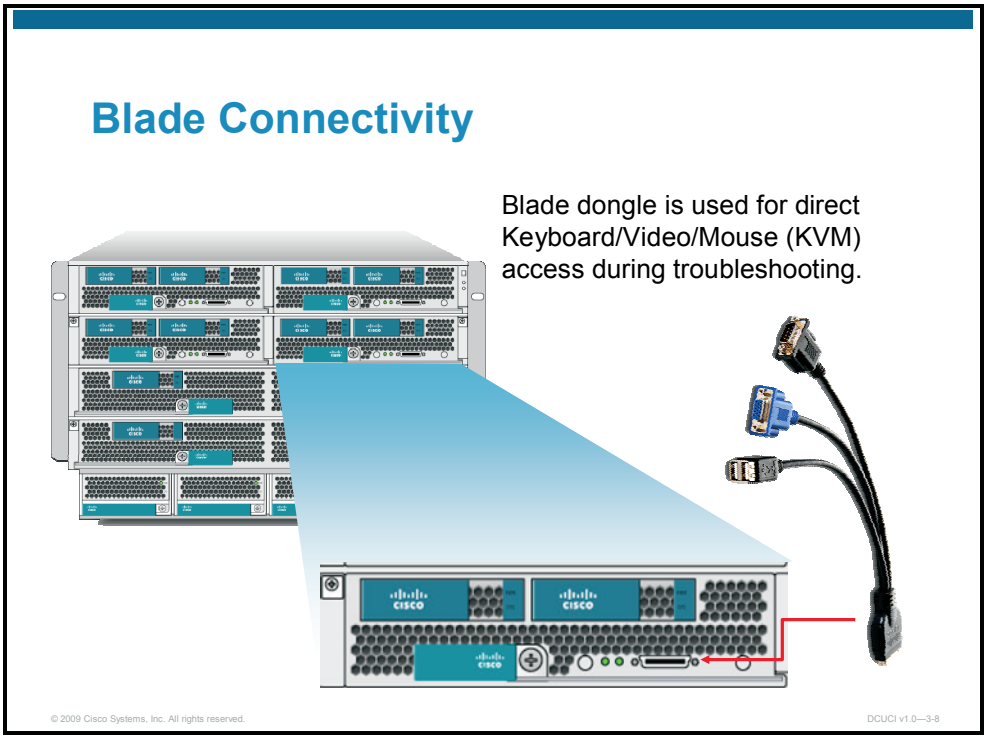

The dongle is used for troubleshooting the host system and includes a DB9 serial connector, VGA connector, and a dual-USB connector. A VGA monitor and USB keyboard/mouse can be connected to provide access to the individual blade server host system.

https://t.me/learningnets

Cisco UCS Chassis

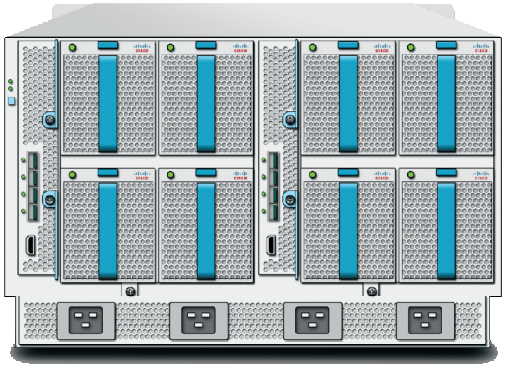

The blade chassis is managed from the central Cisco UCS management system and has no direct management points.

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS 5108 Blade Server Chassis is internally connected to the blades via a midplane. This is where the internal management and data networks are established.

Externally, the chassis has four power connection points for 220V power. There are no management points or connections within the blade chassis.

https://t.me/learningnets

High Availability Connectivity

This section describes the connectivity between redundant Fabric Interconnects in a high-availability configuration.

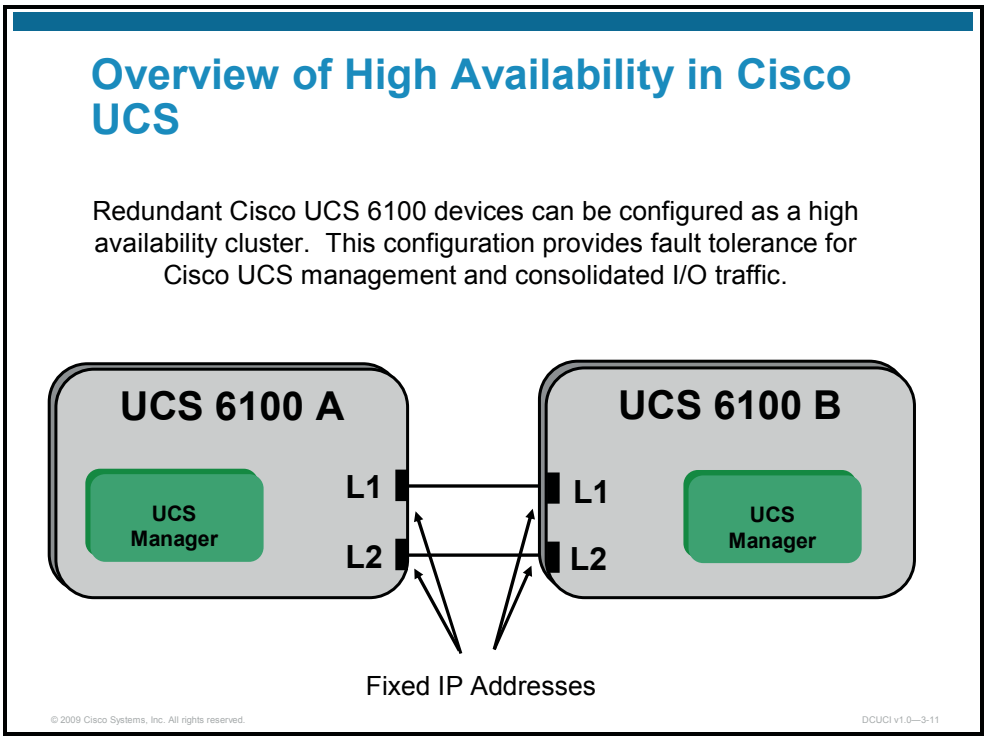

Configuring two Fabric Interconnects in your Cisco UCS system can provide redundancy for both the Cisco UCS Manager functionality and the switching functionality that the Cisco UCS Fabric Interconnect provides.

High Availability Requirements

- 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 1-4 uplinks
- Connect Fabric Interconnect B to IOM B of each Chassis using 1-4 uplinks

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

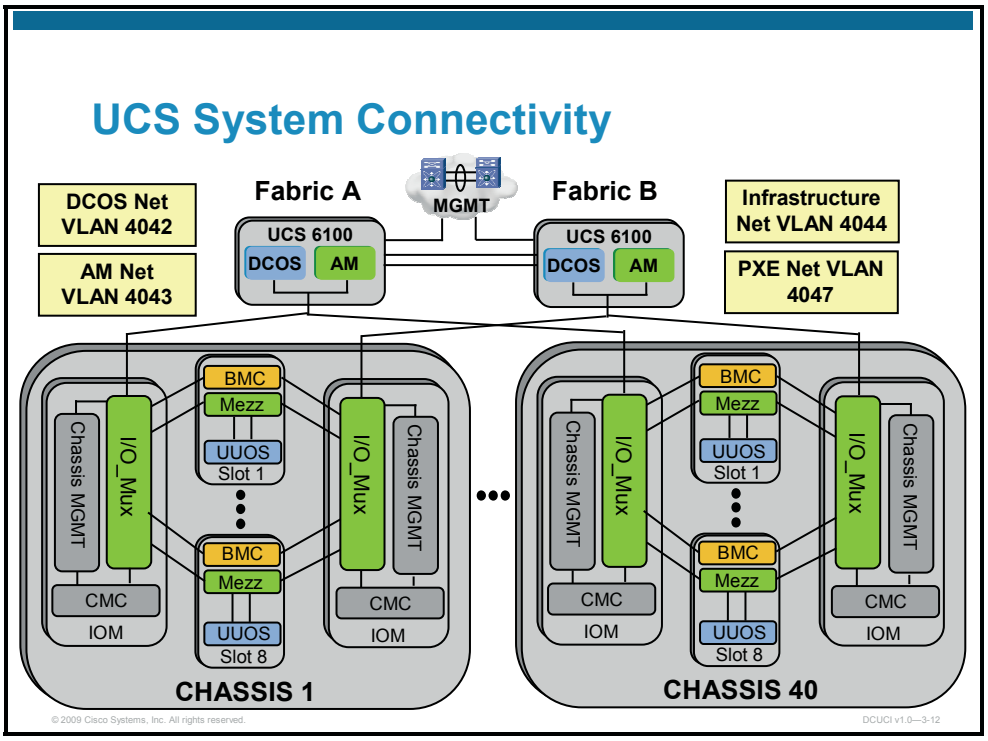

This figure shows the connectivity paths for Cisco UCS management systems.

Fabric A/B Datacenter Operating System (DCOS) Networks

These networks provide connectivity between the DCOS (of Cisco UCS Manager) instance in each Fabric Interconnect, Fabric A/B chassis management controllers (CMC), and the management entities on mezzanine cards. The network traffic includes Layer 2 protocols, such as the ones used to bootstrap the IOM and the Virtual Interface Card (VIC) protocol used to support Network Interface Virtualization (NIV).

Frames on this network are tagged for VLAN 4042.

Fabric A/B Adapter Management (AM) Networks

These networks provide connectivity between the DCOS (of 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 Managers for general adapter functionality, and Layer 3 (IP) is used for remote login to the adapter operating system.

Frames on this network are tagged with a VLAN ID of 4043 on the adapter and the tagging is carried through the Cisco UCS Fabric interconnect. IP addresses are algorithmically determined by Cisco UCS Manager and passed to the adapter during bootstrapping.

Fabric A/B Adapter Management (AM) Infrastructure Networks

These networks provide UCS-wide connectivity between the AM instance in the Cisco UCS Fabric Interconnect, the Fabric A/B CMC, and each blade's baseboard management controller (BMC). The Cisco UCS Manager bootstraps the process by passing out chassis numbers to each CMC as they are discovered and integrated into the system. Each element then derives its own addresses algorithmically. The interfaces in the CMC and BMC are exposed as virtual interfaces within VLAN 4044, and this VLAN ID is carried throughout each fabric.

https://t.me/learningnets

Because of the way this network is connected, traffic between a CMC and the IBMC in the same chassis is forwarded via the UCS Fabric Interconnect. This allows the CMC IP interface to be brought up early in the chassis association process to facilitate CMC image download, if required.

Cisco UCS Fabric Interconnects are statically assigned to be the head of Fabric A or Fabric B during the Cisco UCS Manager cluster configuration. This configuration is stored persistently.

Fabric A/B Cisco UCS Manager PXE Networks

These networks are isolated networks used to boot host operating systems from images hosted by the Cisco UCS Manager. IP addresses in this set are passed out through DHCP as part of the Pre-eXecution Environment (PXE) process. Frames from the host are tagged, with a default VLAN assignment of 4047.

This Layer 2 segregation using VLANs allows traffic isolation and limits the required IP address scopes.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Fabric Interconnect Network Configuration

This section describes the Fabric Interconnect network configuration.

Profile Prerequisites: Networking

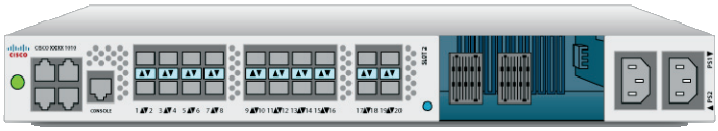

- VLANs:
  - Configure globally within Cisco UCS
  - Out-of-box configured with default VLAN (VLAN 1)
  - Trunk all VLANs defined in profiles across uplinks
- Port Channels:
  - Port Channel uplinks created using Cisco UCS Manager
  - Must match port channel configuration upward from Cisco UCS
  - Uses LACP 802.3ad / 802.1ax

© 2009 Cisco Systems, Inc. All rights reserved.

You can configure Port Channels on the uplink ports from the Cisco UCS to the next system. This is configured using the Cisco UCS Manager.

https://t.me/learningnets

Global VLAN Configuration

- Create VLAN under “LAN Cloud”:
  - Will work across entire Cisco UCS cloud
- Optionally create a per Fabric Interconnect VLAN

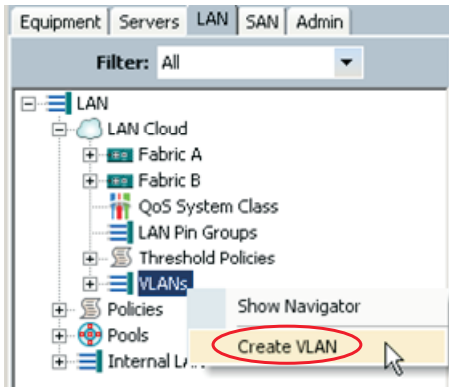

© 2009 Cisco Systems, Inc. All rights reserved.

For any VLAN that you want to support on any blade servers, a VLAN object must be created in the global Cisco UCS configuration (under the LAN tab on the navigation panel, as shown in the example).

You can also create VLANs that are associated with only Fabric Interconnect A or Fabric Interconnect B, but this is typically not recommended.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Support for New VLAN

- Provide name and VLAN ID
- Values need not match – abstraction is provided by the Cisco UCS Manager

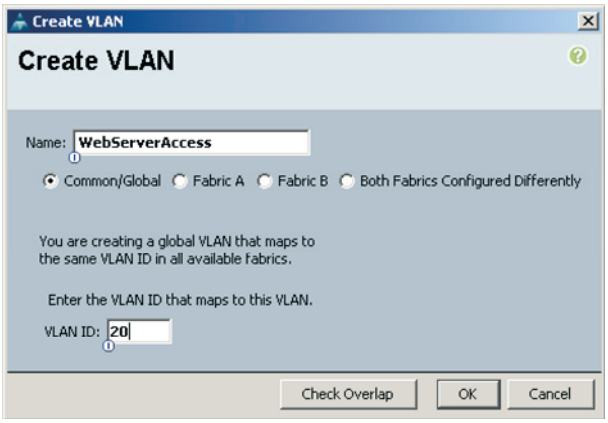

© 2009 Cisco Systems, Inc. All rights reserved.

This example shows the creation of a global VLAN with an ID of 20. The name is used to reference this object from within the Cisco UCS Manager. There need not be a specific relationship between the name and the VLAN ID.

https://t.me/learningnets

Creating Support for New VLAN (Cont.)

Possible naming conventions:

- By function:
  - Name VLAN objects by function rather than number
    - "WebServerAccess", "DatabaseClusterVLAN", and so on
  - Eases VLAN changes in future
- By VLAN number:
  - Match VLAN name to VLAN number
    - "Vlan5", "Vlan155", and so on
  - Simplifies discussions between functional groups
  - Reduces confusion

© 2009 Cisco Systems, Inc. All rights reserved.

There are several possible conventions for naming VLAN objects. Some administrators prefer to abstract the name of the object from the VLAN number and reference it by function. This approach does not require any network-specific knowledge to assign a VLAN, and allows an easy change of the VLAN number in the future if the network topology changes.

Other administrators prefer to keep the name the same as the VLAN number, to reduce confusion as to which VLAN is being referenced. This approach keeps the same naming scheme in use between Cisco UCS and the network teams.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

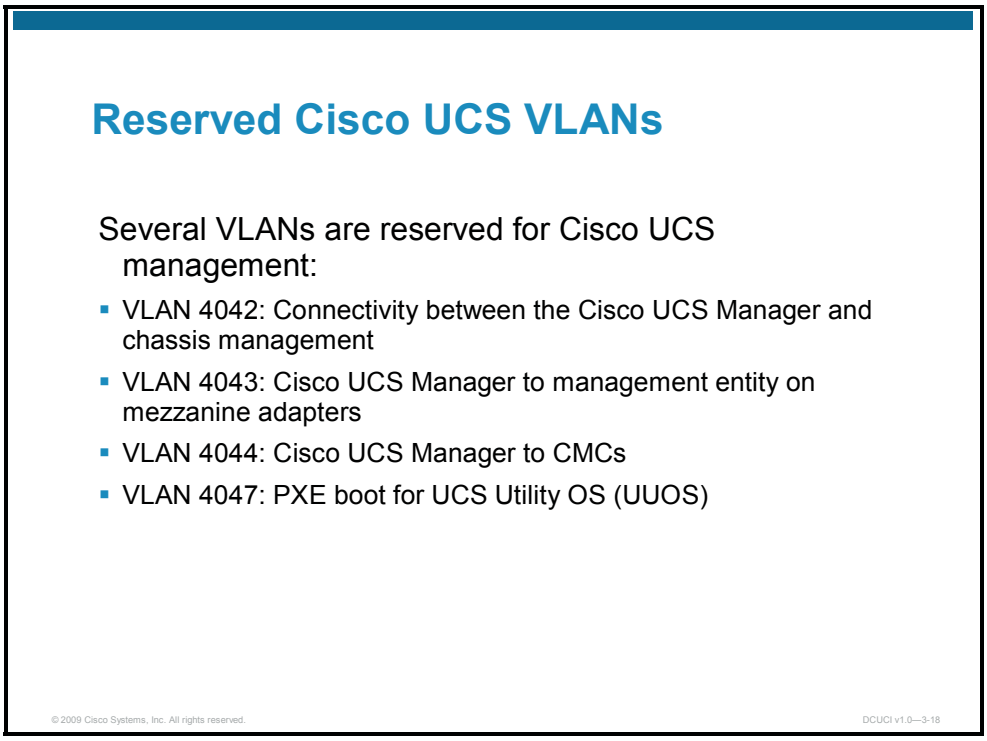

The figure summarizes the reserved VLANs in Cisco UCS, as described in the previous slide.

https://t.me/learningnets

Creating Port Channels

- Always Fabric Interconnect-specific:

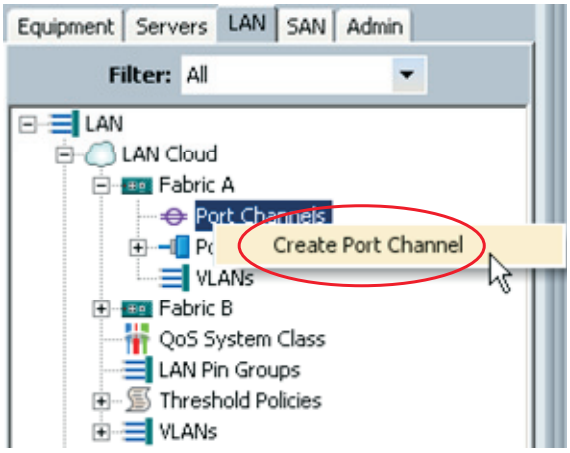

© 2009 Cisco Systems, Inc. All rights reserved.

Port channel creation on the Cisco UCS side is always specific to a particular Fabric Interconnect and specific ports.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Port Channels (Cont.)

- When adding ports, all Fabric Interconnect ports will be shown
- Currently configured server ports will be reconfigured if selected!

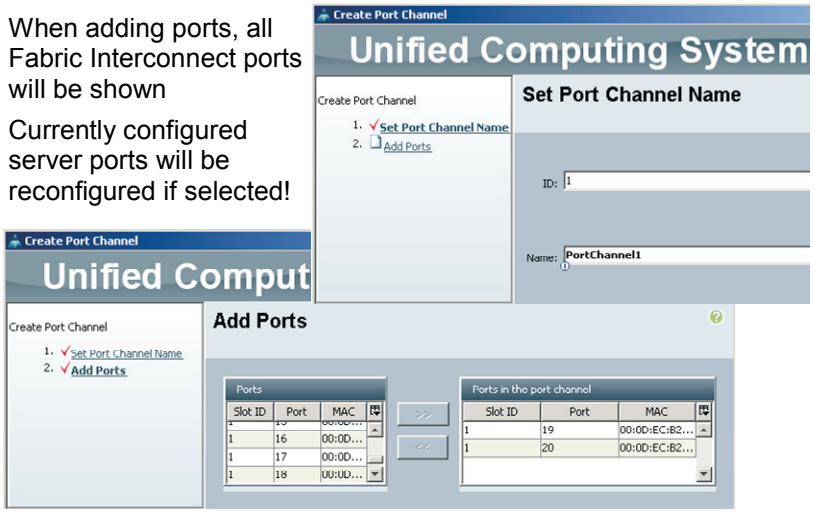

© 2009 Cisco Systems, Inc. All rights reserved.

In this example, a port channel has been created with ID 1. Particular uplink ports (19 and 20) are chosen on a given Cisco UCS Fabric Interconnect. A matching configuration must be configured on the uplink switch. Note that this wizard shows all ports on the Fabric Interconnect regardless of current usage or configuration. A warning is displayed before reconfiguration.

https://t.me/learningnets

Configuring VSANs

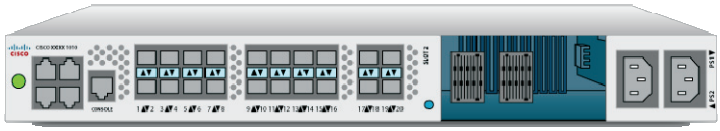

- Similar to VLANs:
  - May declare VSAN:
    - Globally
    - Per Fabric Interconnect
  - Default VSAN 1 is preconfigured and not recommended for data traffic
  - Each uplink from Cisco UCS switch currently supports only one VSAN

© 2009 Cisco Systems, Inc. All rights reserved.

The logic for configuring VSANs is similar to that of configuring VLANs. VSANs may be configured globally or per Fabric Interconnect in the Cisco UCS Manager.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

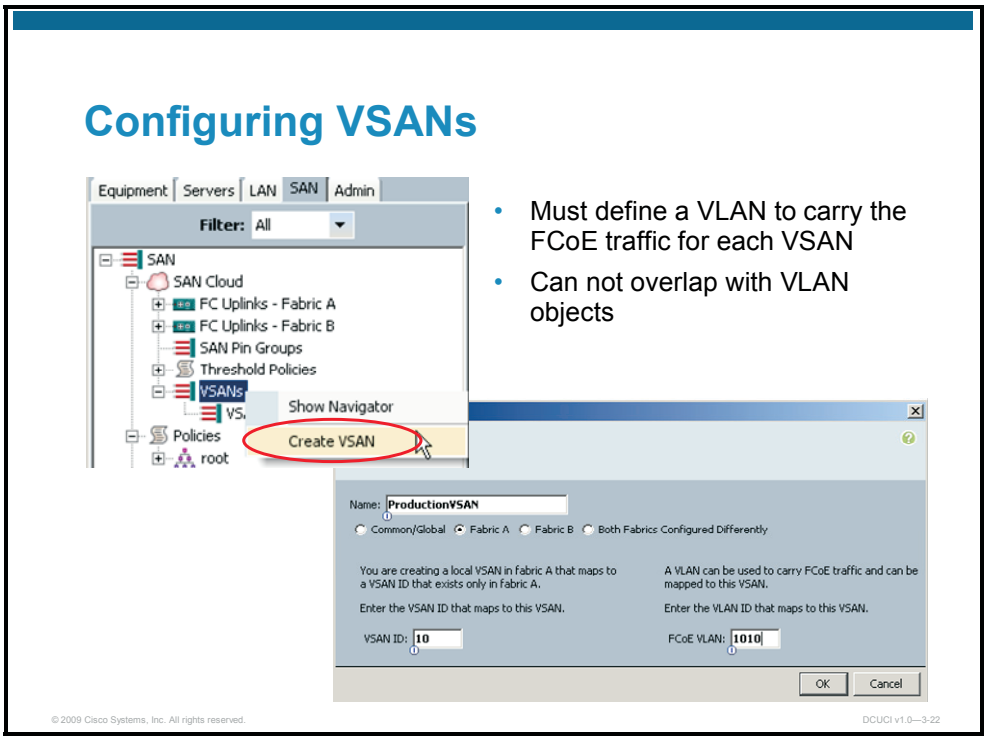

The figure shows the configuration of VSAN ID 10 in the global Cisco UCS SAN configuration. You must define a VLAN to carry the Fibre Channel over Ethernet (FCoE) traffic for this VSAN. The FCoE VLAN selected must not overlap with the VLANs used for traditional Ethernet access. After this is accomplished, that VSAN can be associated with a Cisco UCS uplink connected to that VSAN.

https://t.me/learningnets

FCoE VLANs

- Must define a FCoE VLAN for each VSAN
- All server Fibre Channel traffic is carried via FCoE in dedicated VLANs
- FCoE VLANs must not conflict with Cisco UCS Manager VLAN objects
- Possibly select unused range of VLANs and dedicate that range to FCoE

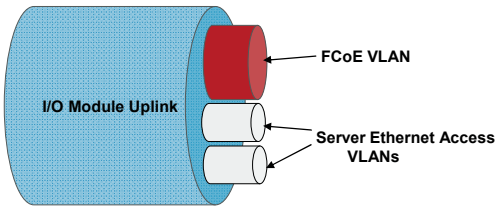

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes FCoE VLAN configuration.

https://t.me/learningnets

Summary

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

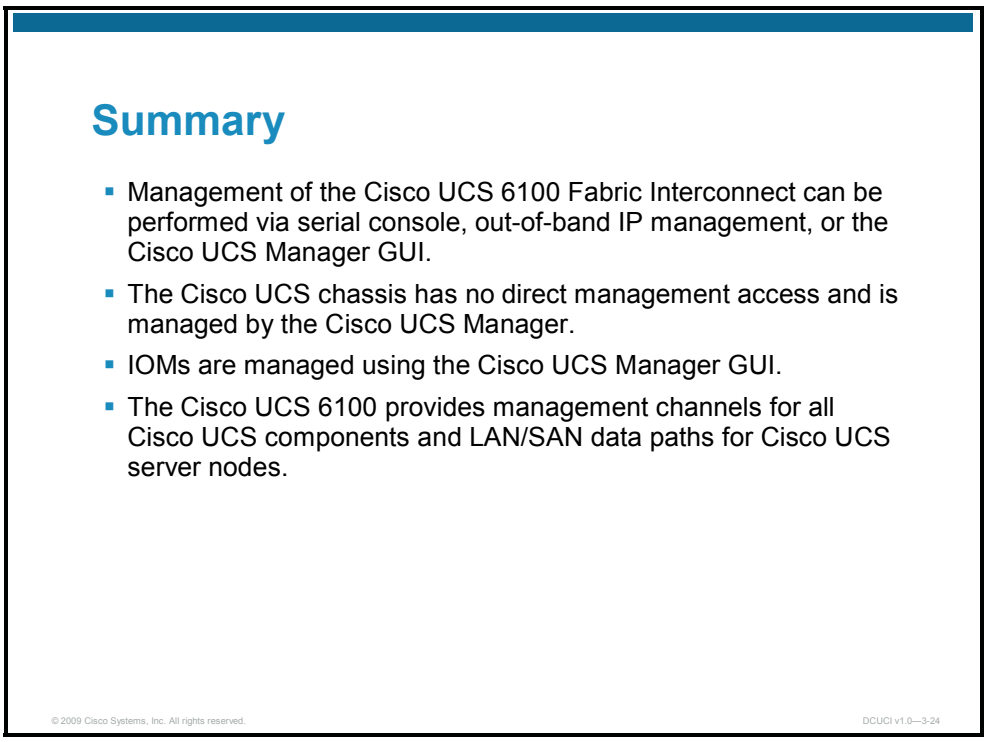

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Exploring the Unified Computing User Interfaces

Overview

This lesson explores the purpose and use of the unified computing user interfaces.

Objectives

Upon completing this lesson, you will be able to explore the unified computing user interfaces. This includes being able to meet these objectives:

- Describe the graphical user interface
- Describe the command-line interface
- Describe the connect commands

https://t.me/learningnets

The Graphical User Interface

This topic describes the graphical user interface (GUI).

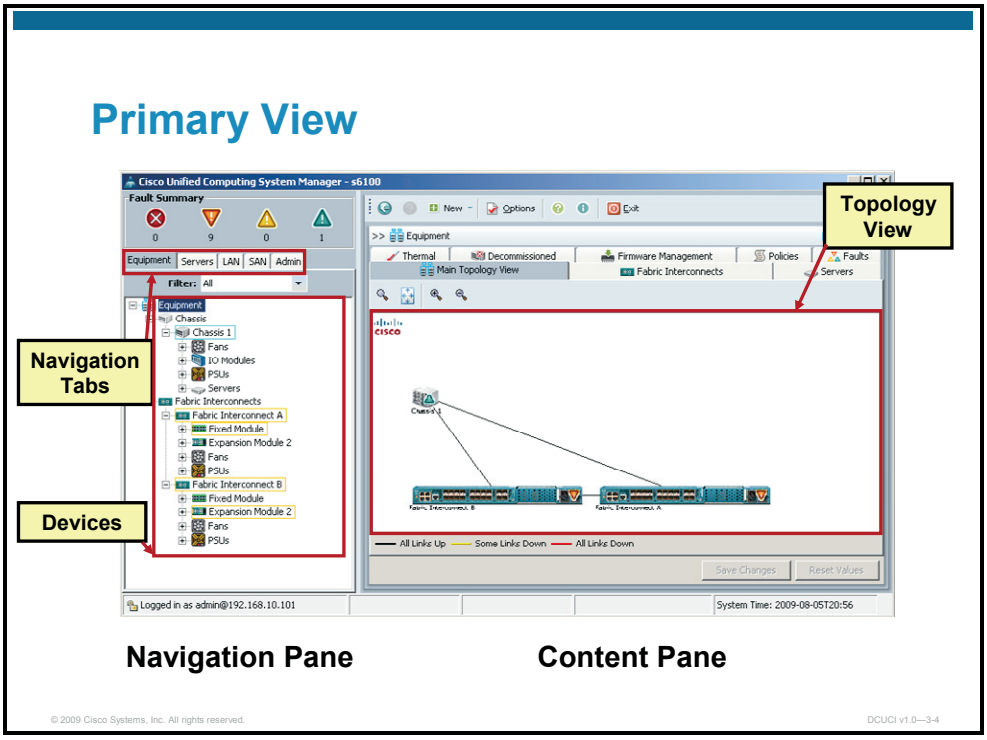

The navigation pane at the left of the screen includes a fault summary bar across the top view, and a series of five tabs that offer different views of the various managed components in the Cisco Unified Computing System (Cisco UCS) system.

The fault summary identifies four conditions: critical, major, minor and warning. The faults summary contains the cumulative totals for the entire Cisco UCS system.

An expandable branch or tree function allows the operator to traverse the various components located within the five tabs.

The content pane at the right of the screen includes a top toolbar with a Back button, a New object creation pull-down menu, Options, questions, and information buttons, and a Debug pull-down menu.

The second toolbar in the content pane offers hierarchies that have already been traversed, providing the ability to rapidly return to a previous screen. At the right-most portion of this bar is the current location.

The largest part of the content pane offers granular details associated with the objects that are highlighted in the navigation pane. And at the very bottom of the content pane are function buttons for committing, saving, or discarding the changes made by the operator.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

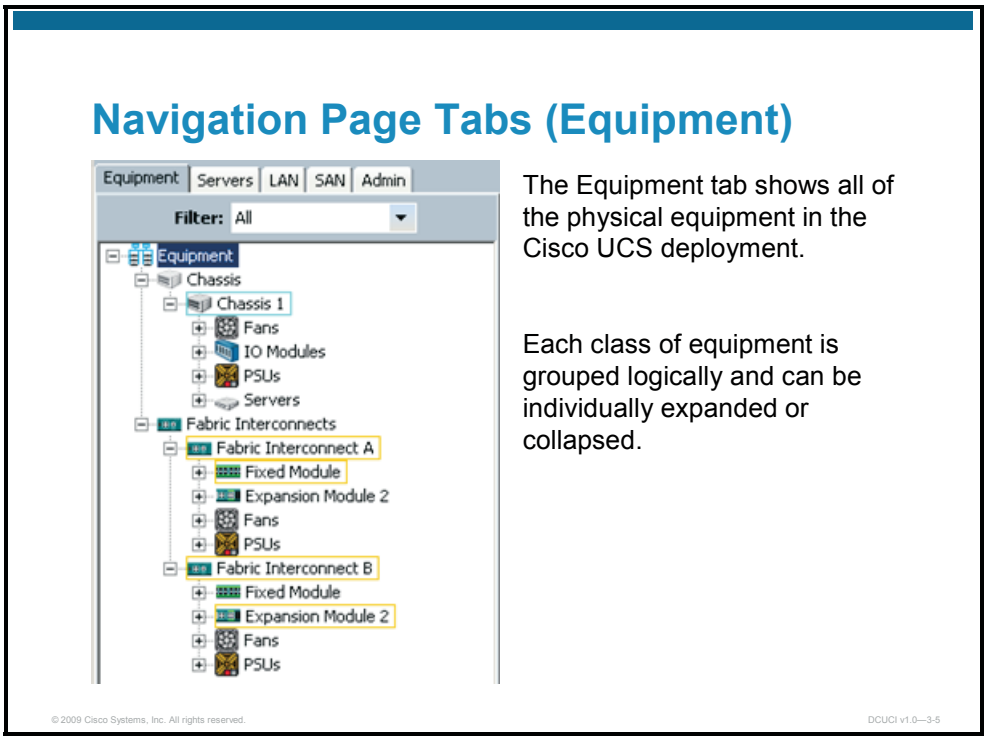

The figure shows the Equipment tab of the navigation pane. The display is in a hierarchical format that lists Chassis 1-40 and then Fabric Interconnect (switch) A and B. Subcomponents such as compute nodes and interface cards can be expanded out from the chassis.

When a hardware object is highlighted, its corresponding attributes are displayed in the content pane.

The content pane can have one or more tabs, depending on the attribute or function of the object under review.

At the top of the navigation pane is a filter pull-down menu for tailoring your view. The filter is set to All by default, but can be set to restrict the object view to the topmost objects in that tab. In the Equipment view, the chassis or switches are the topmost objects.

https://t.me/learningnets

Navigation Page Tabs (Servers)

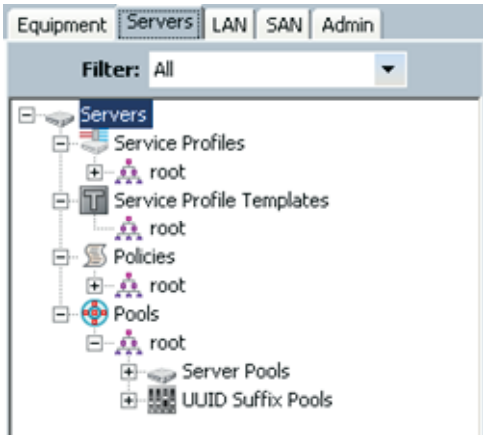

The Servers tab shows Service Profiles, Templates, Policies, and Pools.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the Servers tab. The filter pull-down menu on this tab offers the following selected views:

- All
- Server Profiles
- Service Profile Templates
- Policies
- Pools

The functions associated with this series of logical objects includes the ability to create the following:

- Organization
- Server Profile
- Server Profile Expert
- Service Profile Template
- Pools
- Policies
- Many others

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Navigation Page Tabs (LAN)

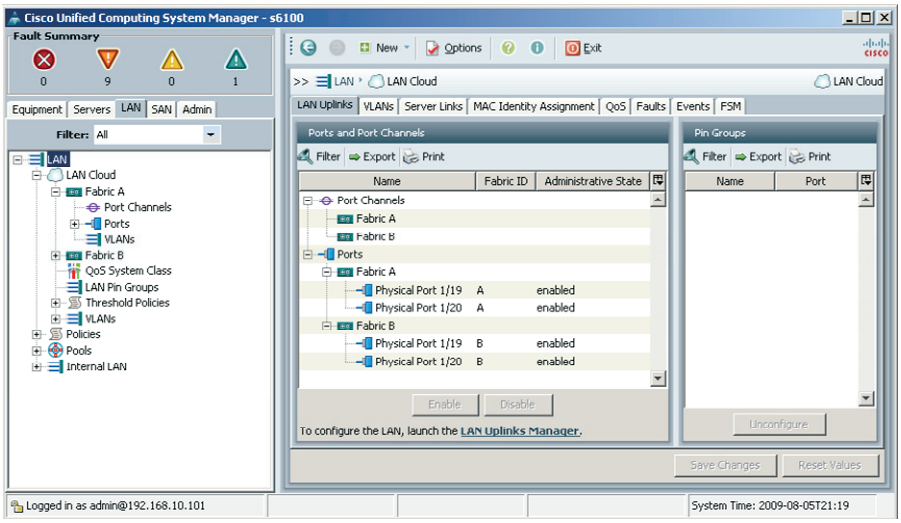

The LAN tab offers the following filterable views:

- All
- Policies
- Pools
- Internal LAN Domains

https://t.me/learningnets

Navigation Page Tabs (SAN)

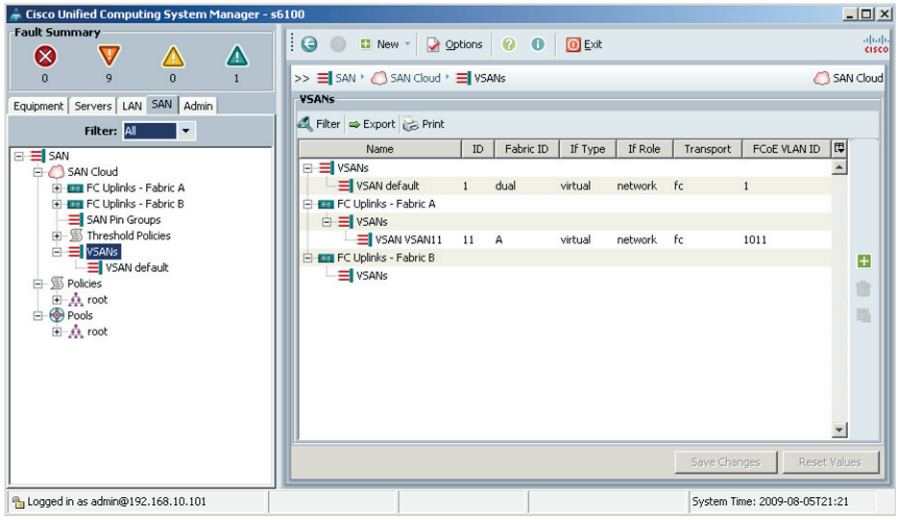

© 2009 Cisco Systems, Inc. All rights reserved.

The SAN tab on the navigation pane offers the following filterable views:

- All
- SAN Cloud
- Policies
- Pools

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Navigation Page Tabs (Admin)

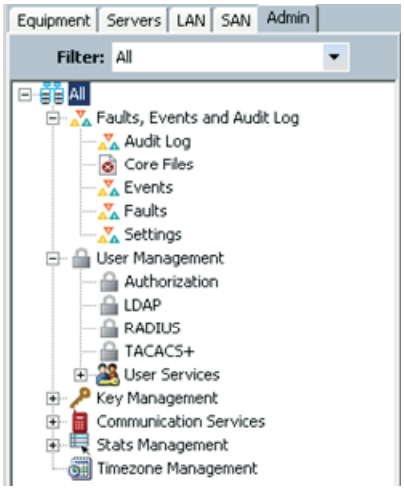

The Admin tab displays system configuration and reporting options.

The filter may be adjusted to reduce the number of options displayed at one time.

© 2009 Cisco Systems, Inc. All rights reserved.

The Admin tab offers nine filter views:

- All
- Fault, Events, Audit Log
- User Management
- Key Management
- Communication Services
- Stats Management
- Timezone Management

https://t.me/learningnets

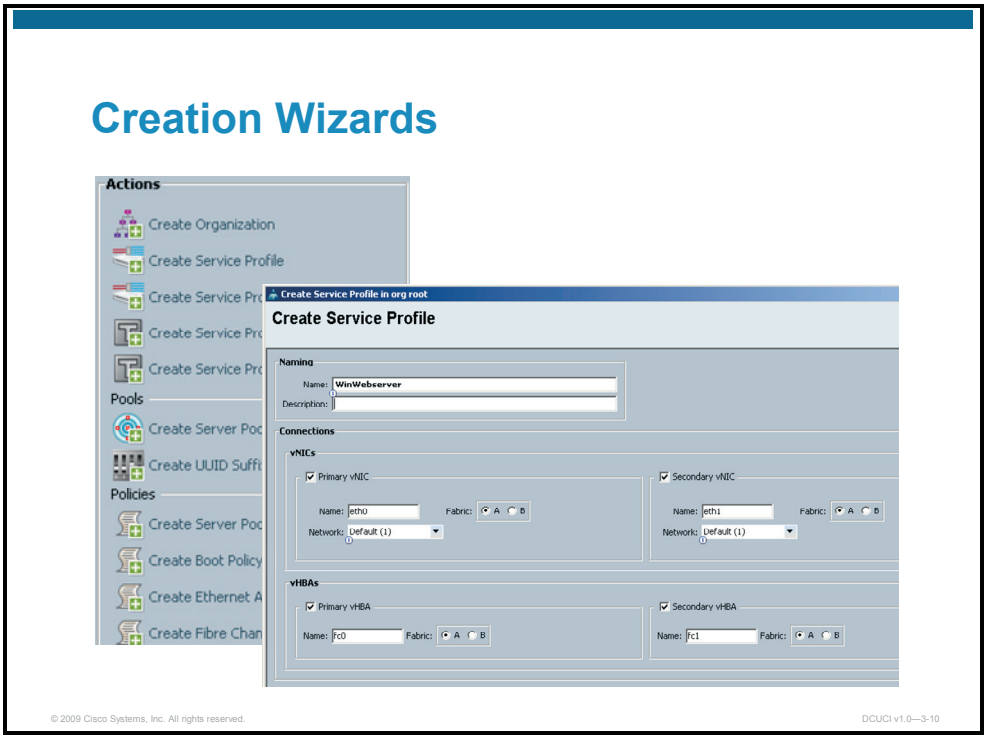

There are numerous creation activities that are assisted by creation wizards. Each requires the definition of necessary attribute values in order to create the requested object.

Required fields are shaded red. Filling in the required fields is usually enough to create the object; however, this information may not be sufficient to render the object functional.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Command-Line Interface

This topic describes the command-line interface (CLI).

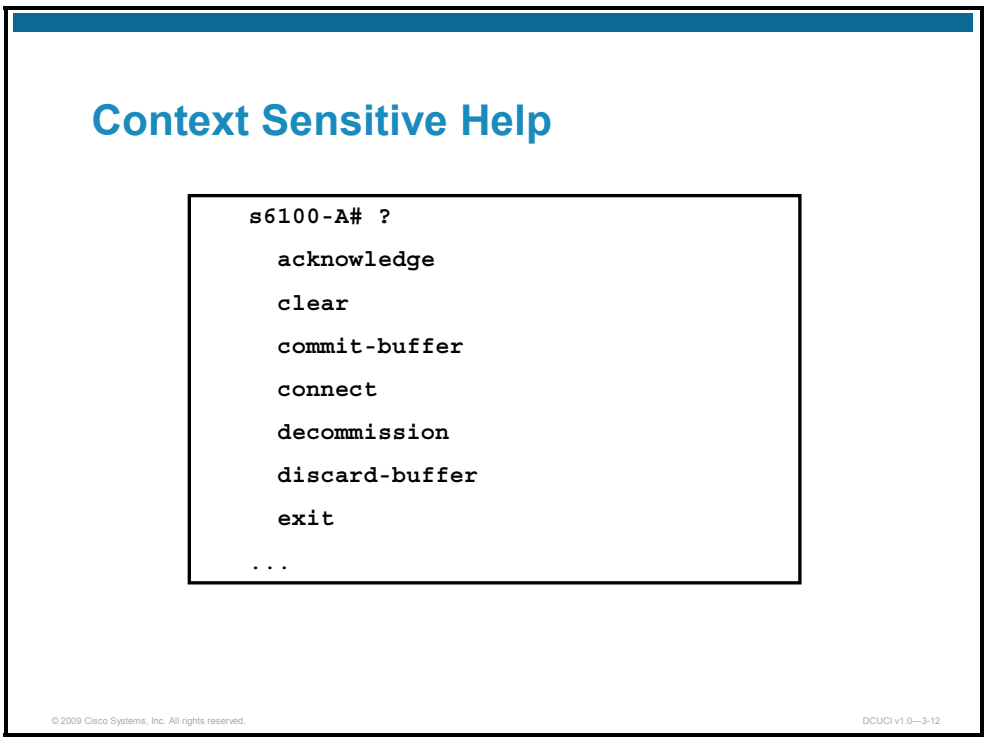

The question mark is familiar to most Cisco command-line interface (CLI) users. It shows all of the commands available within the present context. The figure shows some of the most common commands.

`UCS-6100-A# ?`

- acknowledge Acknowledge
- clear Reset functions
- commit-buffer Commit transaction buffer
- connect Connect to another CLI
- decommission Decommission managed objects
- discard-buffer Discard transaction buffer
- exit Exit from command interpreter
- recommission Recommission server resources
- remove Remove
- scope Changes the current mode
- set Set property values

https://t.me/learningnets

- zshow — Show running system information
- terminal — Set terminal line parameters
- top — Go to the top mode
- up — Go up one mode
- where — Show information about the current mode

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

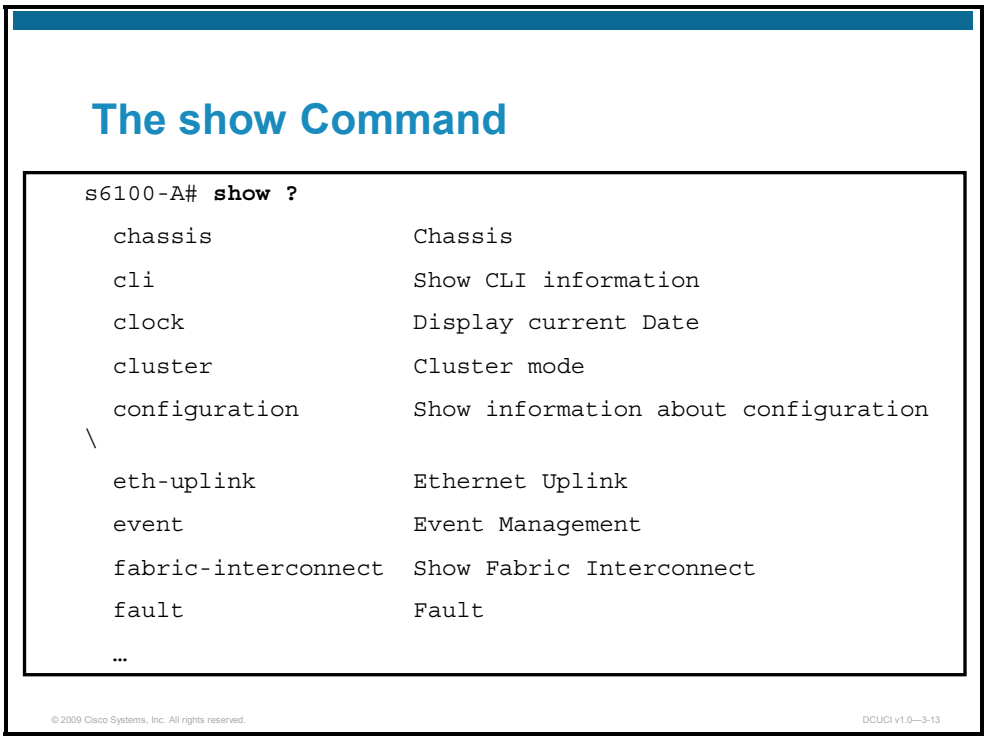

The **show** command, followed by a question mark, displays the attributes or parameters available for viewing.

This figure shows a few of the commands displayed using the context sensitive help for the **show** command.

https://t.me/learningnets

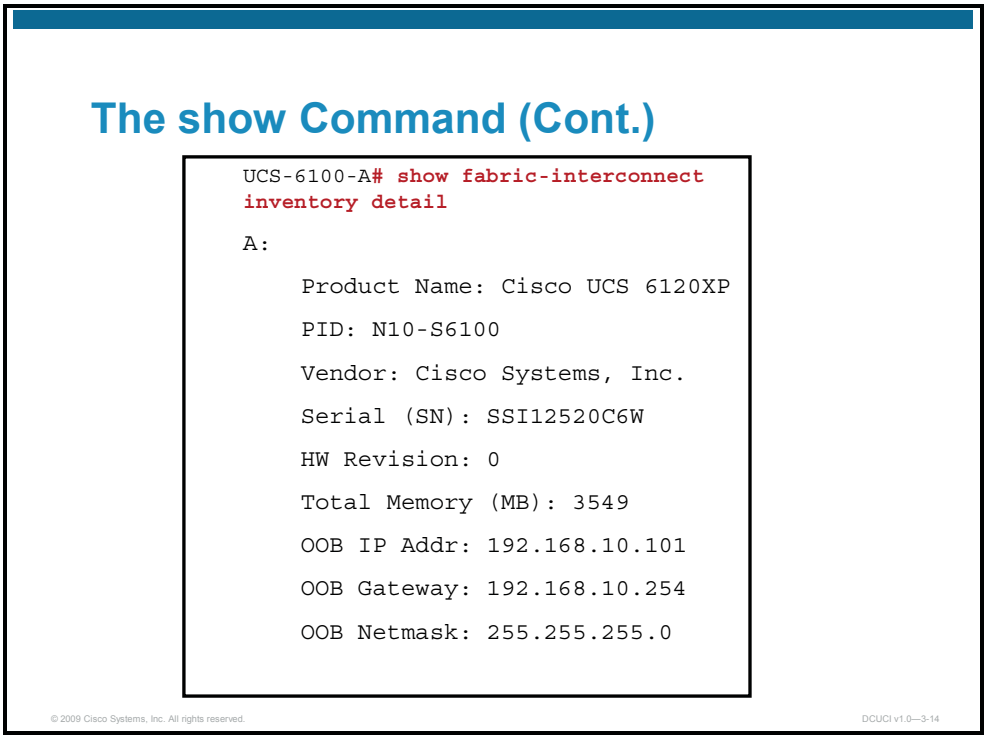

This example illustrates the **show** command with the arguments of *fabric-interconnect inventory detail*. The resulting display shows the inventory components of the switch. This is one example of the **show** command.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

The show Command (Cont.)

```
UCS-6100-A# show configuration
 recommission chassis "Cisco Systems
Inc" N20-C6508 FOX1307H0M8
scope org
enter server-qual all-chassis
    enter chassis 1 40
    exit
exit
enter service-profile TestP1 instance
    associate server 1/3
    enter boot-definition
        enter storage
            enter san-image primary
                set vhba fc0
…
```

© 2009 Cisco Systems, Inc. All rights reserved.

The **show configuration** command displays the commands to rebuild the configuration. Note that the **show configuration** command pulls its configuration information from the configuration database. This command, as well as other commands, can be run in scripted fashion using the **run-script** command.

https://t.me/learningnets

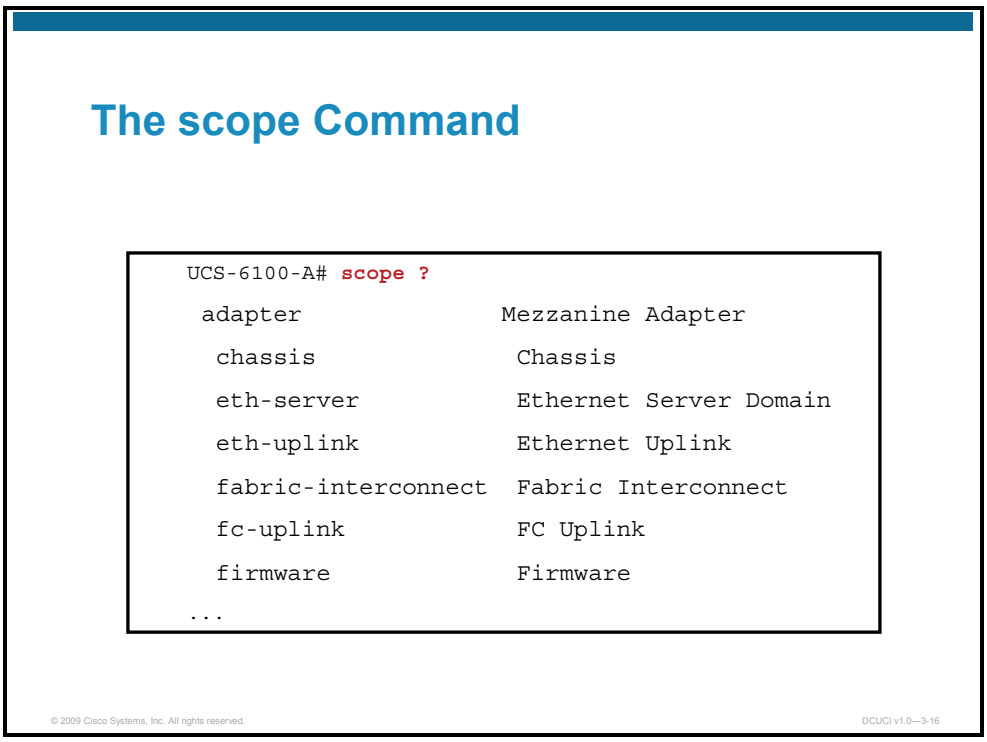

The **scope** command is used to change the CLI focus of the object being configured. This command allows the traversing of the hierarchical tree of objects and their child or parent objects. The **show** command can be used in conjunction with the **scope** command to immediately see the commands available in the new position of the tree.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

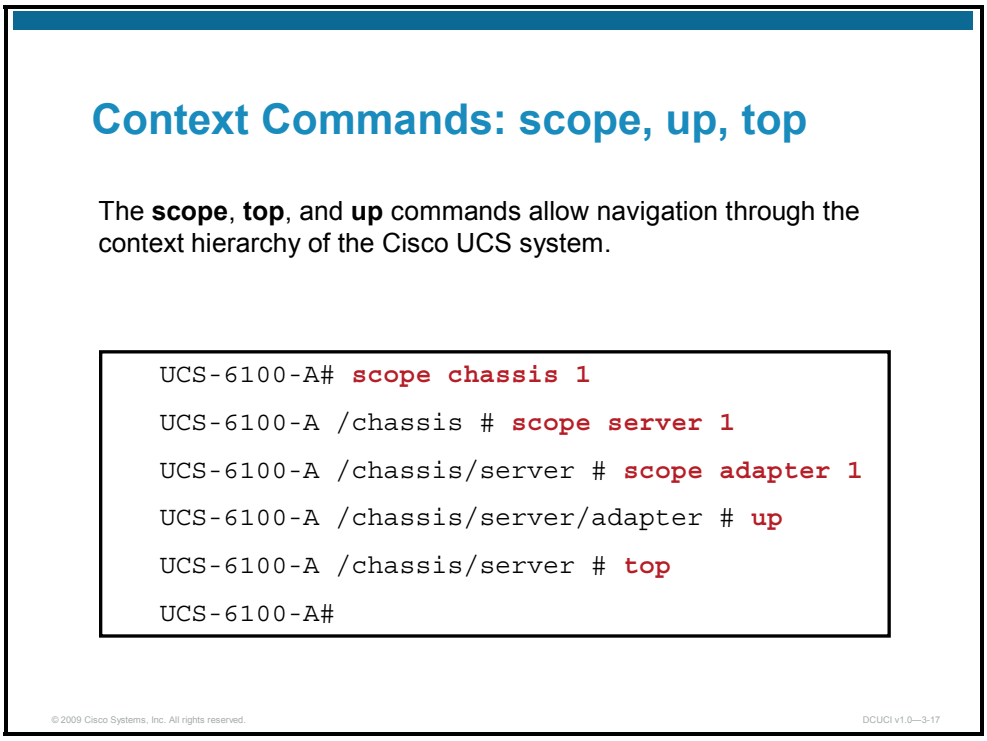

The figure shows the use of the **scope**, **up**, and **top** commands. Levels are traversed using the **scope** command. The **up** command is used to go up one level, and the **top** command is used to return to the root level.

https://t.me/learningnets

Management Commands: where

The **where** command allows you to identify your current scope.

```
UCS-6100-A /chassis/server # top
UCS-6100-A# scope chassis 1
UCS-6100-A /chassis # scope server 1
UCS-6100-A /chassis/server # scope adapter 1
UCS-6100-A /chassis/server/adapter # where
Mode:   /chassis/server/adapter
Mode Data:
        scope chassis 1
        scope server 1
        scope adapter 1
UCS-6100-A /chassis/server/adapter #
```

© 2009 Cisco Systems, Inc. All rights reserved.

The **where** command shows the current scope and all scopes associated with the current view. The **where** command also displays the mode (the exec mode in this case).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

The connect Commands

This topic describes the **connect** commands.

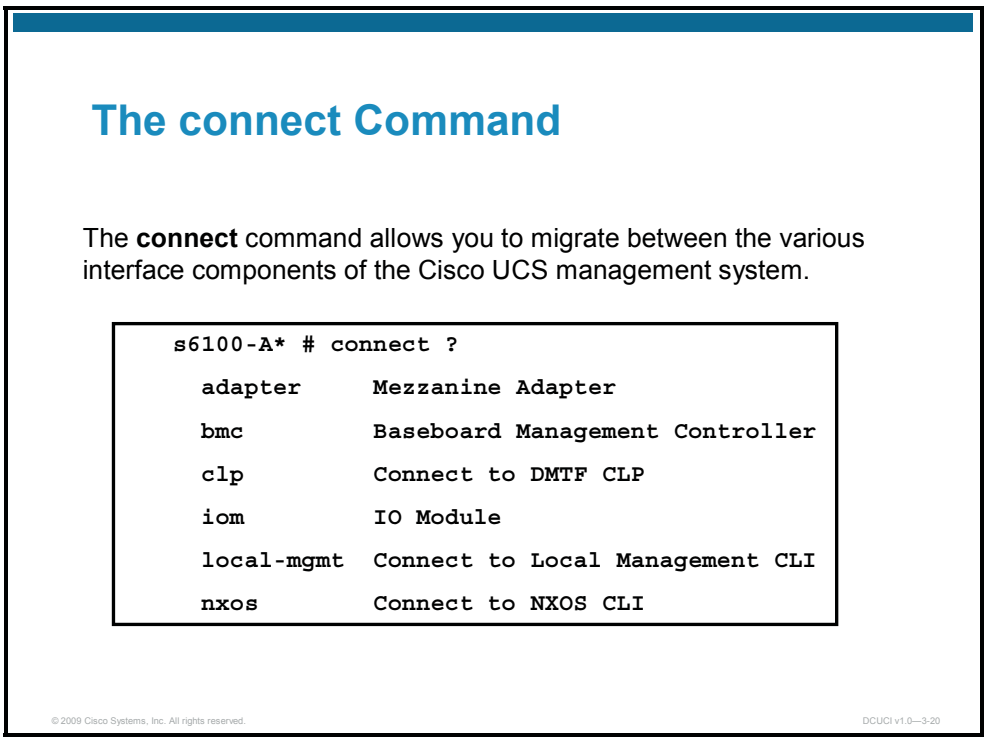

When connecting to the Cisco UCS system with Telnet or Secure Shell (SSH), you are placed in the default UCSM shell.

Enter **connect ?** to display connect options.

- connect adapter Connect to a mezzanine adapter
- connect bmc Connect to a blade BMC
- connect clp Connect to DMTF CLP
- connect iom Connect to a chassis IOM
- connect local-mgmt Connect to Local Management CLI
- connect nxos Connects to Cisco Nexus Operating System (Cisco NX-OS) CLI and offers underlying Nexus operating system commands

https://t.me/learningnets

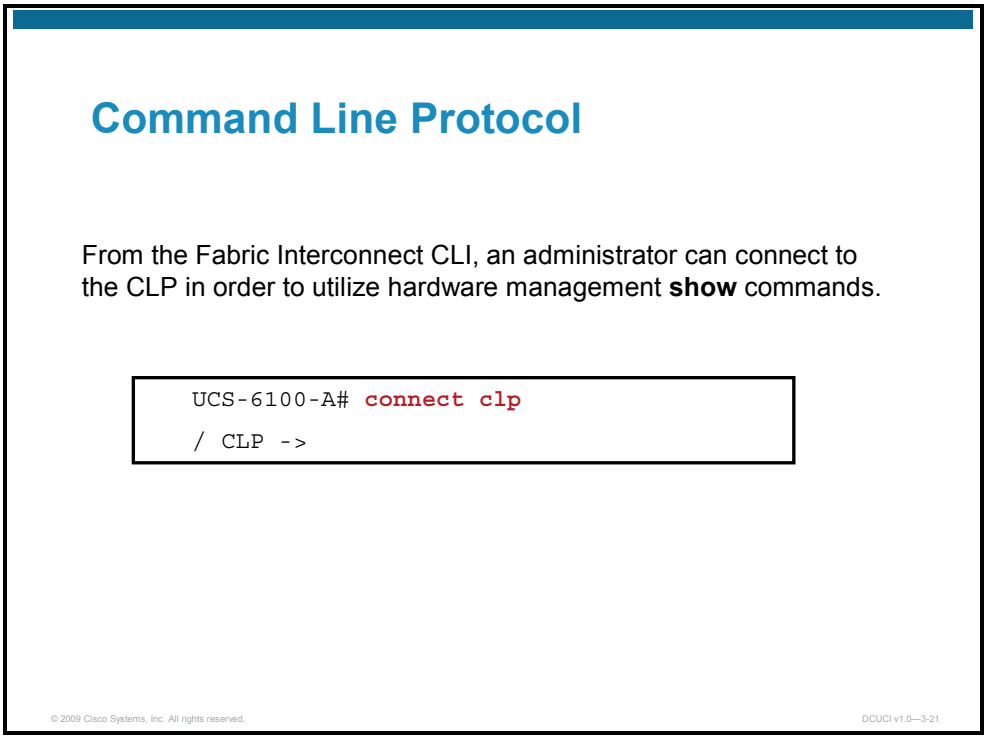

The Command Line Protocol (CLP) uses the smash CLP format. At present, it is mostly limited to **show** commands.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

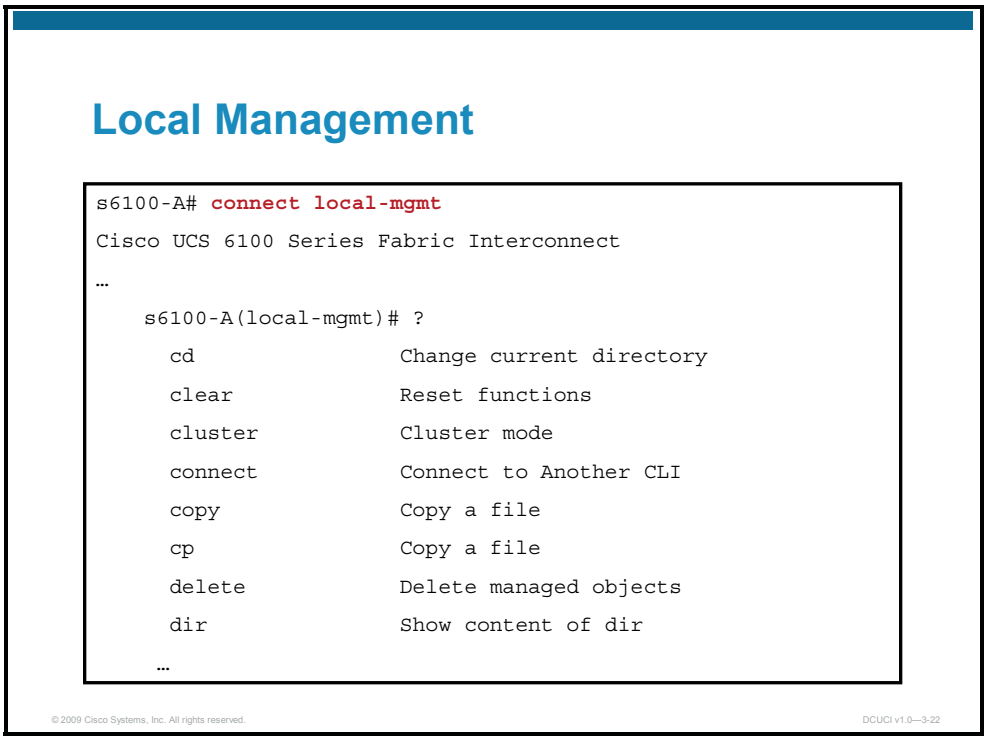

The figure shows the **connect local-mgt** command to attach to the local management CLI. The local management (local-mgmt) CLI contains the **reboot** and **cluster** management commands, as well as directory navigation and file manipulation commands.

https://t.me/learningnets

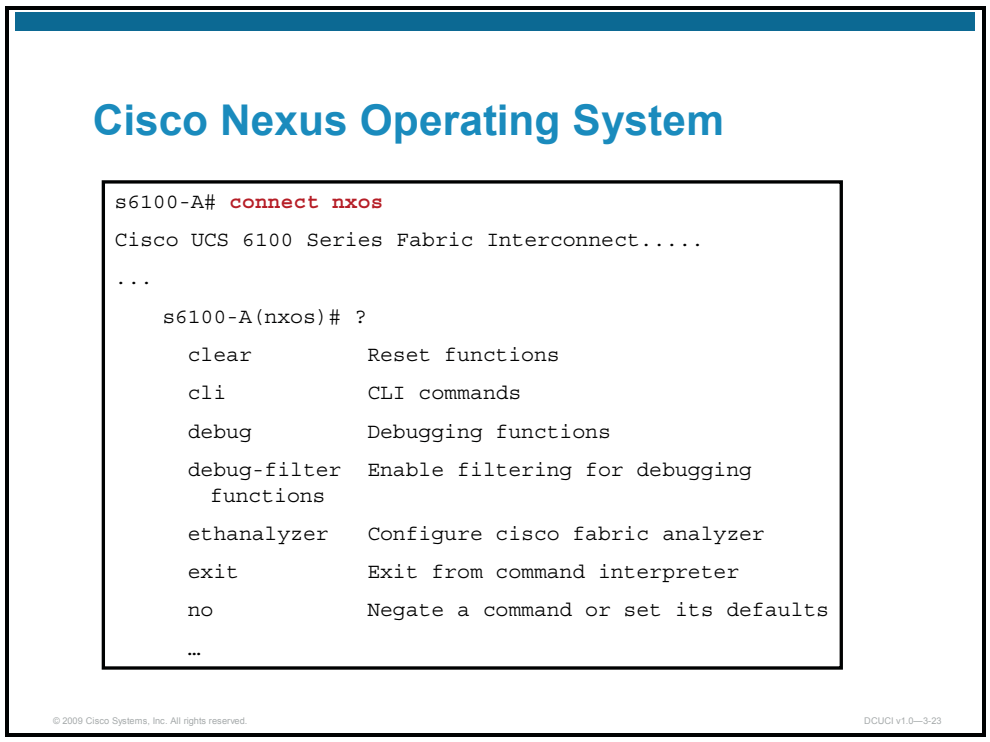

The figure shows a subset of commands available through the Cisco Nexus Operating System (Cisco NX-OS) CLI. An example of a command found in the Cisco NX-OS CLI is **ethanalyzer**, a Cisco fabric analyzer. Use **show ?** for an additional list of commands.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

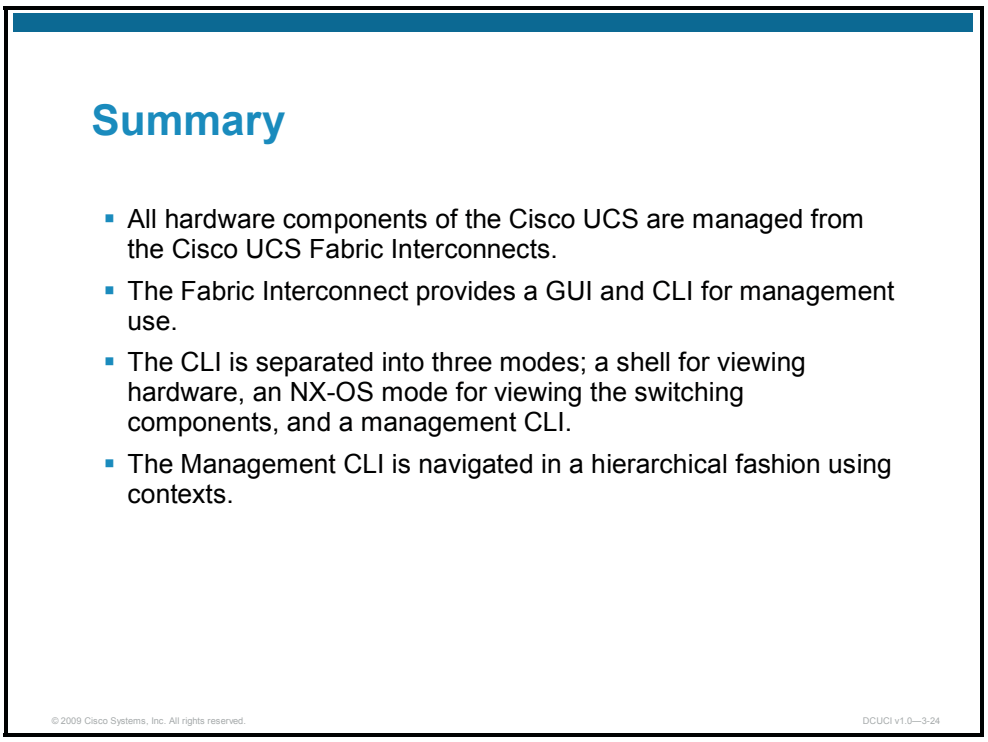

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Describe Cisco UCS Startup and Shutdown Procedures

Overview

In this lesson you will learn how to perform Cisco Unified Computing System (Cisco

.

Objectives

Upon completing this lesson, you will be able to describe Cisco Unified Computing System startup and shutdown procedures. This includes being able to meet these objectives:

- Describe the Cisco UCS 6100 startup and shutdown procedures

https://t.me/learningnets

Cisco UCS 6100 Startup and Shutdown

This topic describes the Cisco UCS 6100 startup and shutdown procedures.

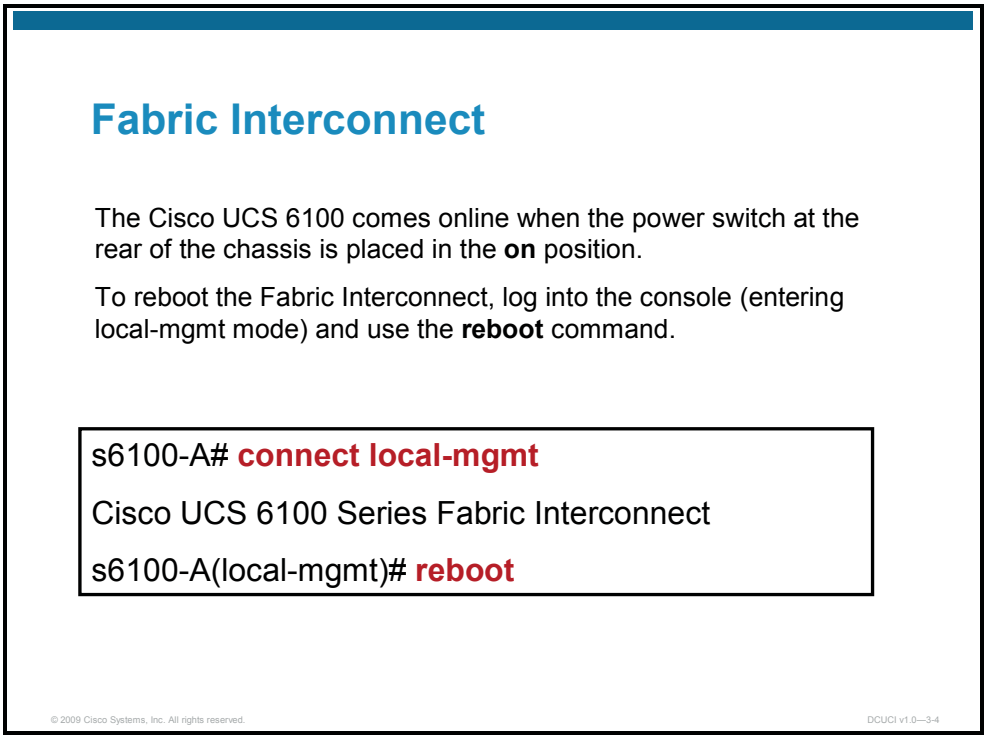

This example describes power on and reset procedures for the Cisco UCS 6100 Fabric Interconnect.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Booting a Compute Node

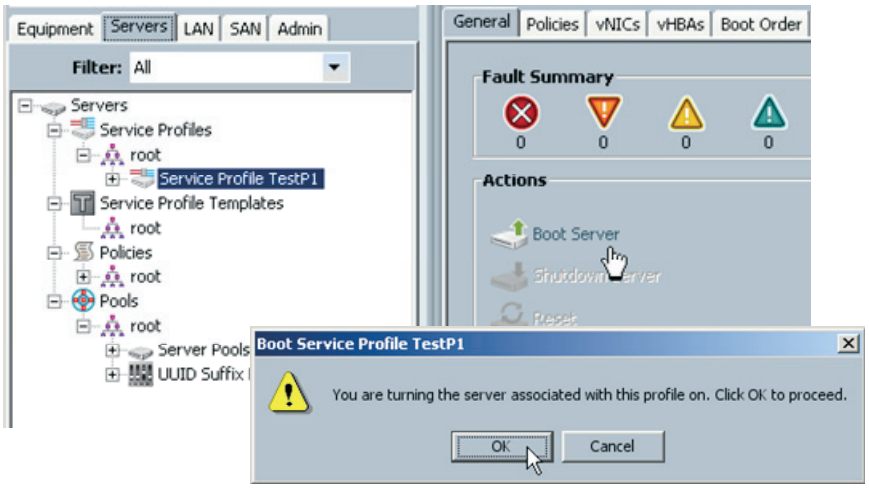

© 2009 Cisco Systems, Inc. All rights reserved.

To boot a computer node:

Step 1 Select the **Servers** tab in the navigation pane, and highlight the **Service Profile** of the node to be started.

Step 2 Select the General tab of the content pane, and click the **Boot Server** link.

https://t.me/learningnets

Shutting Down a Compute Node

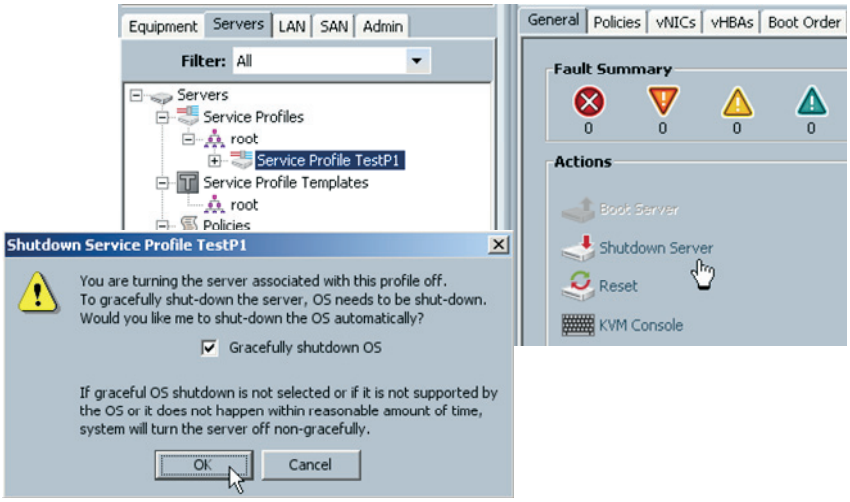

© 2009 Cisco Systems, Inc. All rights reserved.

DCUCI v1.0—3-6

To shut down a computer node:

Step 1 In the Servers tab, highlight the **Service Profile** of the node to be stopped.

Step 2 Ensure that you have gracefully shut down any running operating systems.

Step 3 In the General tab of the content pane, click the **Shutdown Server** link.

https://t.me/learningnets

Compute Node Hardware Restart

Each blade has a reset button on the front of the device. This reset button should only be used when a software reset through the Cisco UCS Manager is not possible.

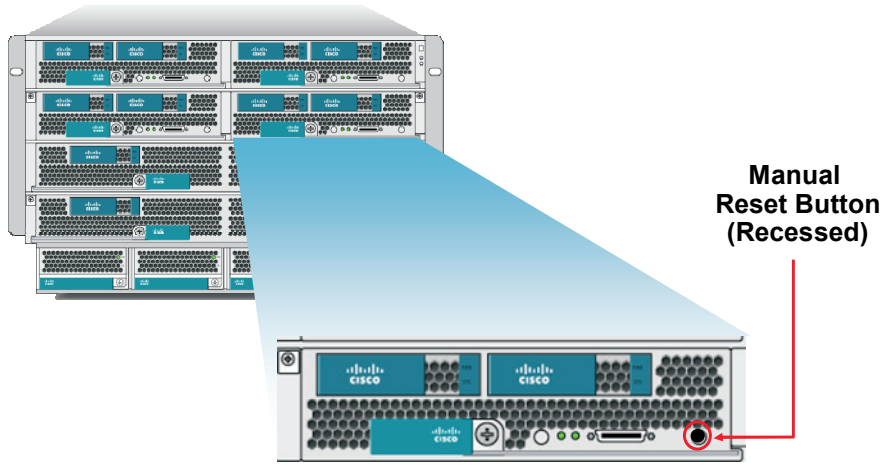

© 2009 Cisco Systems, Inc. All rights reserved.

In Cisco UCS, the compute node front panel power button is de-bounced (routed) though the baseboard management controller (BMC), which allows power button lock-out. Since the BMC has a direct connection to the chipset power line, it can initiate both graceful restarts (pressing the power button for less than four seconds) and forced restarts (holding the power button down for more than four seconds).

https://t.me/learningnets

I/O Module Power Options

The I/O Modules within the blade chassis are not managed devices and therefore have no independent power options. These devices act as an extension of the Cisco UCS 6100 Fabric Interconnect.

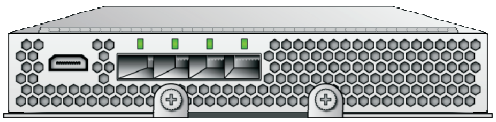

These devices are powered on with the Cisco UCS Chassis when power cables are plugged into the four available power ports.

© 2009 Cisco Systems, Inc. All rights reserved.

Within a Cisco UCS environment, the I/O Modules act as extensions to the Cisco UCS 6100 Fabric Interconnect. These devices have no direct power management.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

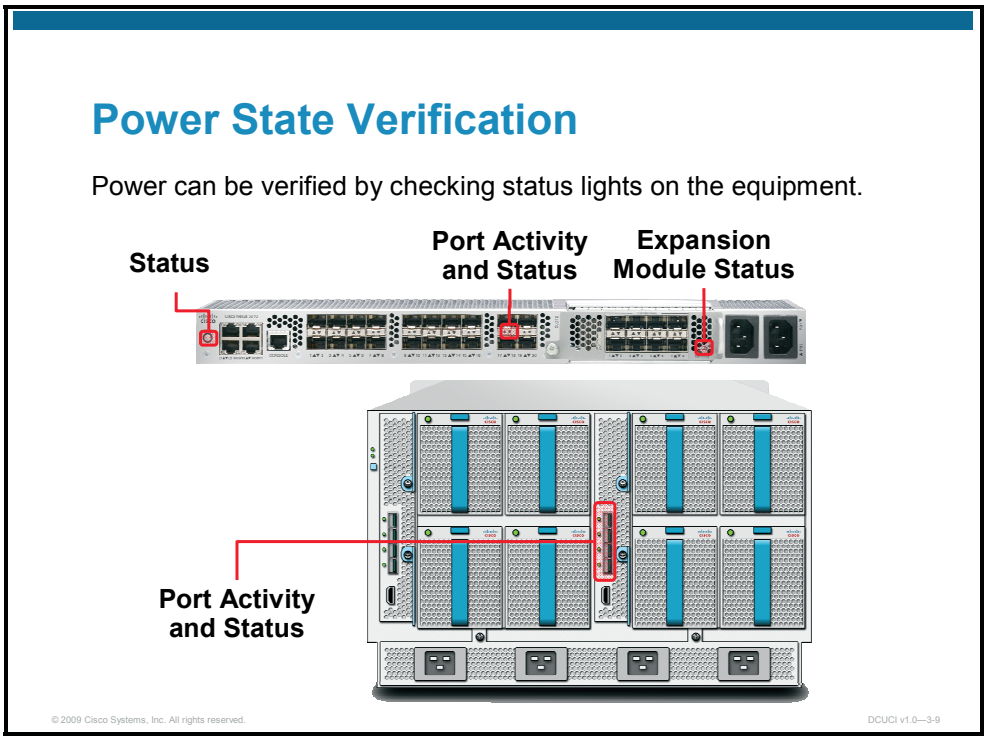

This figure shows the rear view of the Cisco UCS Fabric Interconnect and chassis. Power and status indicators are identified.

https://t.me/learningnets

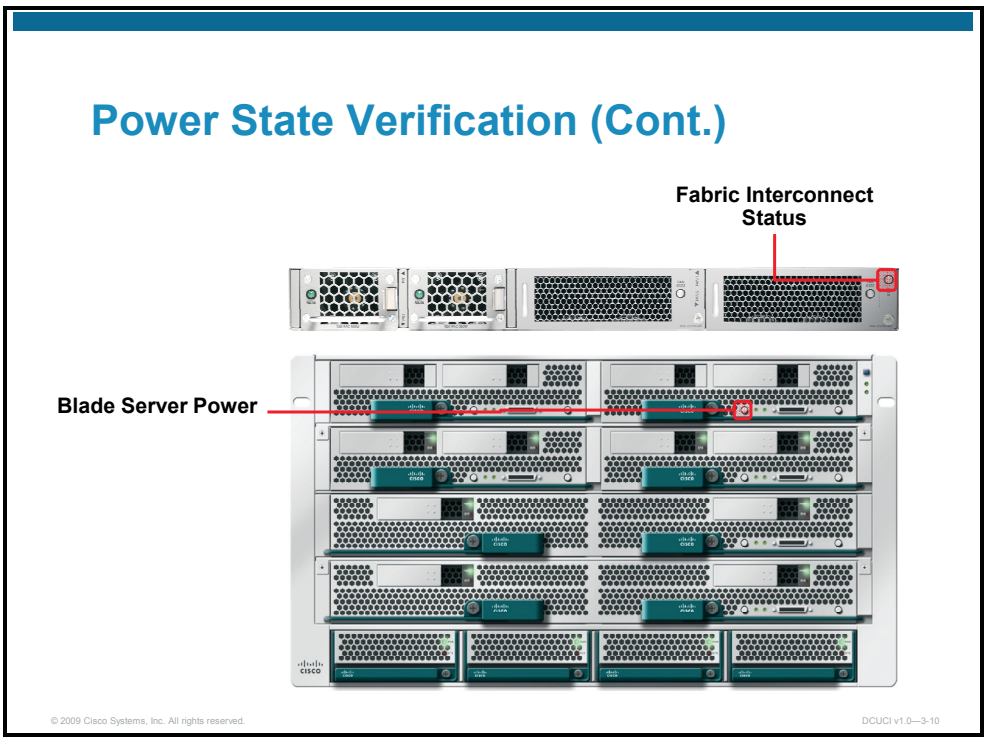

This figure identifies the power status indicators on the front of the Cisco UCS 6100 and chassis.

https://t.me/learningnets

Summary

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

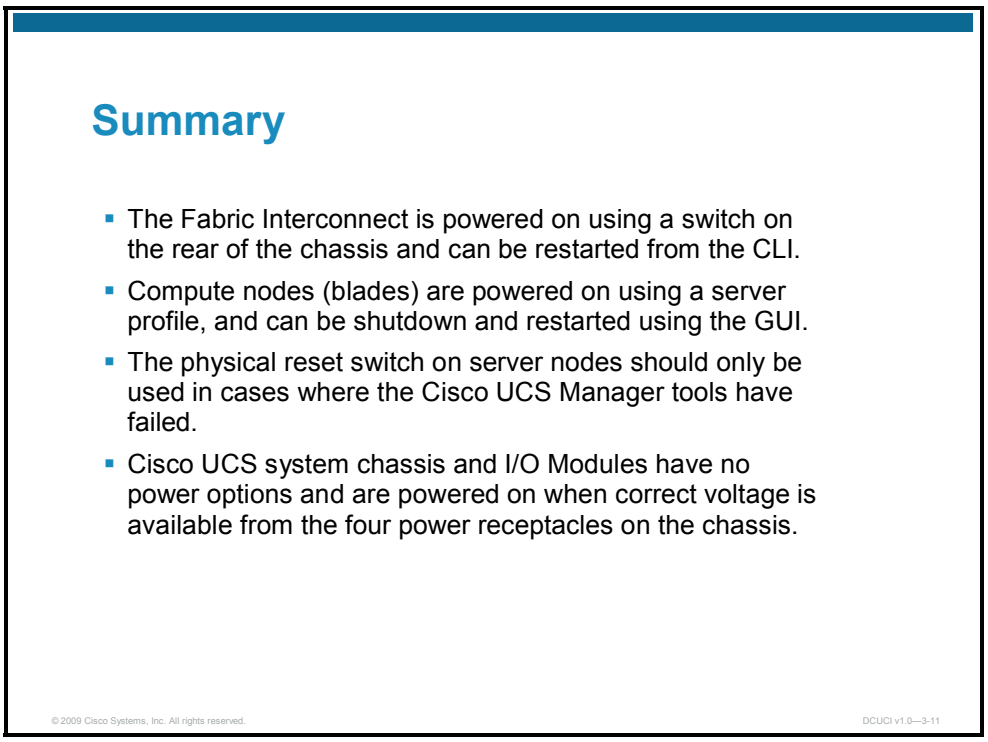

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

Upgrading Software in the Cisco Unified Computing System

Overview

In this lesson you will learn how to upgrade software components in the Cisco Unified Computing System (Cisco UCS).

Objectives

Upon completing this lesson, you will be able to upgrading software in the Cisco Unified Computing System. This includes being able to meet these objectives:

- Describe the Cisco UCS software upgrade process
- Describe how to upgrade Cisco UCS Fabric Interconnect software
- Describe how to upgrade CMC software
- Describe how to upgrade BMC software
- Describe how to upgrade adapter software

https://t.me/learningnets

Cisco UCS Software Upgrade Overview

This topic describes the Cisco UCS software upgrade process.

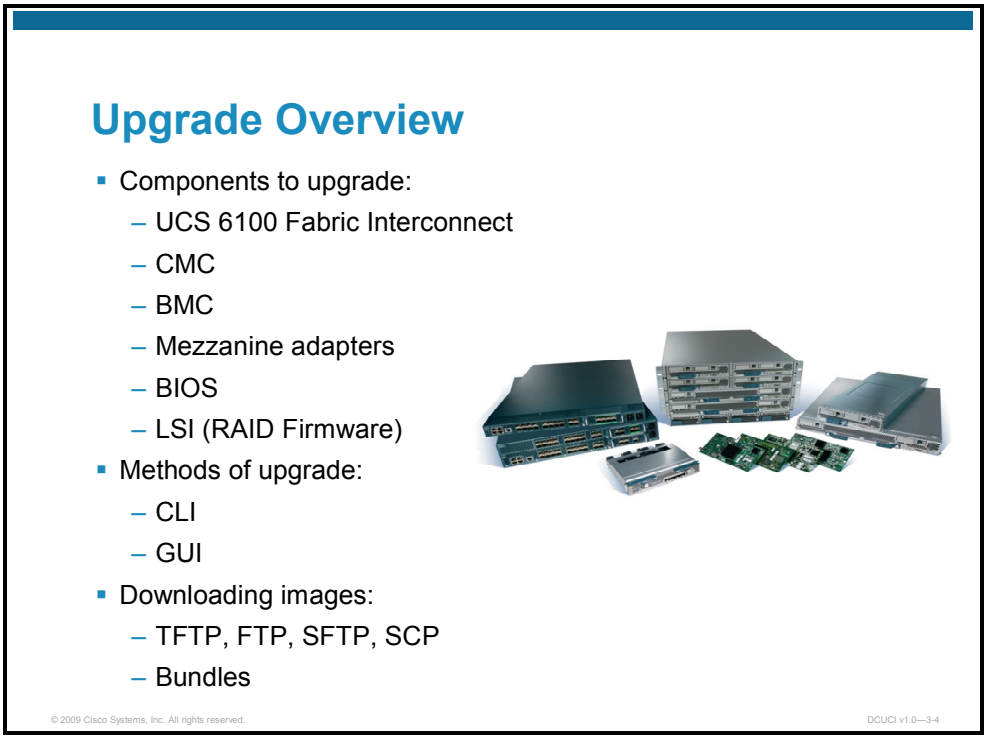

Components to Upgrade

The Cisco UCS components that can be upgraded include the following:

- Fabric Interconnect software:
  - Kernel – basic kernel
  - System – Cisco Nexus Operating System (Cisco NX-OS)
  - Management – Cisco UCS Manager
  - CMC software – Chassis management controller (CMC); integrated into IO Module
- BMC software – Baseboard management controller (BMC); integrated into blade servers
- Mezzanine adapter firmware – add-on cards attached to blade servers
- BIOS firmware – blade server BIOS code
- LSI firmware – blade server RAID controller

Methods of Upgrade

There are two methods for upgrading the components in a Cisco UCS system:

- Using the Cisco UCS Manager command-line interface (CLI)
- Using the Cisco UCS Manager GUI

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

One exception is the CMC, which may be automatically upgraded by the Fabric Interconnect. When upgrading the entire Cisco UCS system, you should upgrade the “leaves” of the system first (in other words, the mezzanine cards first, followed by the BMC, the BIOS, the CMC, and the Fabric Interconnect last).

Downloading Images

You can use the CLI or GUI to download and install images. The Cisco UCS Manager can use TFTP, FTP, SFTP, and SCP protocols for downloading images from file servers on the network. You can either download and install images for individual component types (for example, Converged Network Adapters), or you can download a single file that bundles all component images.

Downloading individual component images is useful for debugging purposes. You can store multiple image versions on your Fabric Interconnect, as long as there is enough space in the /bootflash file system. Currently, the Cisco UCS 6100 Series Fabric Interconnect ships with either 4-GB or 16-GB flash space that is dedicated to /bootflash.

https://t.me/learningnets

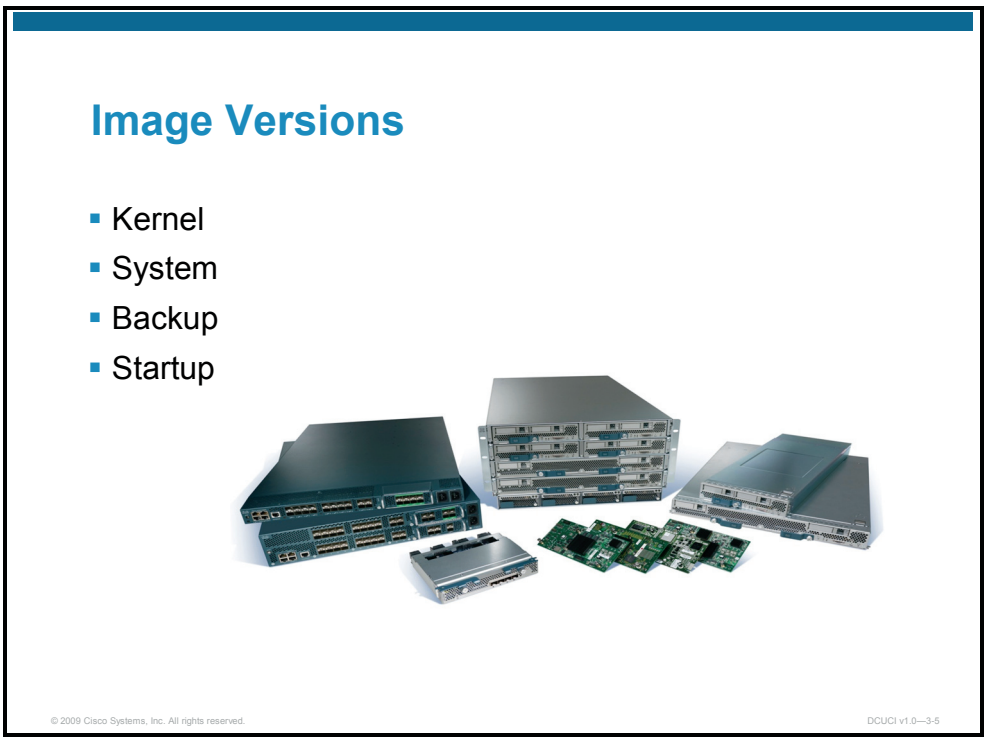

Each Cisco UCS component can store two images (each with a different version). This can be useful for staging an upgrade, where the first stage involves installing the new version (which can be performed without service interruption) and the second stage involves activating the new version (which requires a component reboot).

Kernel

The kernel is the underlying operating system, upon which the Cisco NX-OS and Cisco UCS Manager (UCSM) packages run.

System

The system version is the name of the image that is currently booted and running now.

Backup

The backup version is the name of the image that is not currently booted and running now.

Startup

The startup version is the name of the image that the bootloader will load the next time this component boots.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

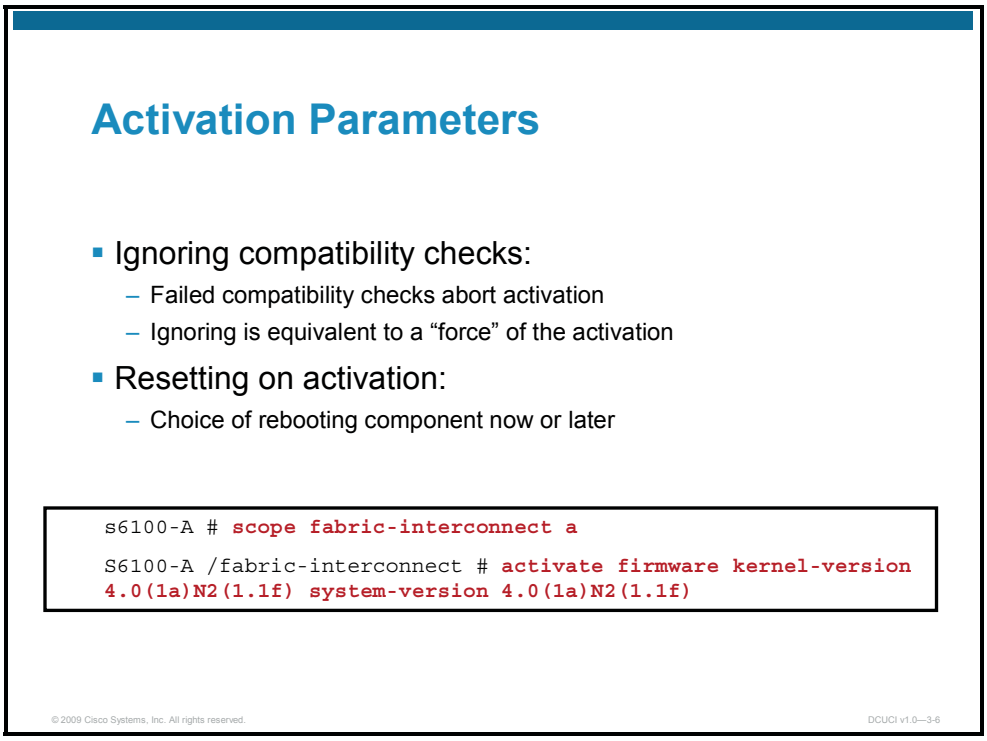

Ignoring Compatibility Checks

When performing an upgrade of a Cisco UCS component, the Cisco UCS Manager software will attempt to determine if the new version is compatible with the version of the Cisco UCS Manager that is currently running. If the Cisco UCS Manager determines the new version is incompatible, the upgrade aborts. You can instruct the upgrade process to ignore the result of this compatibility check in order to force the upgrade.

Resetting on Activation

When performing an upgrade of a Cisco UCS component, there are two images stored for each component; the startup version and the backup version. You can optionally stage the activation to a later time when a reboot can be performed to apply the changes, or you can reboot the component immediately to activate the upgrade image. Verify the components that require reboot prior to making changes to those components.

https://t.me/learningnets

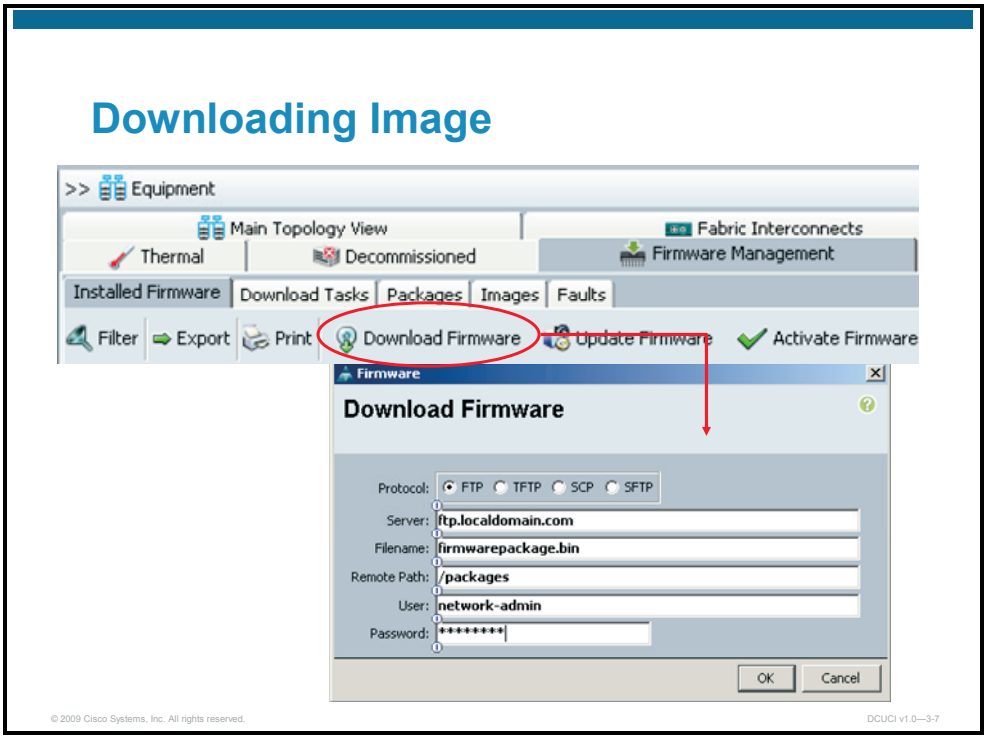

To download an image, select the **Equipment** tab from the navigation pane. Then, select the **Firmware Management** tab from the content pane. Click the **Download Firmware** link to display a download form. Complete the following fields in the form: Protocol, Server, File Name, Remote Path, User, and Password.

Depending on the protocol and current user, the remote path to use will differ; when using FTP or TFTP, use the path that is relative to the directory configured on the FTP/TFTP server (for example, /tftpboot or /pub). FTP, SFTP, or SCP all require a user name and password. Ensure that the credentials entered in this form are valid. Depending on the selected protocol, some of the fields in the download form may not display.

During the download, the Cisco UCS Manager places a copy of the image in the /bootflash/tmp directory. After the download is complete, the Cisco UCS Manager places the image under the /bootflash/installable directory and create links under the /bootflash/distributable directory. The Cisco UCS Manager interface creates a *.hdr file to accompany each installable image. Without the *.hdr file, images can not be installed through the Cisco UCS Manager interface.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

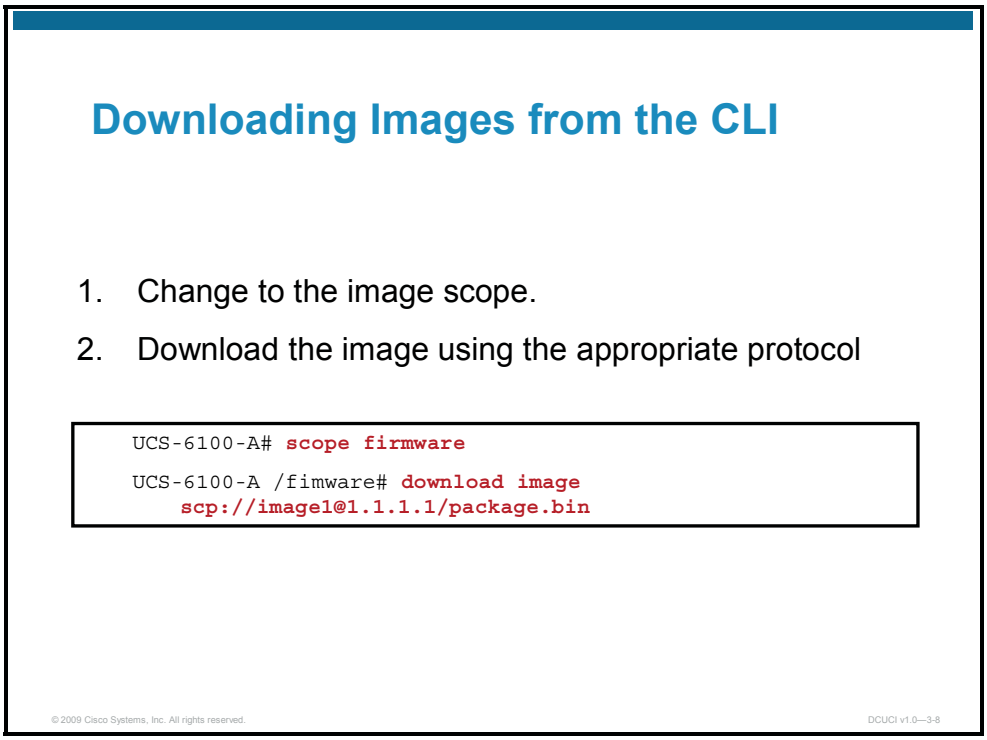

On the File Server

When using the SCP protocol to download images, you must use the absolute path from the root directory to the image named in the **download** command. Use the **cd** command to change to the image directory and run the **pwd** and **ls** commands to obtain the file information.

From the Cisco UCS Manager CLI

Ensure proper syntax when using the **download** command. It is hard to distinguish download failures between incorrectly entered command syntax, path and file names, and user credentials.

Once the download completes, use **show package** to verify that the new images can be installed. Images may now be installed from within the bundle to the Cisco UCS components. Bundles are approximately 275 MB and download duration may vary significantly.

If the download fails, the download object must be deleted before retrying the download process.

https://t.me/learningnets

Verifying Installable Images

Use the **show package** command to verify available images.

```
UCS-6100-A /firmware# show package expand
Package ucs-k9-bundle.1.0.1e.bin:
    Images:
        ucs-2100.1.0.1e.bin
        ucs-6100-k9-kickstart.4.0.1a.N2.1.1e.bin
        ucs-6100-k9-system.4.0.1a.N2.1.1e.bin
        ucs-b200-m1-bios.S5500.86B.01.00.0036-
191.061320091126.bin
        ucs-b200-m1-k9-bmc.1.0.1e.bin
    …
```

© 2009 Cisco Systems, Inc. All rights reserved.

Cisco UCS image bundles include the following:

- Kernel, Management, and System images for the switch
- BMC image
- CMC image
- Two images for the converged network adapter (Emulex and Qlogic)
- Virtual Interface Card image
- BIOS image
- LSI image

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Displaying Installable Images

Installable images and packages can be viewed in the Firmware Management tab of the equipment content pane.

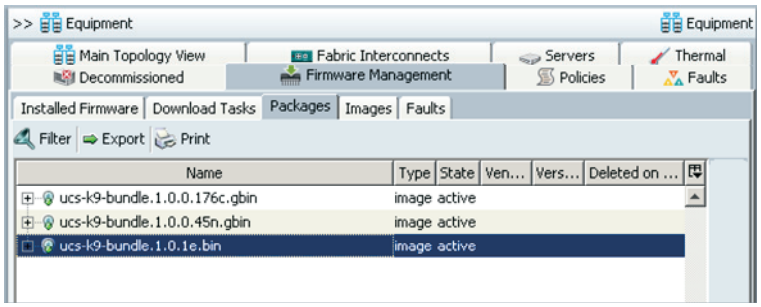

© 2009 Cisco Systems, Inc. All rights reserved.

You can also display installable images in the GUI. From the navigation pane, select the **Equipment** tab. Then, select the **Firmware Management** tab in the content pane.

https://t.me/learningnets

Upgrading Cisco UCS Fabric Interconnect Software

This topic describes how to upgrade Cisco UCS Fabric Interconnect software.

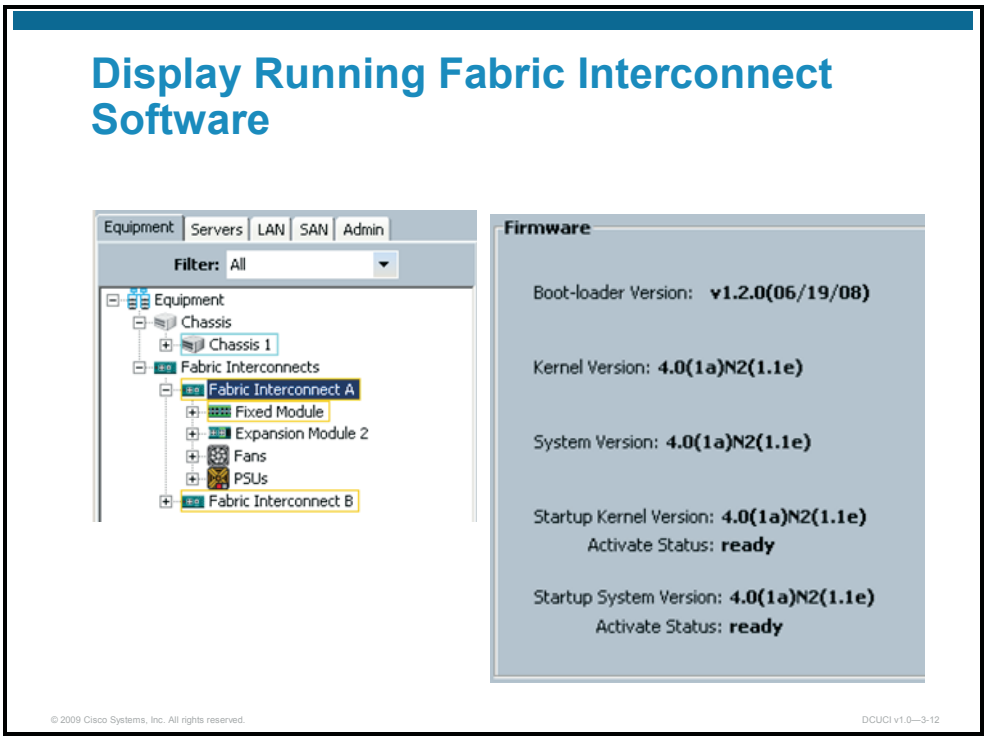

When displaying the currently running software, there are two different images that the component can load at startup (running version, and startup version), as shown in the figure.

In the navigation pane, select the **Equipment** tab. Then, select the **Fabric Interconnect** object.

For the switch, there are three different images:

- The Kernel Version (the kickstart image)
- The System Version (the Cisco NX-OS image)

Note that the management image version is not displayed in the GUI. It is installed along with the System Version.

The kernel, or kickstart, is the underlying Linux-based operating system upon which the System (NX-OS) is run.

In this example, the Kernel and System versions listed indicate which version is currently running. The Startup versions indicate which version would be loaded were the Fabric Interconnects to be rebooted. These might be different if an upgrade is in process or was loaded by not yet activated.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Display Running Fabric Interconnect Software

```
UCS-6100-A /image # scope fabric-interconnect a
UCS-6100-A /fabric-interconnect # show firmware
Fabric Interconnect A:
    Running-Kern-Vers: 4.0(1a)N2(1.1e)
    Running-Sys-Vers: 4.0(1a)N2(1.1e)
    Startup-Kern-Vers: 4.0(1a)N2(1.1e)
    Startup-Sys-Vers: 4.0(1a)N2(1.1e)
    Act-Kern-Status: Ready
    Act-Sys-Status: Ready
    Bootloader-Vers:   v1.2.0(06/19/08)
```

© 2009 Cisco Systems, Inc. All rights reserved.

You can also display the currently running software from the CLI.

https://t.me/learningnets

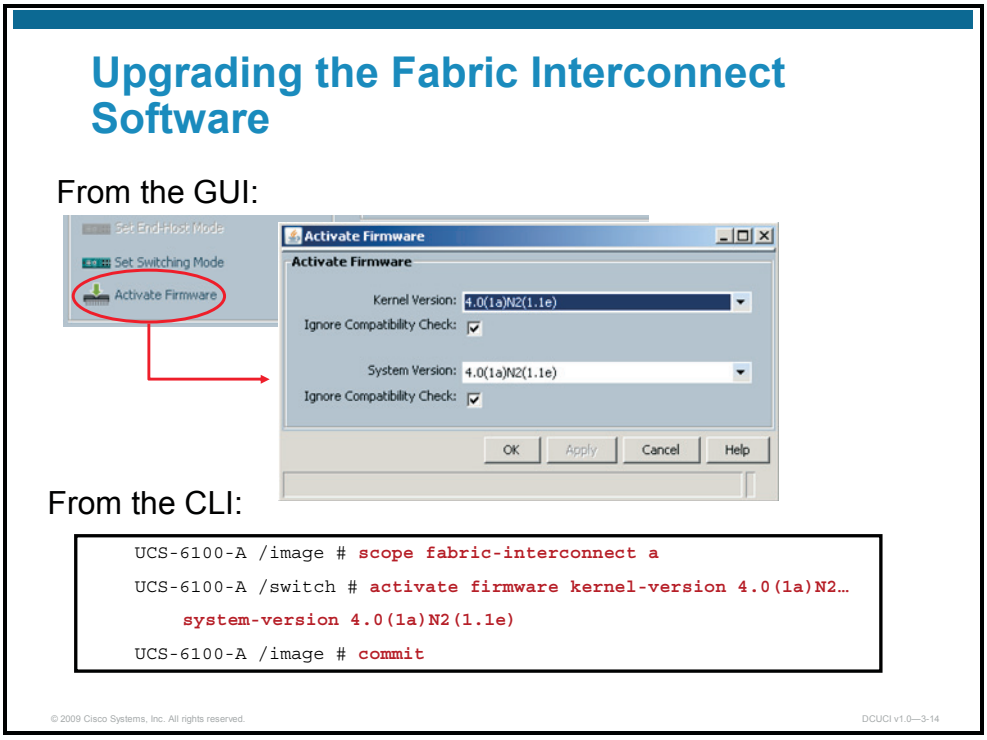

When activating new firmware on the Cisco UCS Fabric Interconnect using the GUI or CLI, the switch is configured to automatically boot that image on power-on or reboot.

Upgrading the Fabric Interconnect Software using the Cisco UCS Manager GUI

The currently running Kernel and System versions are the old versions, and the startup Kernel and System versions are the new versions. The switch must be rebooted for the changes (upgrade) to take effect.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Upgrading CMC Software

This topic describes how to upgrade CMC software.

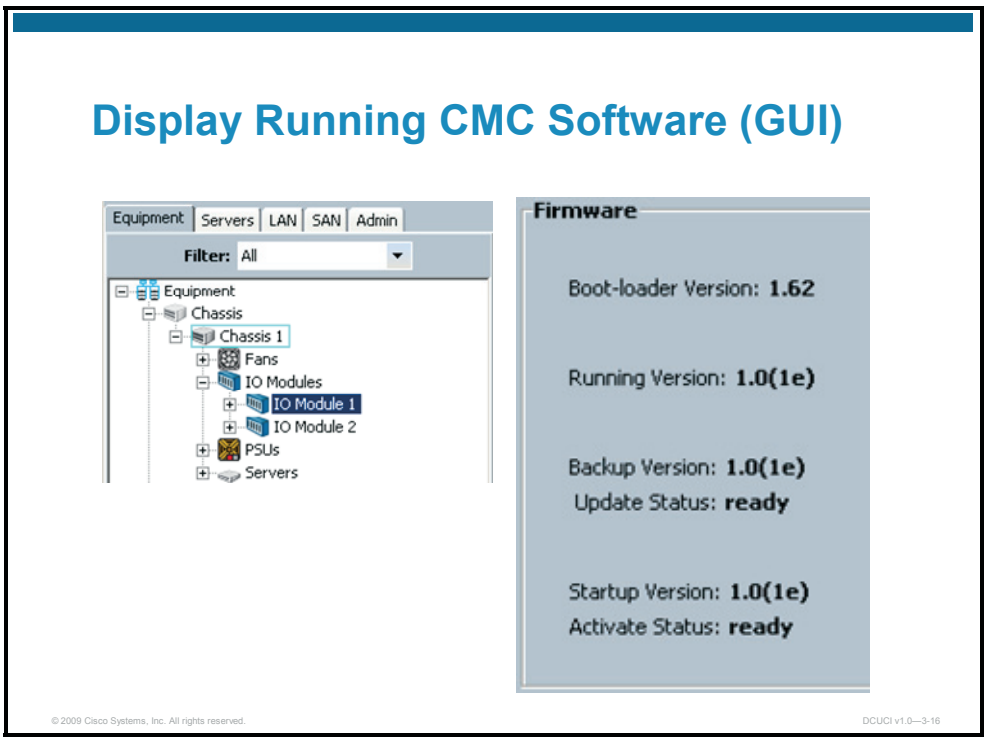

To view the currently running CMC software, select the **Equipment** tab from the navigation pane. Then, select the **Chassis** and **IO Module** objects.

The resulting Running Version field identifies the currently running version. The Update Status field identifies the progress of the upgrade (if an upgrade is in progress.) The upgrade is not complete until the new image has been loaded on the CMC. The user can optionally choose to activate the updated version, and later toggle back to the original version, as both are stored on the CMC. The activation is complete when a pointer has been modified to indicate the selected image.

https://t.me/learningnets

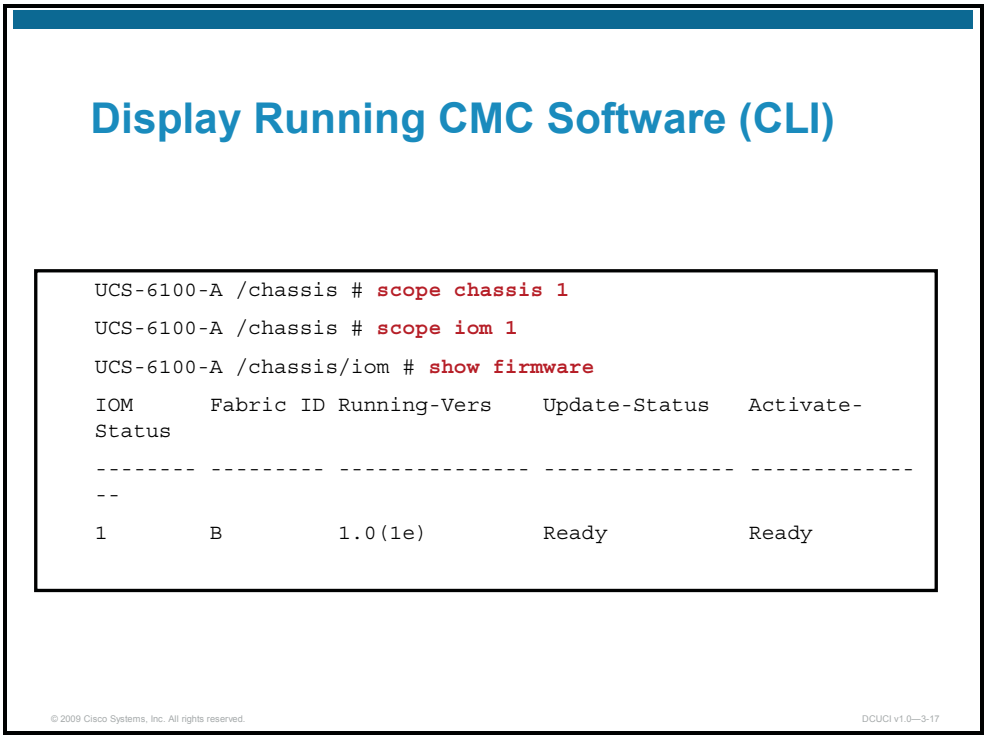

You can also use the CLI to display the software that is currently running on the CMC, as shown in the figure.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Upgrading the CMC Software

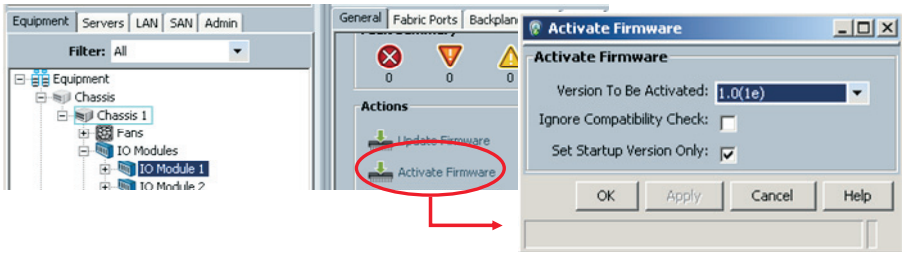


```
# scope chassis 1
/chassis # scope iom 1
/chassis/iom # update firmware 1.0(1e)
/chassis/iom # activate firmware 1.0(1e) -resetonactivate
```

© 2009 Cisco Systems, Inc. All rights reserved.

To upgrade the CMC software, select the **Equipment** tab from the navigation pane. Then, select the **Chassis** and **Fabric Extender** objects. Click the **Update Firmware** icon.

https://t.me/learningnets

Upgrading BMC Software

This topic describes how to upgrade BMC software.

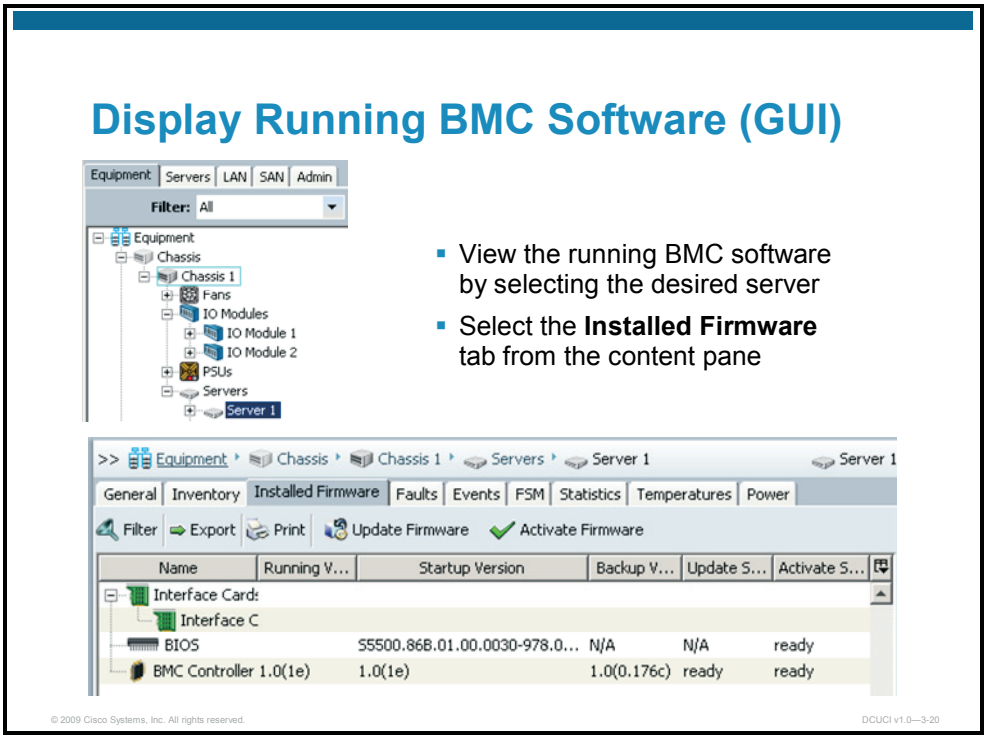

To view the currently running BMC firmware, select the **Equipment** tab from the navigation pane. Then, select the **Chassis** and **Server** objects. From the content pane, select the **Installed Firmware** tab.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Display Running BMC Software (CLI)

```
UCS-6100-A# scope server 1/1
UCS-6100-A# /chassis/server # scope bmc
UCS-6100-A# /chassis/server/bmc # show firmware detail
BMC:
    Running-Vers: 1.0(1e)
    Update-Status: Ready
    Activate-Status: Ready
    Startup-Vers: 1.0(1e)
    Backup-Vers: 1.0(0.176c)
    Bootloader-Vers: uboot1.0.1
```

© 2009 Cisco Systems, Inc. All rights reserved.

You can also view the currently running BMC firmware from the CLI, by using the **show firmware** command.

https://t.me/learningnets

Upgrading the BMC Software (GUI)

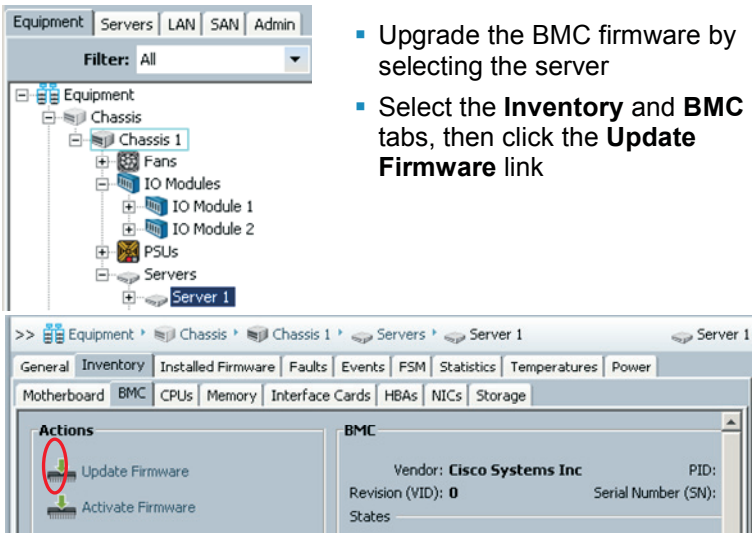

- Upgrade the BMC firmware by selecting the server
- Select the **Inventory** and **BMC** tabs, then click the **Update Firmware** link

© 2009 Cisco Systems, Inc. All rights reserved.

To view the currently running BMC firmware, select the **Equipment** tab from the navigation pane. Then, select the **Chassis** and **Server** objects. Select the **Inventory** and **BMC** tabs, and then click the **Update Firmware** icon.

You can force the activation of an image by ignoring the result of the compatibility check during the activation process. Resetting on activation amounts to a reboot of the BMC after the activation is complete.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Upgrading the BMC Software (CLI)

```
UCS-6100-A# scope server 1/1
UCS-6100-A# /chassis/server # scope bmc
UCS-6100-A# /chassis/server/bmc # update firmware …
```

© 2009 Cisco Systems, Inc. All rights reserved.

From the CLI, use the **update firmware** command to upgrade the BMC code.

https://t.me/learningnets

Upgrading Adapter Software

This topic describes how to upgrade adapter software.

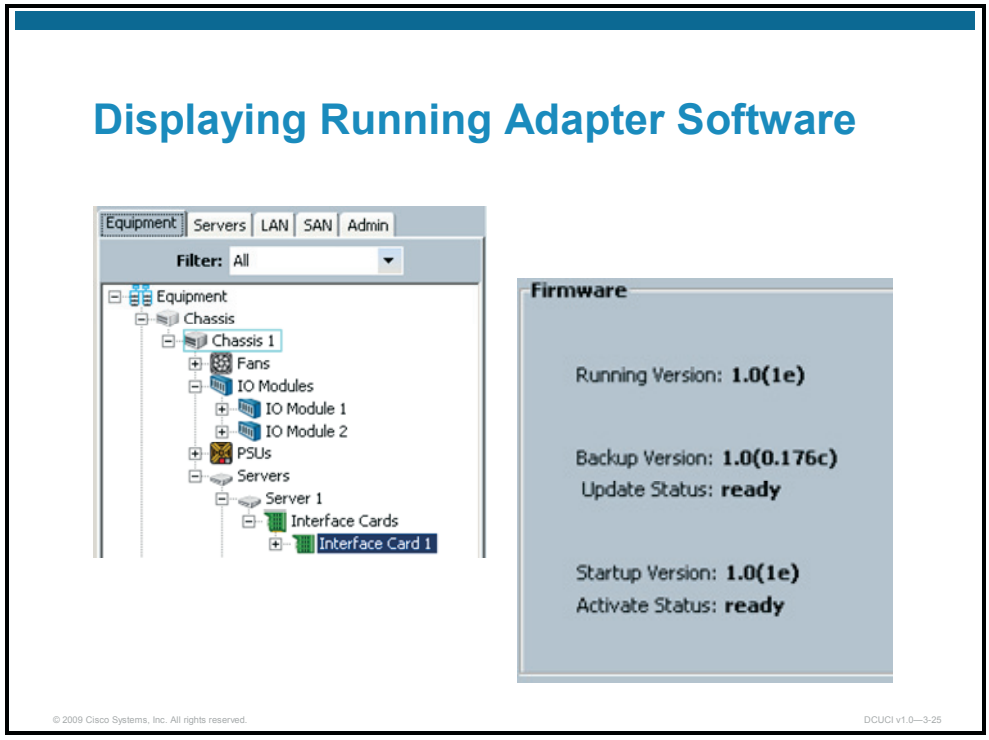

To view the currently running adapter software, select the **Equipment** tab from the navigation pane. Then, select the **Chassis**, **Server**, and **Interface Card** objects.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Display Running Adapter Software (CLI)

```
UCS-6100-A /chassis/server # scope server 1/3
UCS-6100-A /chassis/server # show firmware
Server 1/3:
    BMC:
        Running-Vers: 1.0(1e)
        Update-Status: Ready
        Activate-Status: Ready
    Adapter 1:
        Running-Vers: 1.0(1e)
        Update-Status: Ready
        Activate-Status: Ready
        …
```

© 2009 Cisco Systems, Inc. All rights reserved.

You can also view the currently running adapter software from the CLI. Here, the **scope** command is used to change the current CLI command scope to **Chassis 1**, **Server 3**.

The syntax is:

scope chassis/server

https://t.me/learningnets

Upgrading the Adapter Software

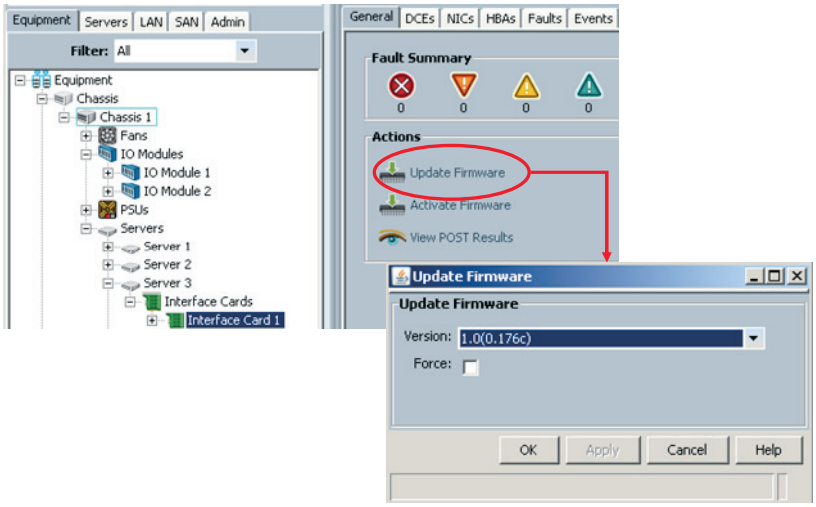

© 2009 Cisco Systems, Inc. All rights reserved.

To upgrade the adapter software, select the **Equipment** tab from the navigation pane. Then, select the **Chassis**, **Server**, and **Interface Card** objects. Click the **Update Firmware** or **Activate Firmware** link, as appropriate.

You can force the activation of an image by ignoring the result of the compatibility check during the activation process. Resetting on activation amounts to a reboot of the BMC after the activation is complete.

The Adapter firmware is for the “management entity” on the mezzanine card itself. It is this firmware that allows the UCS Manager to communicate with the mezzanine card out-of-band for status information, to configure the card, or to upgrade the vendor-specific BIOS or firmware (such as HBA firmware for a CAN).

https://t.me/learningnets

Upgrading the Adapter Firmware (CLI)

```
UCS-6100-A# scope server 1/1
UCS-6100-A /chassis/server # scope adapter 1
UCS-6100-A /chassis/server/adapter # update firmware 1.0(1.0e)
UCS-6100-A /chassis/server/adapter # activate firmware 1.0(1.0e)
```

© 2009 Cisco Systems, Inc. All rights reserved.

You can also upgrade the adapter firmware from the CLI, as shown in this example.

https://t.me/learningnets

Suggested Upgrade Practices

- Upgrade all components to the latest level available at initial installation, before deploying operating systems.
- Consult with operating system and application vendors for supported combinations.
- If possible, stage new code on spare compute nodes for testing.
- Upgrade from “outside-in” – upgrade components furthest from Fabric Interconnect first

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes suggested upgrade practices.

https://t.me/learningnets

Summary

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

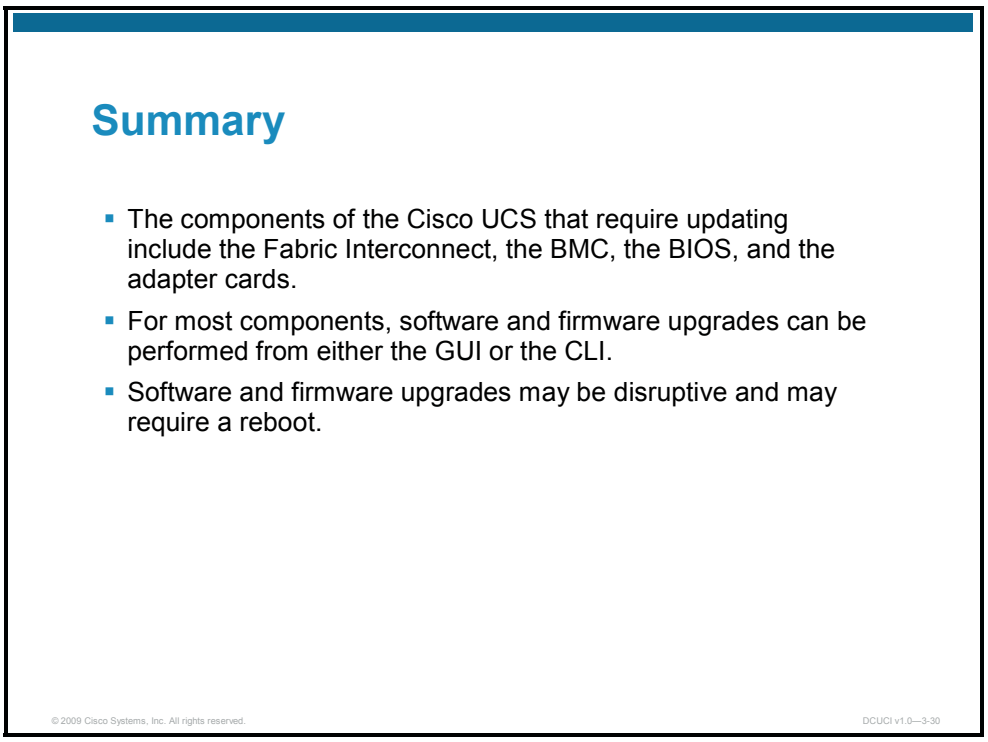

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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

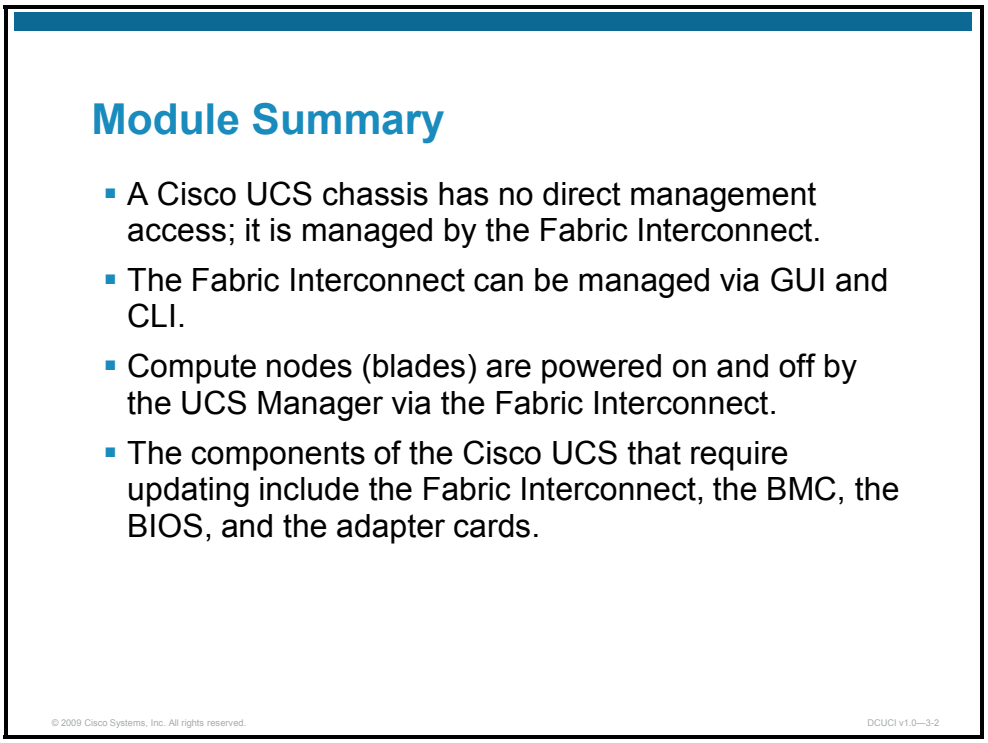

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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) What protocol does the Cisco UCS Fabric Interconnect use to negotiate the creation of PortChannels? (Source: Describing Unified Connectivity)

A) 802.3ad/802.1ax LACP
B) 802.1q
C) EtherChannel
D) 802.1x

Q2) In which navigation-pane tab will you find compute nodes? (Source: Exploring the Unified Computing User Interfaces)

A) Servers
B) Equipment
C) Service Profiles
D) Admin

Q3) In which command-line shell can the UCS 6100-series Fabric Interconnect be rebooted? (Source: Startup and Shutdown Procedures)

A) xos
B) clp
C) default
D) local-mgmt

Q4) Which answer shows the recommended order in which to upgrade UCS components? (Source: Upgrading UCS Software)

A) Fabric Interconnect, IO Module, BMC, Mezzanine Adapter
B) Mezzanine Adapter, IO Module, BMC, Fabric Interconnect
C) Mezzanine Adapter, BMC, IO Module, Fabric Interconnect
D) Fabric Interconnect, Mezzanine Adapter, BMC, IO Module

https://t.me/learningnets

Module Self-Check Answer Key

Q1) A

Q2) B

Q3) D

Q4) C

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 4

Understanding the Cisco UCS Manager

Overview

This module describes the components and functions of the Cisco UCS Manager.

Module Objectives

Upon completing this module, you will be able to explain the components of the Cisco UCS Manager. This includes being able to meet these objectives:

- Discuss the fundamentals of the Cisco UCS Manager
- Discuss Role Based Access Control within the Cisco UCS

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Cisco UCS Manager Fundamentals

Overview

This lesson discusses the fundamentals of the Cisco UCS Manager.

Objectives

Upon completing this lesson, you will be able to discuss the fundamentals of the Cisco UCS Manager. This includes being able to meet these objectives:

- Define the resources managed by Cisco UCS
- Explain the Cisco UCS server deployment model
- Describe Cisco UCS resource pools
- Describe Cisco UCS Service Profiles
- Explain the use of Cisco UCS Service Profile Policies
- Describe Service Profiles in the basic opt-in model
- Describe Service Profiles in the logical server model

https://t.me/learningnets

Cisco UCS Resources

This topic defines the resources managed by Cisco UCS.

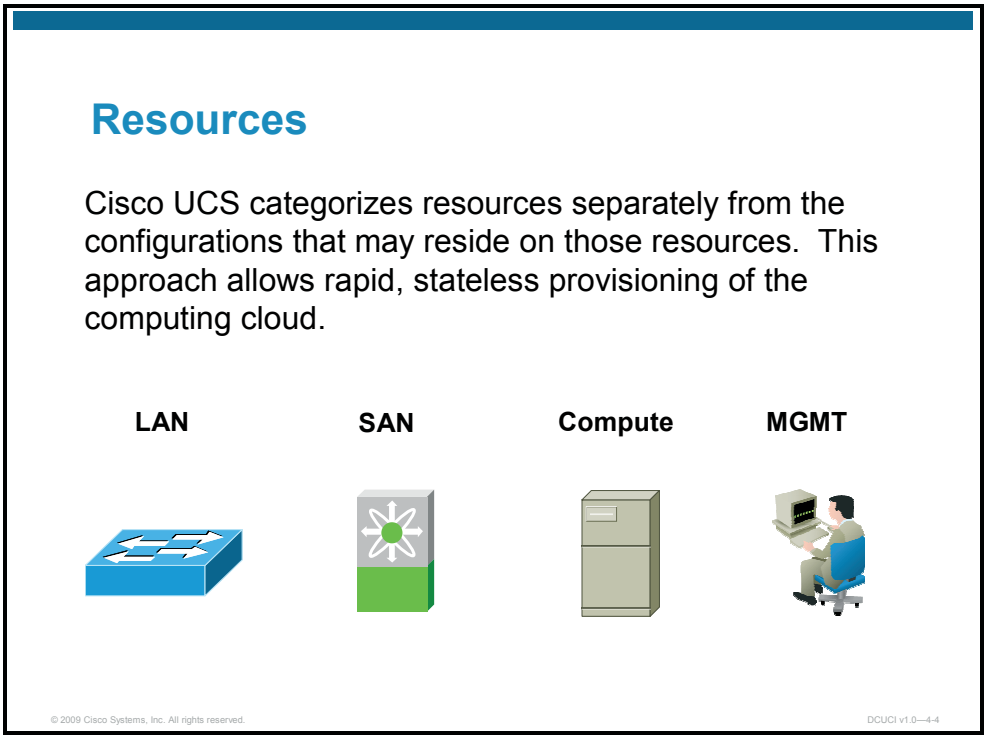

In order to provide a computing infrastructure that is not tied to physical devices, the Cisco UCS platform separates physical devices (resources) from their configurations (personalities). This abstraction provides a flexible environment where resources can be provisioned and migrated rapidly between physical devices.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

LAN Resources

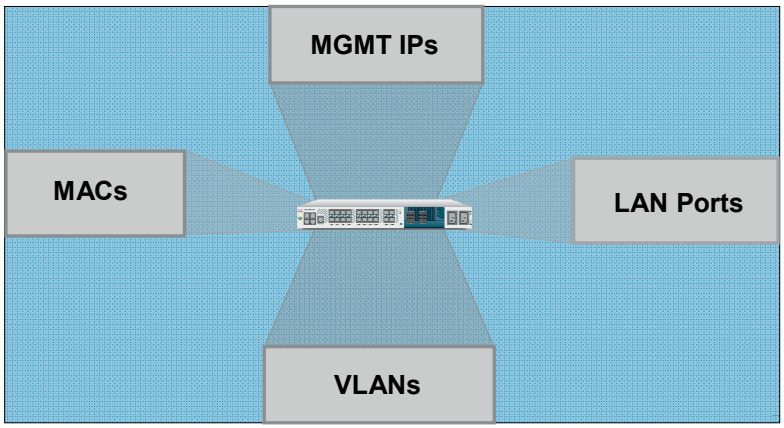

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS Manager controls LAN resources, including MAC addresses, management IP addresses, VLANs, and physical LAN ports.

https://t.me/learningnets

SAN Resources

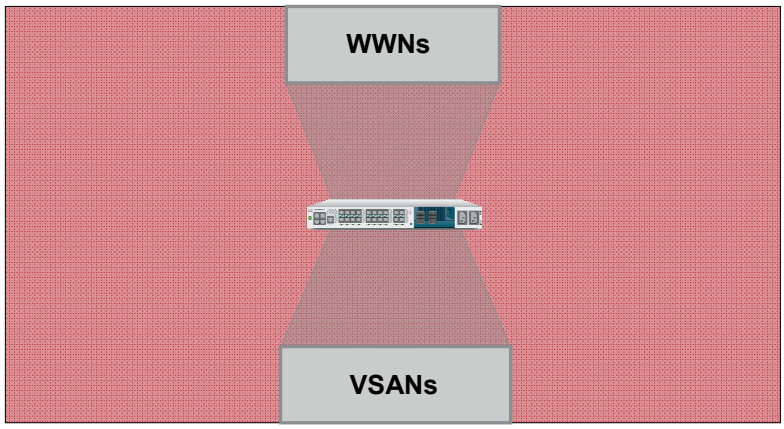

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS Manager is responsible for VSAN and world wide name (WWN) assignment. The central assignment of WWNs is critical to the portability of Service Profiles.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Compute Resources

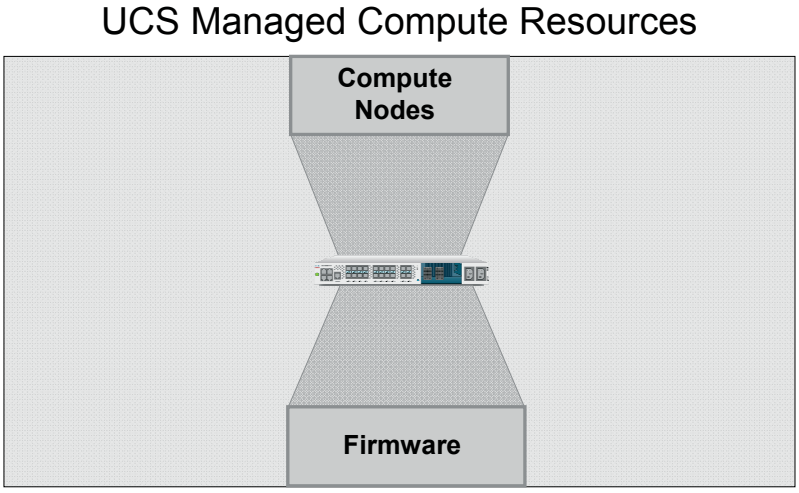

© 2009 Cisco Systems, Inc. All rights reserved.

The compute nodes and firmware are managed as resources within Cisco UCS.

https://t.me/learningnets

Management Resources

UCS Managed Management Resources

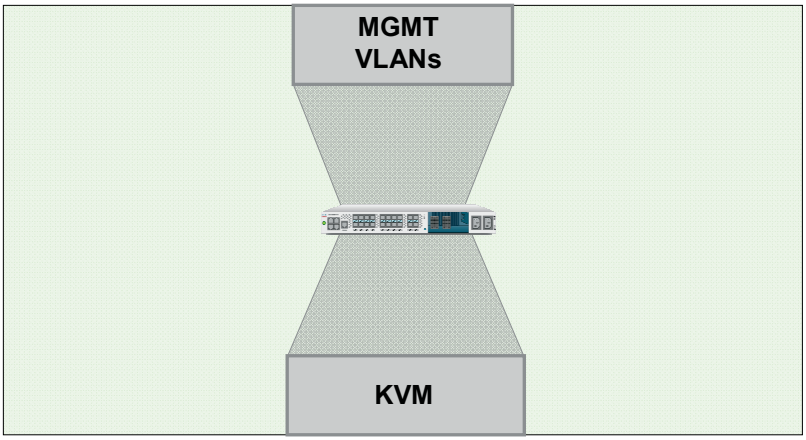

© 2009 Cisco Systems, Inc. All rights reserved.

All system management is handled through the Cisco UCS Manager running on the Fabric Interconnect. The management VLANs, server Keyboard/Video/Mouse (KVM) attachment, and low level control are all handled as Cisco UCS resources.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Service Profile Overview

This topic defines the concept of server profiles and how they enable the Cisco UCS server deployment model.

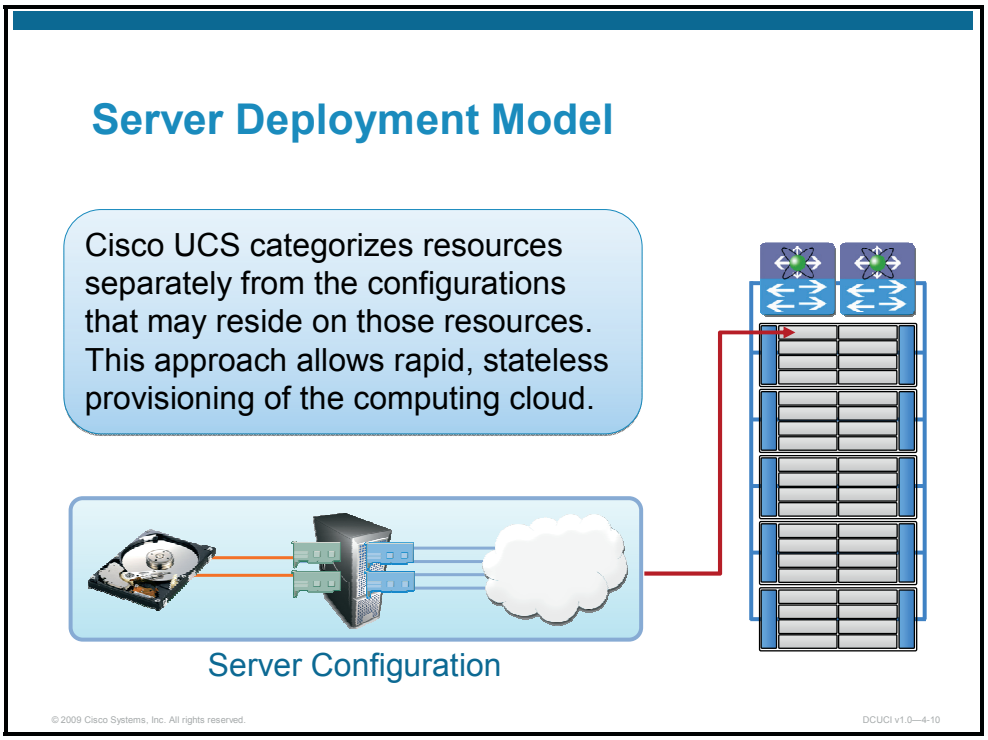

In order to provide a computing infrastructure that is not tied to physical devices, the Cisco UCS platform separates physical devices (resources) from their configurations (personalities). This abstraction provides a flexible environment where resources can be provisioned and migrated rapidly between physical devices.

To allow rapid, stateless provisioning of the computing cloud, UCS categorizes physical resources separately from the configurations that may reside on them:

- The UCS Manager running on the Fabric Interconnect controls various LAN resources, including MAC addresses, IP addresses, VLANs, and LAN ports.
- The UCS Manager also manages SAN connectivity through VSAN and WWN assignment. The central assignment of WWNs is critical to the portability of service profiles.
- The compute nodes (blade servers) and firmware are also managed as compute resources by the UCS Manager.
- All system management is handled through the UCS Manager running on the Fabric Interconnect. The management VLANs, server KVM attachment and low level control are all handled as UCS resources.

https://t.me/learningnets

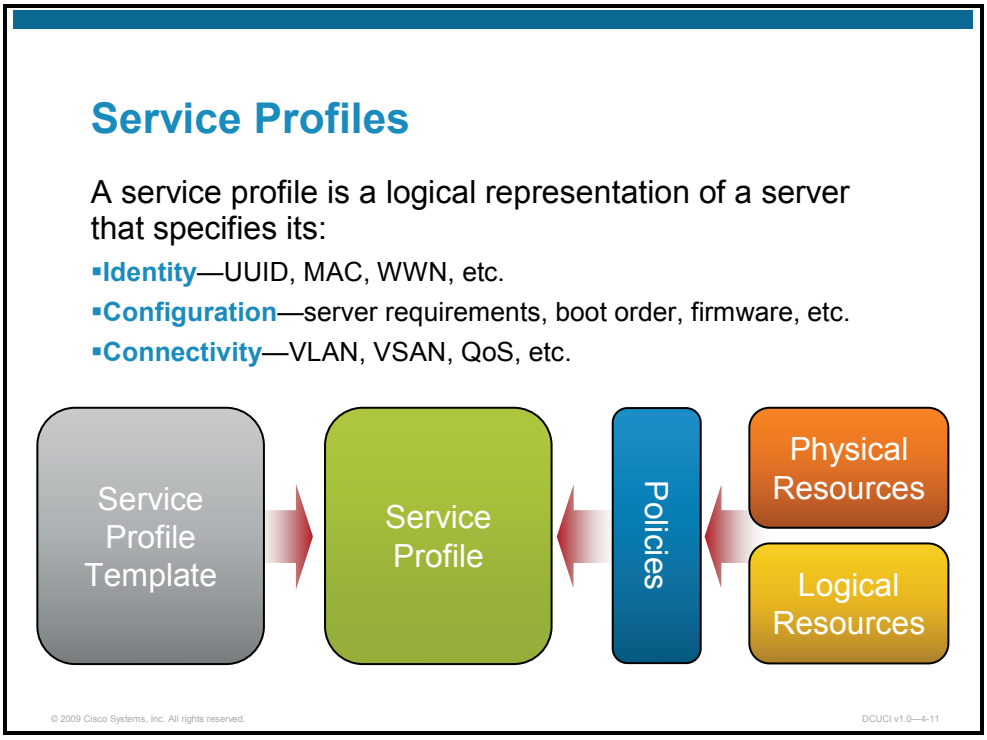

The figure shows the UCS server deployment model in which compute nodes (servers) are stateless. Servers obtain their properties from the UCS Manager, via pools of identity resources, and according to the policies that the UCS administrator specifies.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Resource Pools

This topic describes Cisco UCS resource pools.

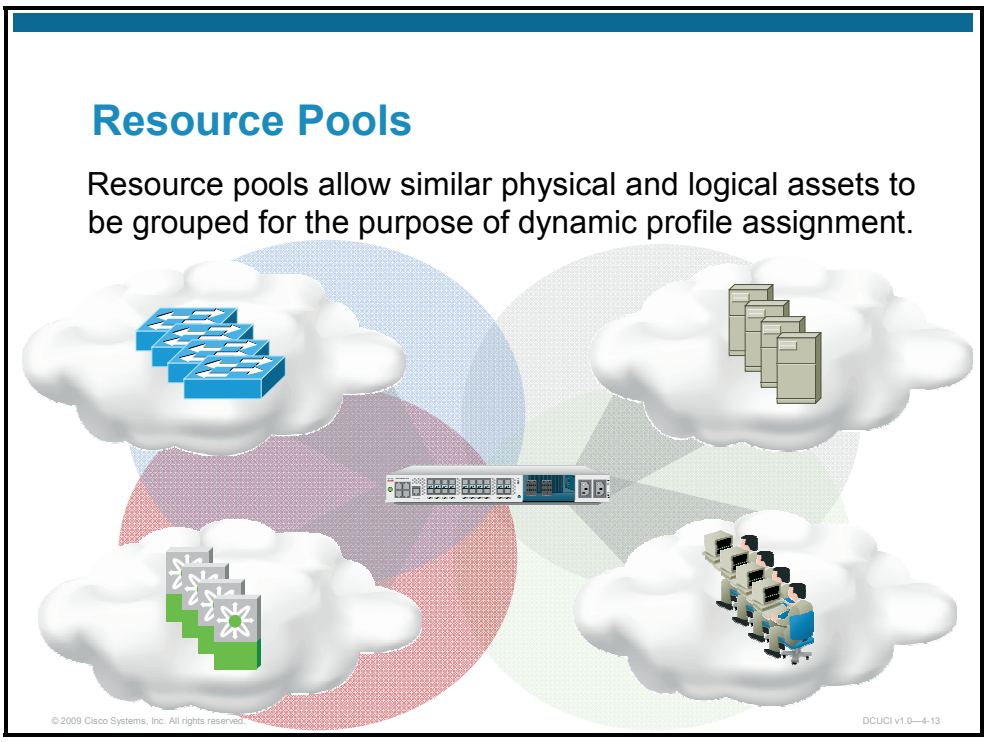

Resource pools allow common physical and logical resources to be grouped together. This grouping allows for rapid dynamic selection during server deployment and redeployment.

https://t.me/learningnets

Network Resource Pools

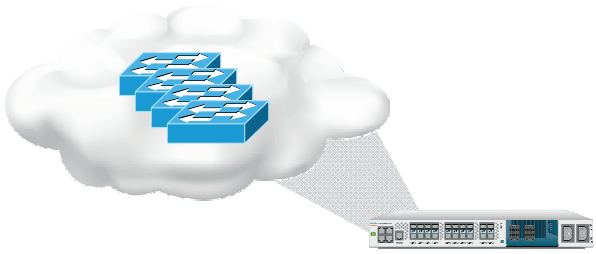

Network resource pools include MAC address pools:

- Pools of MAC ranges that allow server portability

© 2009 Cisco Systems, Inc. All rights reserved.

Network resource pools allow the dynamic assignment of unique server characteristics. MAC pools allow nodes to remain stateless by assigning addressing from pre-assigned pools of address space.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

SAN Resource Pools

- World Wide Node Name (WWNN) pools:
  - WWNN traditionally refers to a physical device, such as an HBA
- World Wide Port Name (WWPN) pools:
  - WWPN usually refers to a port on a physical device
  - A dual-port HBA would have two unique WWPNs, but share a WWNN

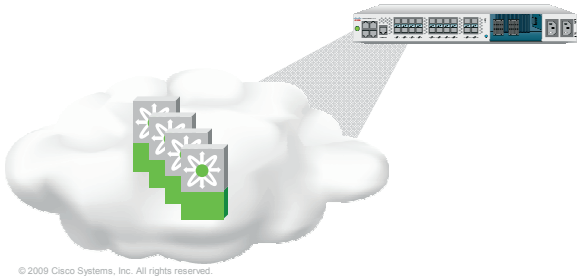

© 2009 Cisco Systems, Inc. All rights reserved.

Pools of WWNs remove the physical dependency of storage resources. Typically, storage LUNs (logical unit numbers) are masked to a physical port name, and zones are created using the same name. This approach forces a change in the SAN and array configuration when the physical hardware is changed. By virtualizing these addresses, Cisco UCS is able to ensure that hardware changes are invisible to the storage network.

https://t.me/learningnets

Compute Node Resource Pools

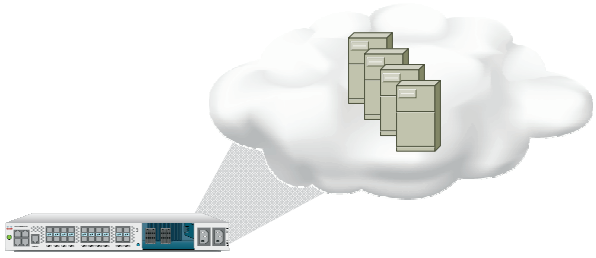

- Compute nodes can be pooled for use by Service Profiles
- Pools can be deterministic by policy
- Pools can be arbitrary by manual assignment

© 2009 Cisco Systems, Inc. All rights reserved.

Groups of compute nodes can be pooled. This pooling is typically based on processors, memory, and I/O device requirements. By selecting a pool instead of a physical blade, a Service Profile can be quickly deployed, based on pre-defined requirements. Administrators can also choose to manually pool servers for organizational or functional reasons.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Management Resource Pools

Management IP address pools provide automatic provisioning of KVM addresses.

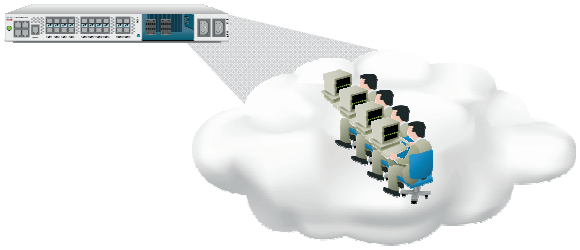

© 2009 Cisco Systems, Inc. All rights reserved.

IP address pools allow for the automatic provisioning of KVM addresses to the compute node KVM function.

https://t.me/learningnets

Service Profiles

This topic describes Cisco UCS Service Profiles.

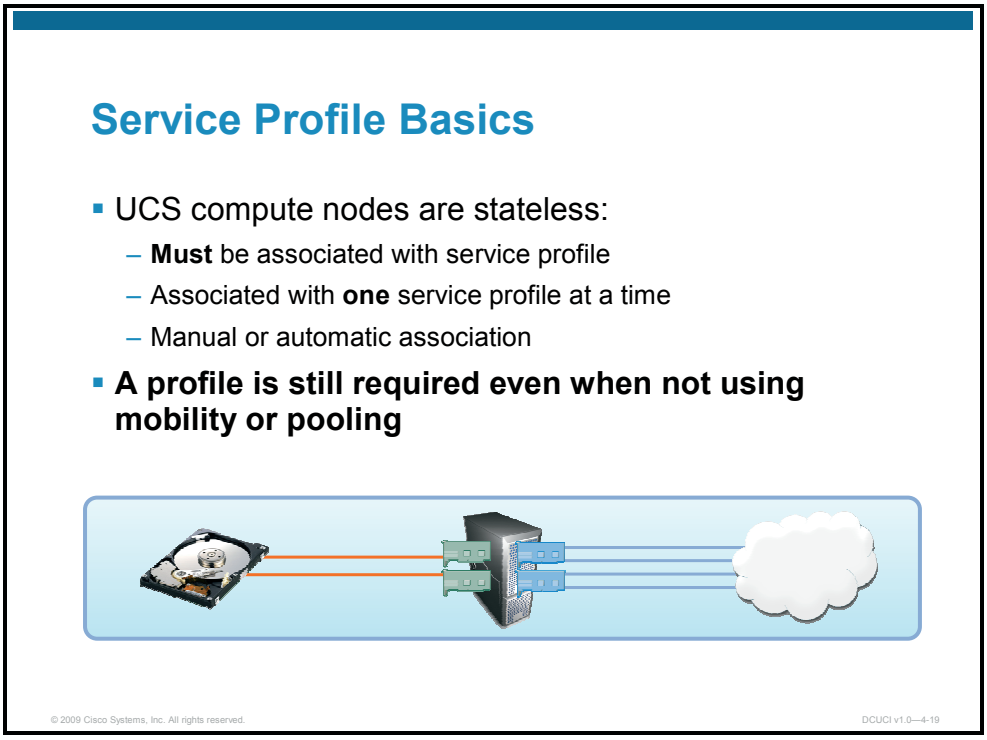

Understanding the Service Profile concept is key to understanding blade server management in Cisco UCS.

The Service Profile represents a logical view of a single blade server, without requiring you to know exactly which physical blade server it might be running on. The profile object contains the server personality (identity, network information, and so on) and connectivity requirements. The profile can then be associated with a single physical blade server (one blade server at a time).

The concept of profiles is important to the concept of mobility—transferring the identity of a logical server transparently from one physical blade server to another—as well as pooling concepts (discussed in the next lesson).

Even if you intend to manage the blade server as a traditional individual server, without taking advantage of mobility or pooling, you still need to create and manage a Service Profile for the blade server. While you could theoretically boot a blade server without a Service Profile, it would have no network or SAN connectivity.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

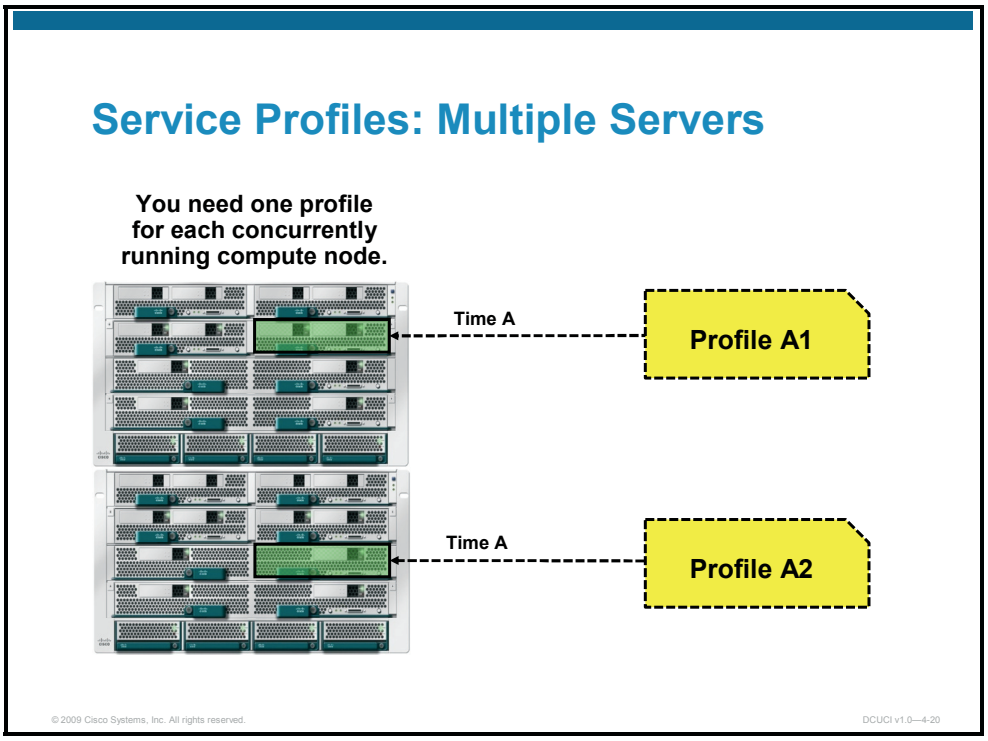

Every physical compute node that is running concurrently must have its own Service Profile, even if their profiles seem very similar (having exactly the same connectivity requirements, for example).

Two easy ways for creating large numbers of very similar Service Profiles are available. These are *cloning* and *templates*. These methods will be discussed later.

https://t.me/learningnets

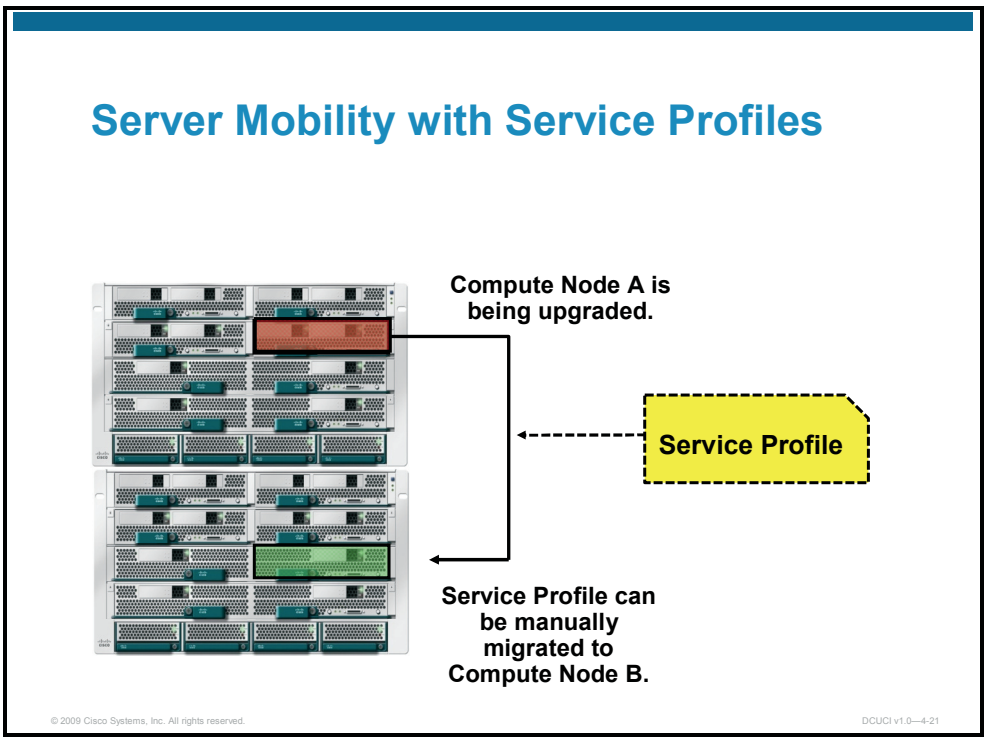

The figure shows the server mobility enabled through Service Profiles and stateless compute nodes. A logical server or Service Profile is defined with an identity including Universally unique identifier (UUID), MAC, WWN, and VLAN/VSAN assignment. The Service Profile is associated with one blade at a time but can be migrated in the event of failure or required maintenance.

Mobile logical servers are most effective when booted from SAN environments where the server's operating system image resides on a shared storage LUN. When the Service Profile is migrated, the new compute node inherits the server-specific properties that are stored within that Service Profile configuration. No configuration changes are required on the LAN or SAN infrastructures due to the portability of unique server characteristics such as WWPN and MAC address. Internal disks on compute nodes can be used for temporary storage or swap files.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

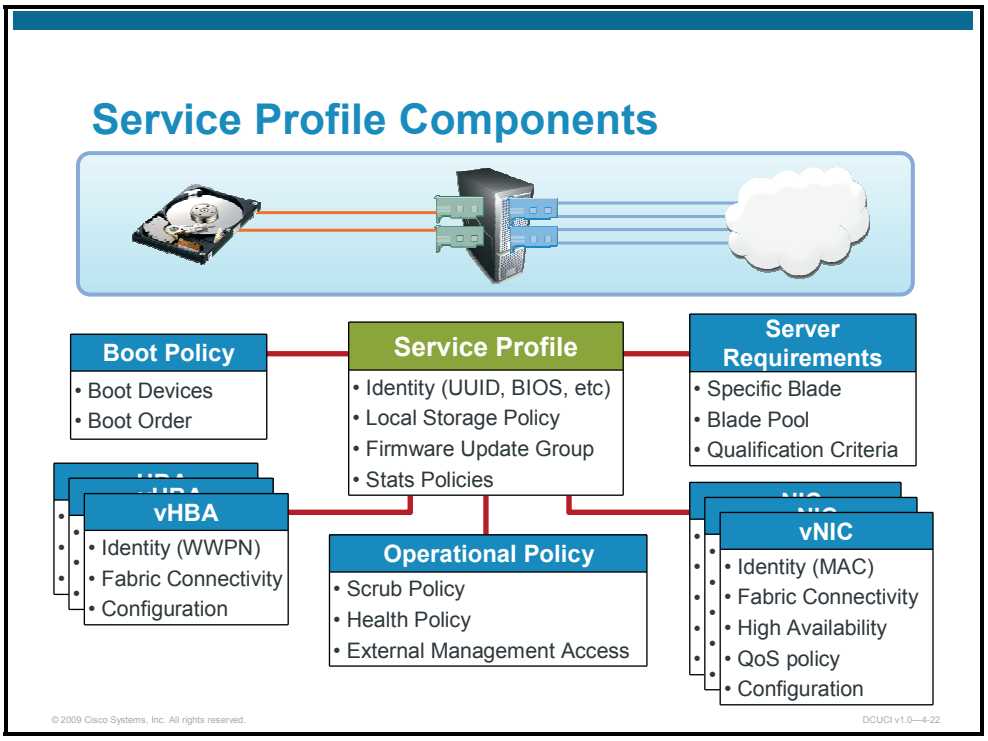

Service Profiles contain the personality (behavior and identity information) for a logical server. In this manner, servers are not tied to a specific physical compute node. The service profile defines server hardware (configuration, firmware, identity, and boot information), fabric connectivity, policies, external management, and high availability information.

The figure summarizes the contents of a Service Profile. Examples of blade server identity information that is unrelated to specific network or disk adapters includes the Universally Unique Identifier (UUID), BIOS boot order, and so on. The UUID for the server can have a significant impact on the portability of operating systems and applications. Some systems tie licensing to the UUID of the authorized server. The Service Profile also holds network adapter identity information. All of this information is physically applied to the blade server that is associated at a particular time with that profile.

In this manner, blade servers that are associated with a particular profile at different points in time are indistinguishable from each another, which can be advantageous when dealing with identity-based licensing schemes or network and external storage configurations.

https://t.me/learningnets

Service Profile Policies

This topic explains the use of Cisco UCS Service Profile Policies.

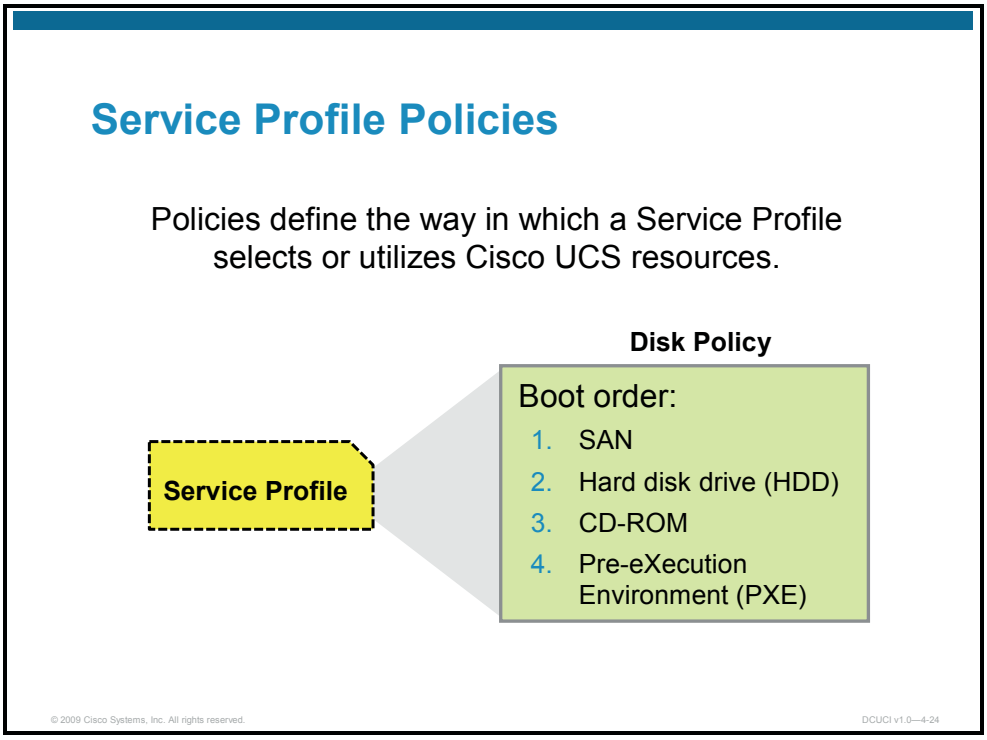

Policies defined within Service Profiles allow specific criteria to be selected during server deployment. These policies also define the specific configurations that will be applied at the time of deployment. In this example, a boot order policy can define how a compute node will be configured to boot.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

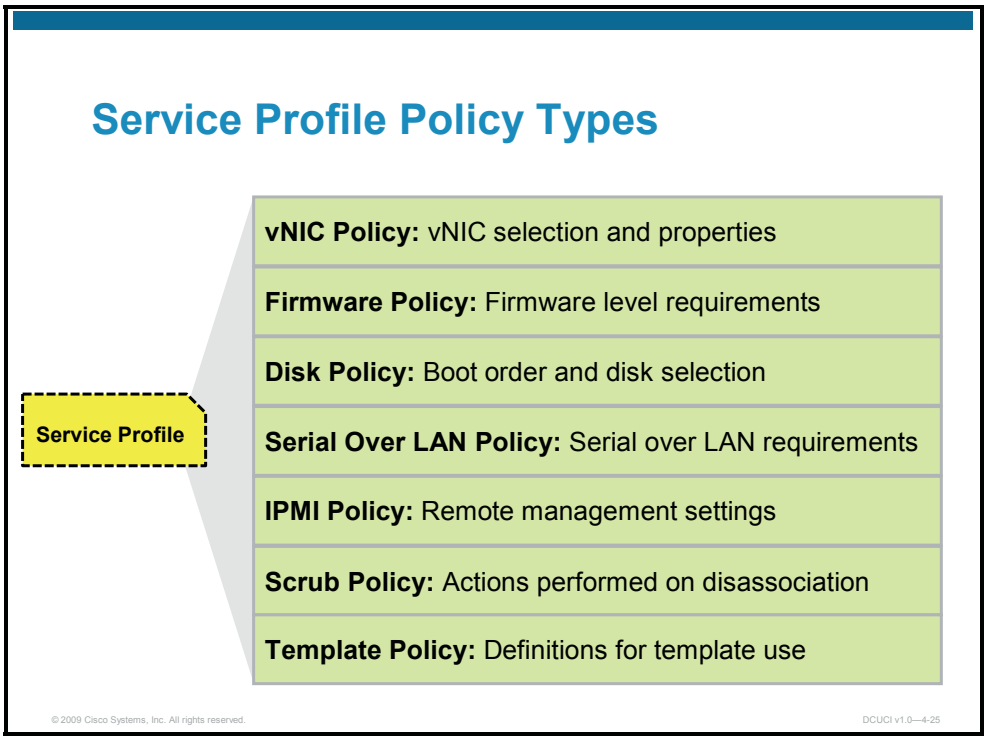

The figure illustrates the various types of Service Policies that can be defined within a Server Profile.

https://t.me/learningnets

Service Profiles in Basic Opt-In Model

This topic describes the Service Profiles in the basic opt-in model.

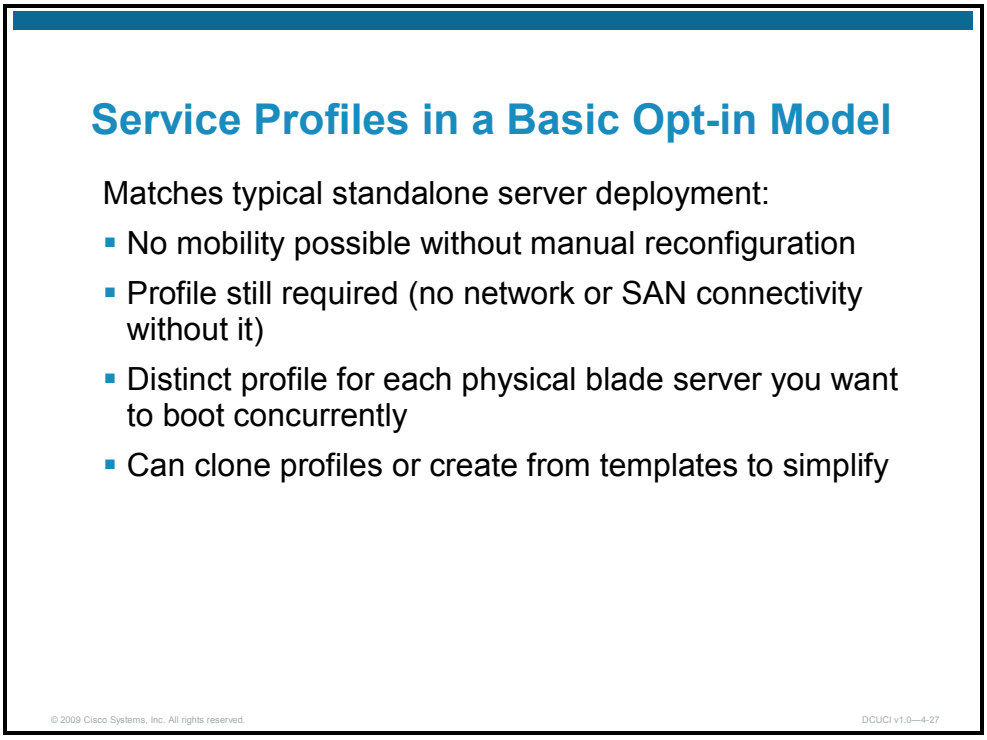

You can not use a compute node and give it network and SAN connectivity (if any) through Cisco UCS without a Service Profile.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Profile Elements in Basic Opt-in Model

- UUID:
  - Uses hardware defaults
  - Called “derived” addressing
- vNIC / vHBA:
  - Must create vNIC/vHBA for every adapter you want connected in blade server operating system
- vNIC/vHBA Identity (MAC/WWN):
  - Specify or use derived
- Boot order:
  - Set through profile
  - Leave empty to control via server BIOS

© 2009 Cisco Systems, Inc. All rights reserved.

In the basic model, the transference of identity information is not an issue.

In a Service Profile, you can use a special value of “derived” or, all 0s for MAC and WWN in the GUI. This implies that the hardware identities in a blade server’s BIOS and adapters will be used “as is”. If you want to supply actual values in the profile, you still can, and they will still be transferred to the blade server when it is associated with the profile. Derived addressing most closely matches traditional server deployments, but will limit portability of Service Profiles. This concept will be explored further in this lesson and in the following lesson.

You must create vNIC objects in your profile to define connectivity of the adapters in your blade servers on the network. This topic is discussed later in this lesson.

You must create vHBA objects in your profile to define connectivity for any Fibre Channel adapter (in the blade server) that needs to access SAN storage.

https://t.me/learningnets

Service Profiles in Logical Server Model

This topic describes Service Profiles in the logical server model.

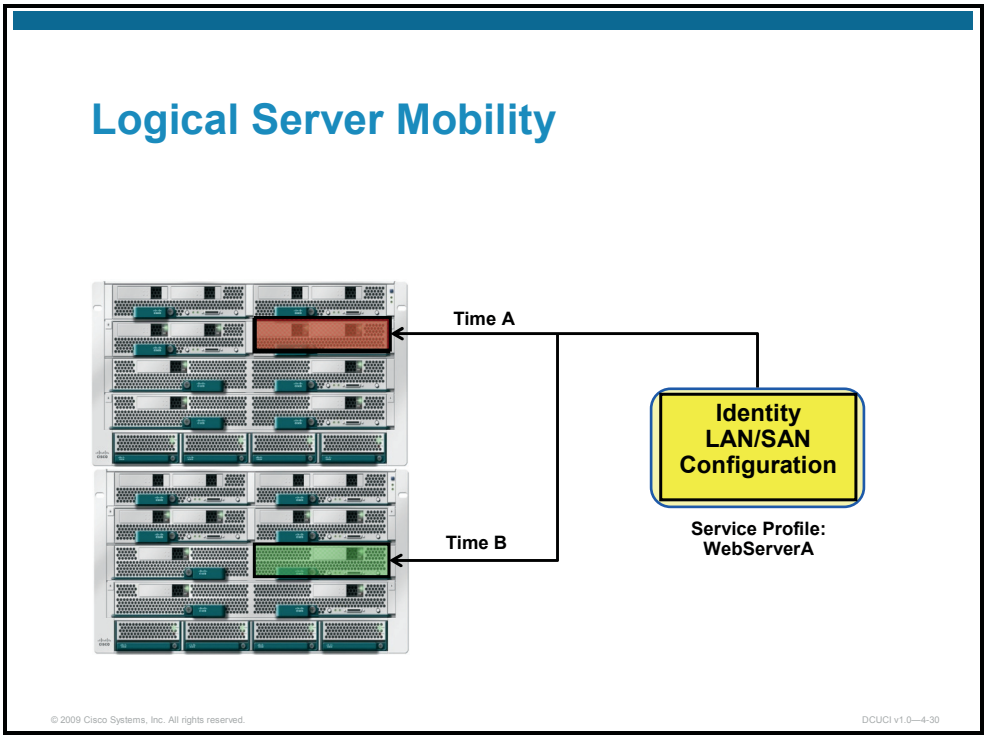

A logical server is defined with an identity that includes UUID, MAC/WWN addresses, and VLAN/VSAN requirements. The profile can be associated with only one blade server at a time, but the association can be changed if there is a problem with a particular blade server, or if hardware maintenance is required on that blade server.

SAN boot can be part of the design of a completely mobile logical server. A “replacement blade server” for the profile will inherit the WWN/VSAN information and boot specifications from the original blade server. From the point of view of the LAN/SAN infrastructure, the new blade server is identical to the old server. It will be exposed to, and boot exactly the same SAN LUNs as the original server, without any further reconfiguration.

Optional blade server local disks in this model could still be used for temporary storage and/or swap space.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Identity Information for Mobile Profiles

- UUID:
  - Do not use derived value
  - Generate UUID that moves with profile
  - Could also use UUID pool:
    - Guarantees uniqueness among all profiles that use same pool
- vNIC, vHBA identities (MAC, WWN):
  - Do not use derived value
  - Generate MAC/WWN that move with profile
  - Could also use MAC/WWN pools:
    - Guarantee uniqueness among all profiles that use same pool

© 2009 Cisco Systems, Inc. All rights reserved.

In the basic opt-in model, it was typical to use the derived values for UUID, and for adapter identities.

In the logical server model, you must define identities with the logical Service Profile that will be applied to the compute node.

These values and pools are discussed in detail in a later lesson.

https://t.me/learningnets

Server Farm Model

- Easy to create unique profiles with identical attributes:
  - Templates
  - Cloning
- Uses same attribute types as logical server model
- Templates allow for rapid deployment of many similar logical servers

© 2009 Cisco Systems, Inc. All rights reserved.

Cisco UCS lends itself to a model where server farms can be easily allocated and deployed, without having to pre-associate specific hardware to specific server roles.

Note that you still need a separate Service Profile for each server that is to run concurrently. If you want to deploy a farm of 50 web servers, you will need 50 separate profiles.

The following Cisco UCS features facilitate easy deployment of server farms:

- Profile cloning
- Profile templates
- Server pools

Note that the server farm model may imply that servers are being booted from SAN LUNs (then every member of the server farm is, itself, fully mobile) or from local disks, with some external operating system provisioning software.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

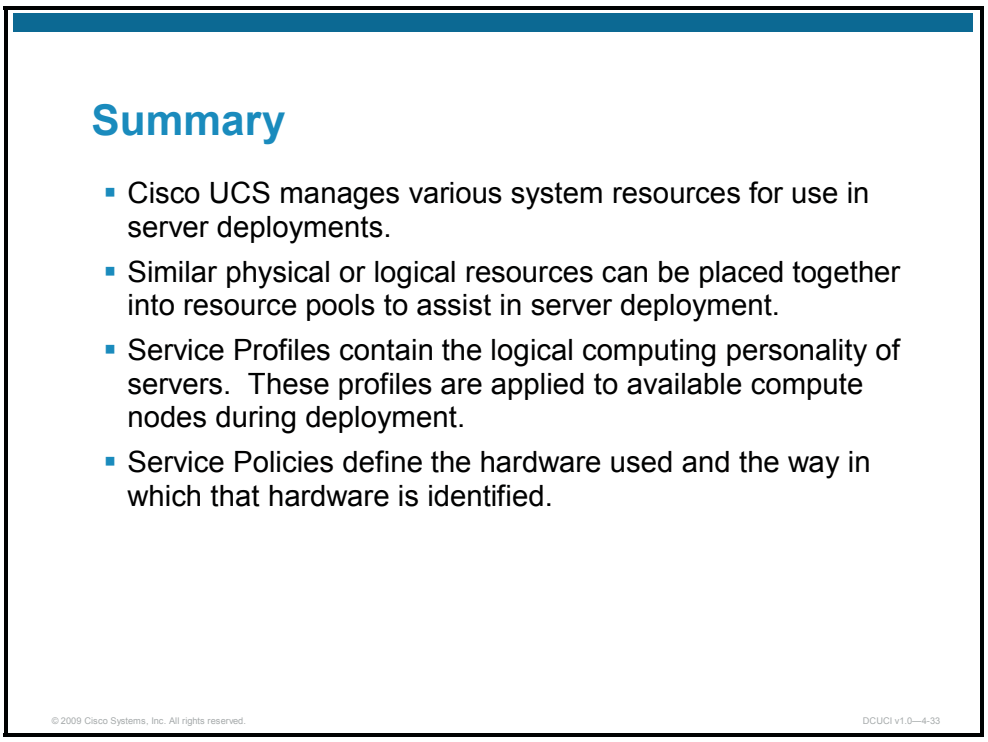

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Role-Based Access Control

Overview

In this lesson you will learn the fundamentals of Role-Based Access Control (RBAC), as they apply to the Cisco UCS.

Objectives

Upon completing this lesson, you will be able to discuss Role-Based Access Control within the Cisco UCS. This includes being able to meet these objectives:

- Explain Cisco UCS Management Groups
- Explain Role-Based Access Control
- Explain RBAC configuration
- Explain remote RBAC configuration
- Describe the configuration of AAA features and options
- Describe multi-tenancy features and options

https://t.me/learningnets

Management Groups

This topic explains Cisco UCS Management Groups.

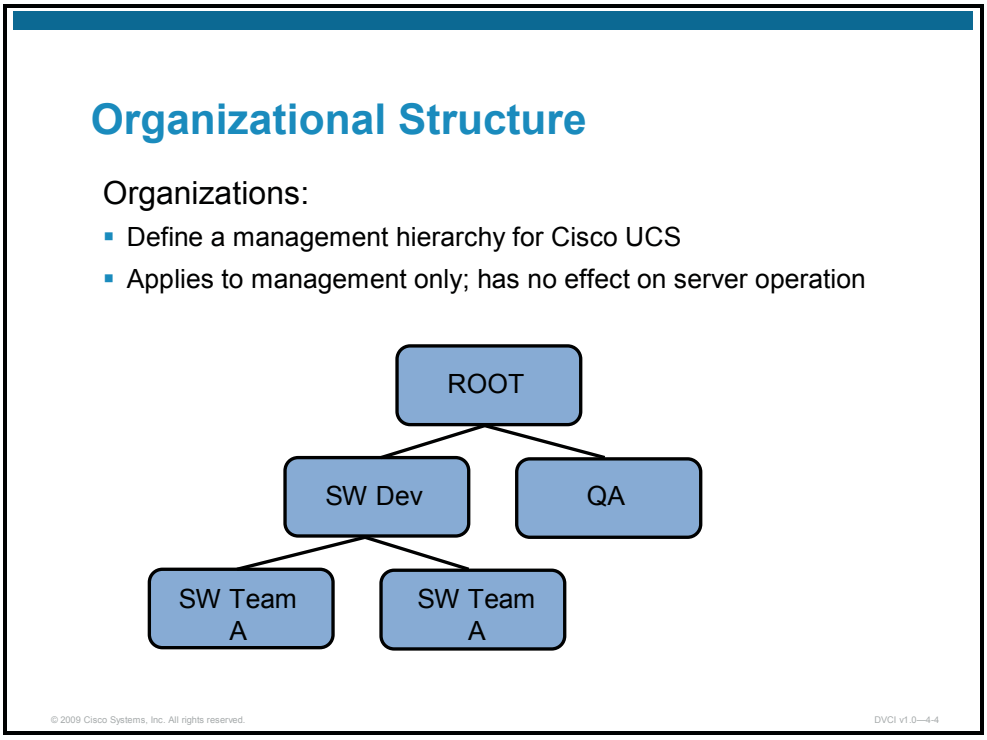

The organization structure defines a management hierarchy within Cisco UCS. The hierarchy is used strictly for management of resources and logical objects inside the system; it has no effect on the operation of servers and compute nodes.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

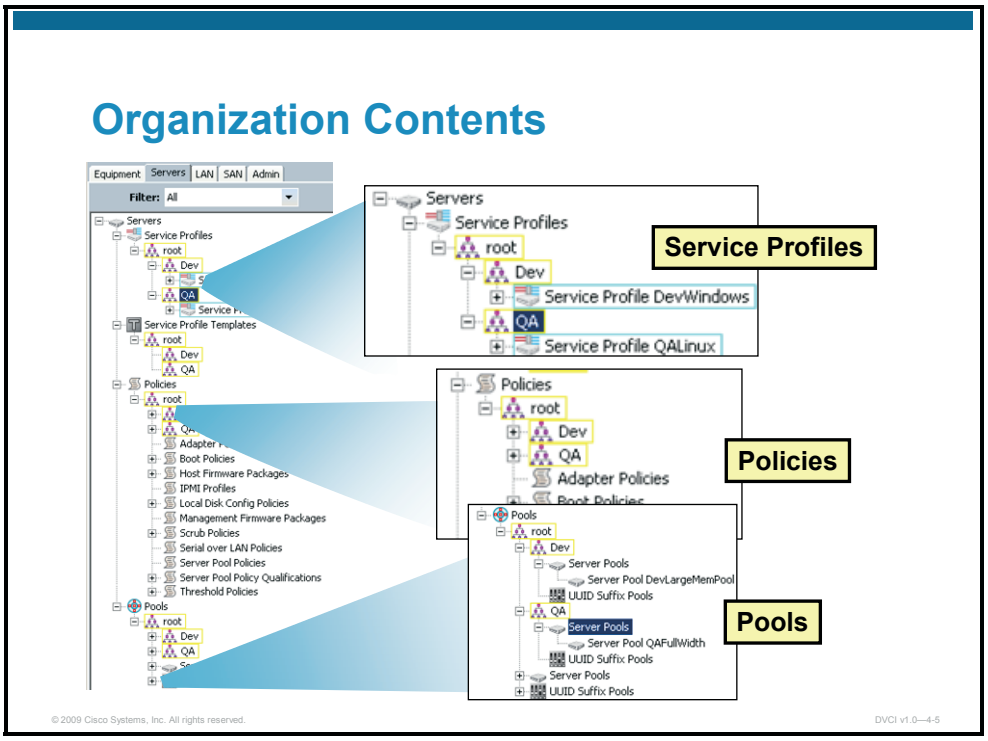

Service Profiles are the main component of an organization. They represent the server identities available for use by that organization. Each Service Profile is a member of only one organization. If no organization is defined, it will automatically be assigned as a member of the default “root” organization. Service Profiles and their member objects do not participate in organization inheritance (discussed later).

Pools are displayed at the bottom of the Servers tab. These include: server, world wide name (WWN), and MAC pools.

There is currently no direct way to move logical Service Profiles or any other object from one organization to another.

The figure shows the profiles and policies contained in an organization. When viewed in the GUI, links to configure all profiles and policies within that particular organization appear underneath that organization.

https://t.me/learningnets

Compute Node Ownership

- Compute nodes are independent of the organizational structure
- Nodes can exist in one or more server pools and organizations
- Nodes can be associated with Service Profiles in any organization
- The first Service Profile deployed on a node takes ownership

© 2009 Cisco Systems, Inc. All rights reserved.

Compute nodes are an exception to the organizational ownership structure. Compute nodes can be a member of multiple server pools and organizations. This allocation method allows for better resource utilization by allowing nodes to be disassociated from a role and redeployed for a priority application in another organization.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Server Pools

- Server pools are inherited by child organizations
- Child organizations have access to all server pools above them in the hierarchy

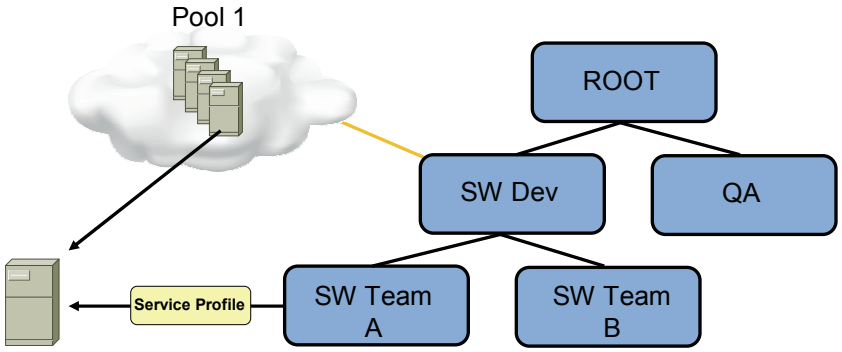

© 2009 Cisco Systems, Inc. All rights reserved.

Server pools are defined within an organization. Child organizations inherit the privileges of parent resource pools. If pools are defined in a organizational hierarchy with the same name, the local pool will always be used first. This means that when deploying a Service Profile that calls for an available resource from “Pool 1”, the Cisco UCS Manager will first search the local pool, then move up the hierarchy searching for “Pool 1” at each level, until that pool is located.

In this example, a Service Profile is deployed from the “SW Team A” organization using the server pool “Pool 1”. Because “Pool 1” does not exist in organization “SW Team A”, the Cisco UCS Manager searches up the tree to find it. “Pool 1” exists in the parent organization “SW Dev”, so that pool will be used to select a physical server to fulfill the requirements of this Service Profile.

https://t.me/learningnets

Server Pool Association Example

- Pools “RootServerPool” and “DevLargeMemPool” are both available to the “Dev” organization

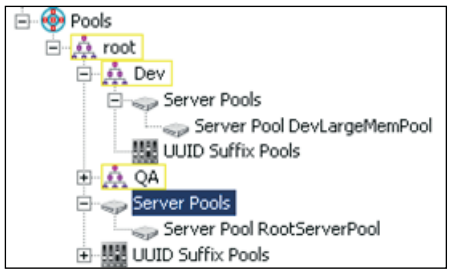

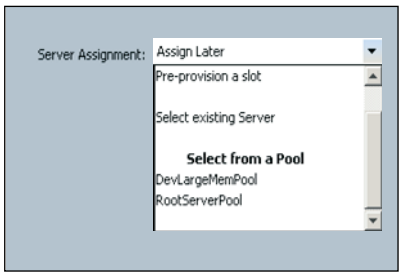

© 2009 Cisco Systems, Inc. All rights reserved.

In this example, a Service Profile is being created in the “Dev” organization, which is a child of root. The “Dev” organization has server pool “DevLargeMemPool” defined within the organization, and root has server pool “RootServerPool”. The new Service Profile can choose between either of these pools because it is a child of the root organization.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Other Resource Pools

MAC, WWN and UUID pools operate in the same manner as server pools.

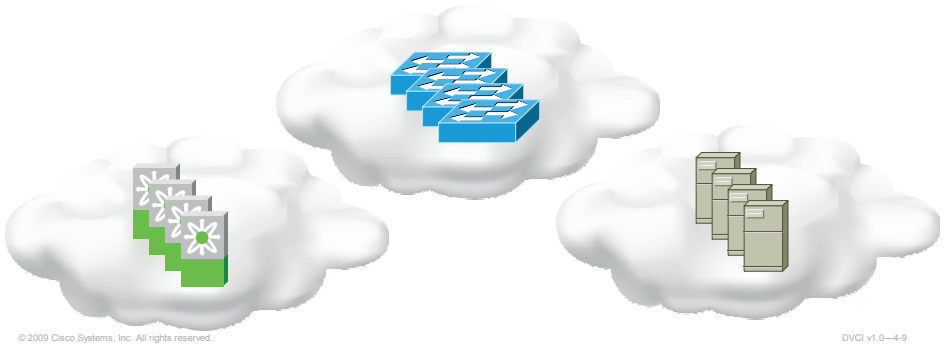

© 2009 Cisco Systems, Inc. All rights reserved.

MAC WWN and universally unique identifier (UUID) pools operate in the same manner as server pools. An organization will have access to its own pools, as well as those of the organizations that exist above it in the hierarchy.

https://t.me/learningnets

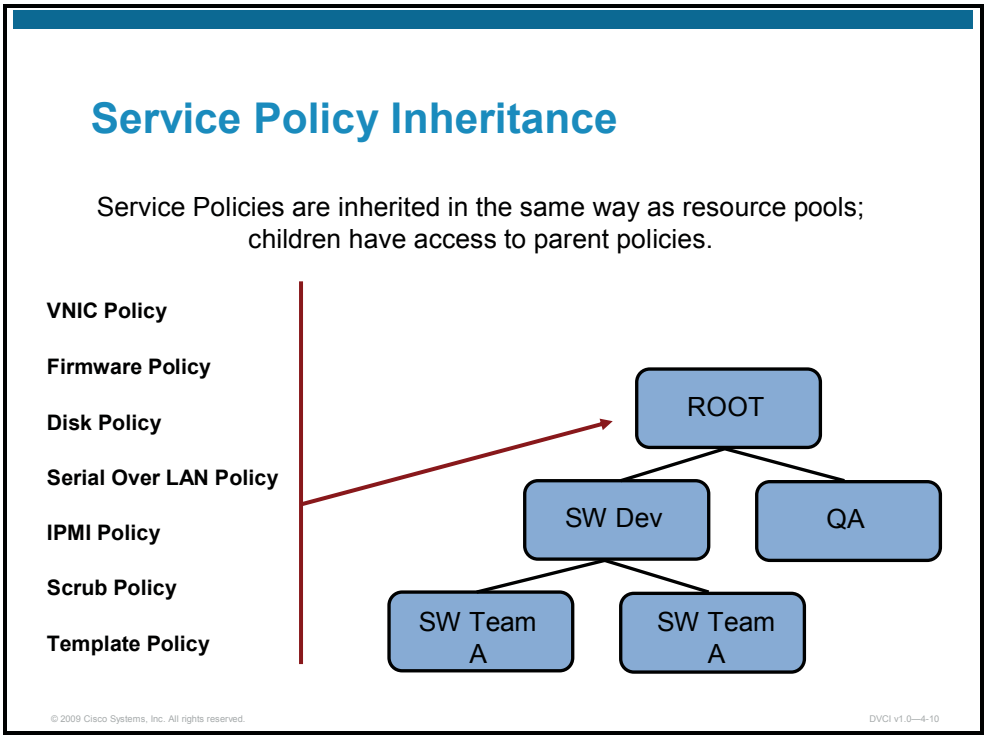

Service Policies operate the same way as resource pools from an organizational standpoint. They are defined within a specific organization and are inherited by all child organizations.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

RBAC Overview

This topic explains Role-Based Access Control (RBAC).

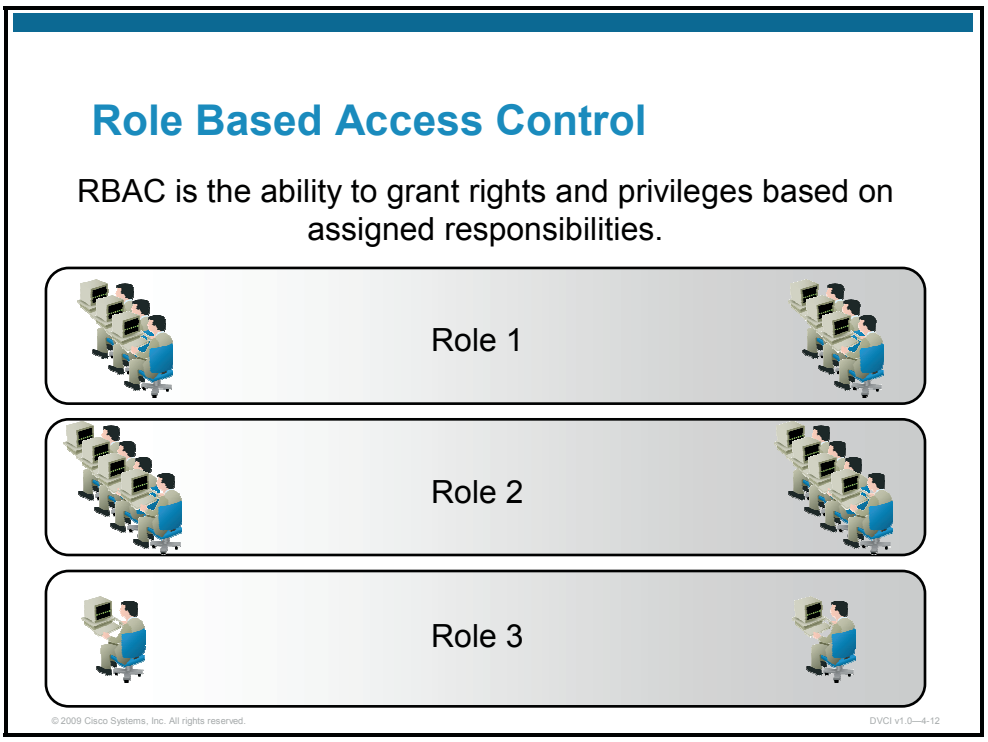

RBAC allows users to be granted granular permission sets, based on their responsibilities. A series of roles are defined by a Super User or system administrator and then assigned to other users.

https://t.me/learningnets

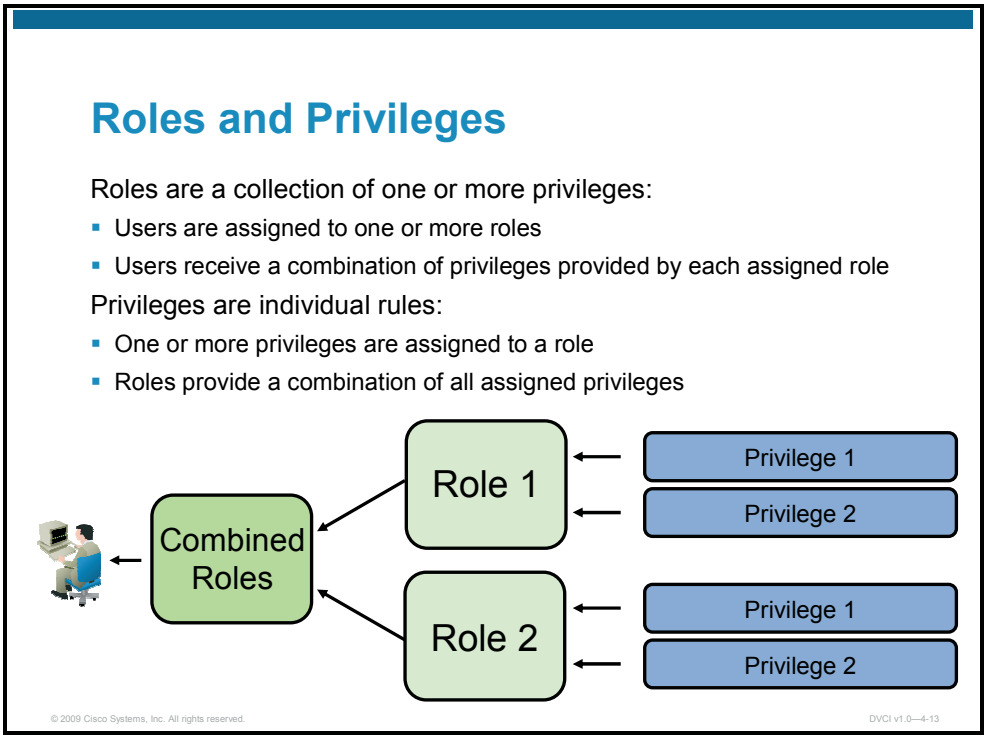

Privileges are the building blocks of roles. Each role is defined with one or more privileges. Users receive rights or privileges based on assigned roles. A user may be assigned one or more roles. When users are assigned more than one role, they receive a combination of the privileges defined in each role. This means they will have all of the privileges defined in each assigned role. This is important to note because vendor RBAC schemes differ and other products or systems may operate in another fashion.

Example privileges:

- Admin - Can do everything on the system (across all organizations), as a Super User
- AAA - Authentication, authorization, and accounting (including administration of the RBAC feature itself)

Note	Neither of these privileges is organization specific.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

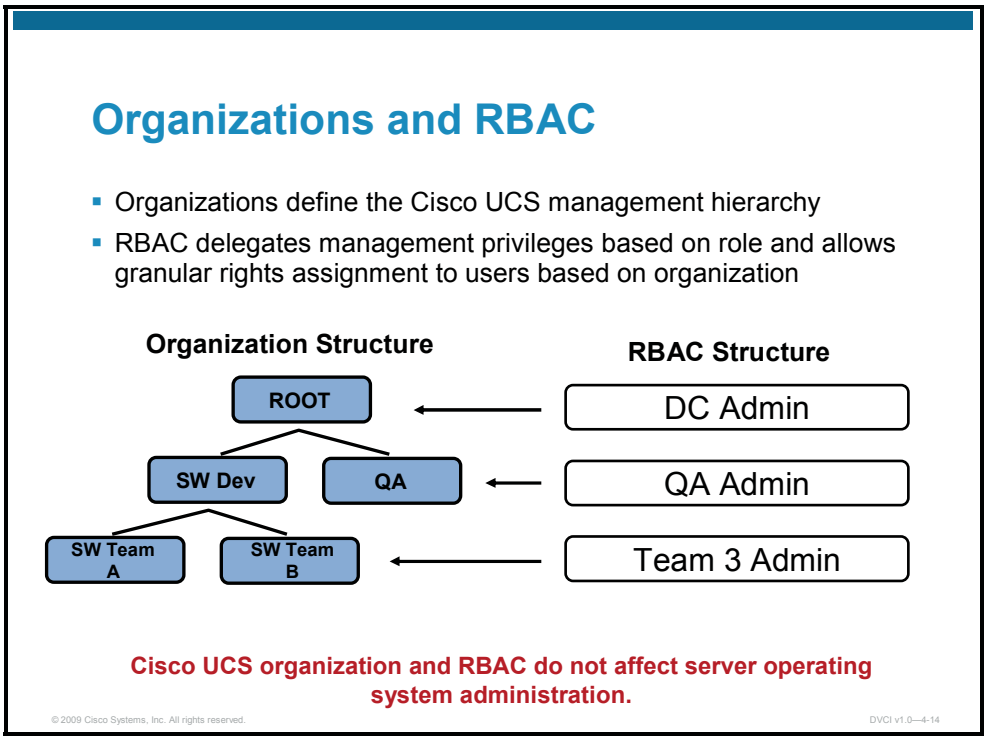

The organization structure defines a management hierarchy within the Cisco UCS. This hierarchy is used to assist the management of the resources and logical objects used within the system. RBAC delegates rights to users based on the organization where they exist. The RBAC feature allows for fine-grained privilege assignment and restricts user privileges to specific organizations.

https://t.me/learningnets

Organizations and RBAC (Cont.)

- Organizations can be used without RBAC:
  - Structural management hierarchy
- RBAC can be used without defined organizations:
  - All rights assigned at the root level
- Locales are groups of organizations

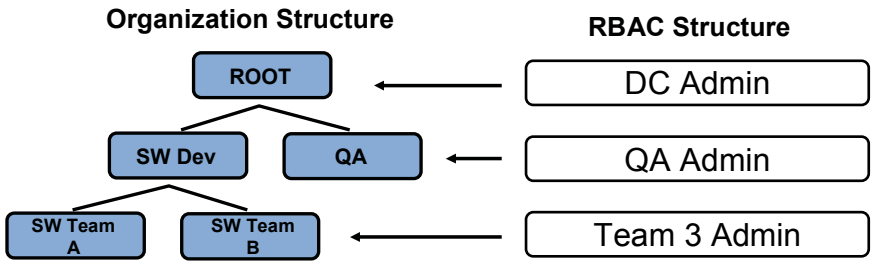

© 2009 Cisco Systems, Inc. All rights reserved.

RBAC and organizations are complementary and not dependant on each another. By themselves, organizations provide a logical management hierarchy allowing granular resource assignment. RBAC (without defined organizations) places all defined roles at the default root level to provide role-based resource control to the entire Cisco UCS system.

https://t.me/learningnets

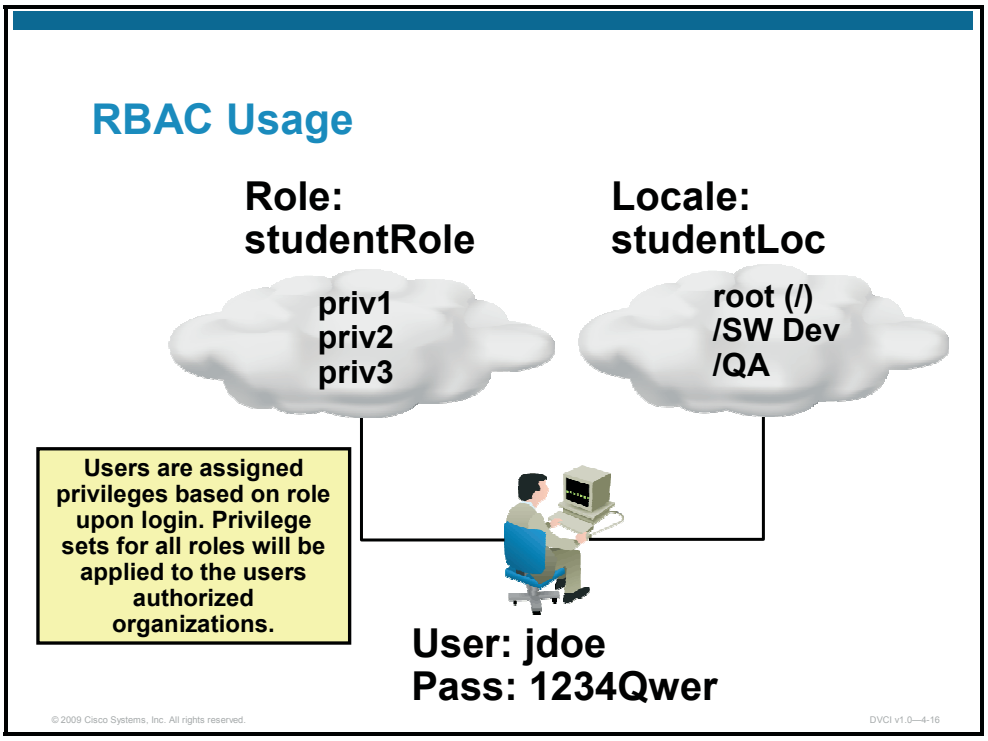

When users are created, they are assigned to one or more roles and one or more locales.

A locale is a collection of one or more organizations. The organizations do not need to be related within the hierarchy. The purpose of a locale is to apply a user's roles to specific subsets of organizations. Locales are only used within the RBAC scheme.

A role is a set of privileges. There are several predefined roles, and custom roles can be created with user-defined privileges.

Note	There is no way a user in two different organizations can have different privileges in those two organizations. All privileges are cumulative and applied in combination to all the organizations in the users' specified locales.

Some privileges are not org-related. These privileges will apply regardless of the users assigned locales.

https://t.me/learningnets

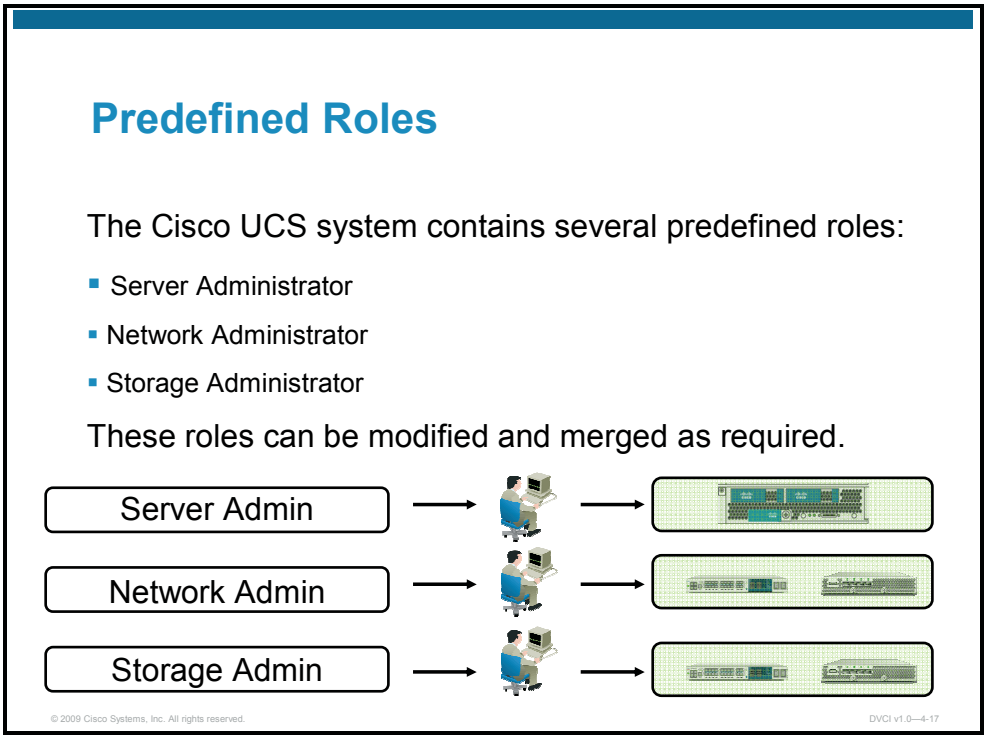

Cisco UCS comes with three pre-defined or default roles. These roles allow out-of-the-box RBAC for the most common Cisco UCS administration tasks:

- The Server Admin role allows administration of physical server resources, Service Profiles, and other activities related to logical and physical servers.
- The Network Admin role allows configuration of VLANs, uplinks, trunks, and other activities related to Ethernet network connectivity.
- The Storage Admin role allows configuration of VSANs, uplinks, and other activities related to Fibre Channel network connectivity.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

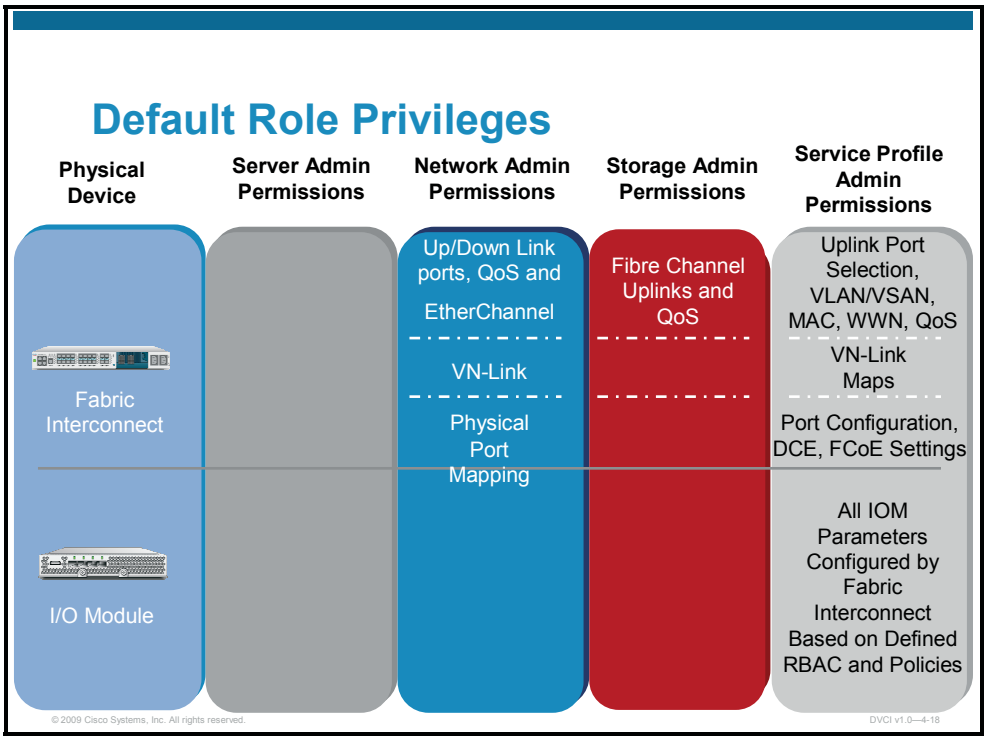

The figure illustrates the physical and logical resources that could be configured by an administrator assigned to one of the pre-defined roles.

https://t.me/learningnets

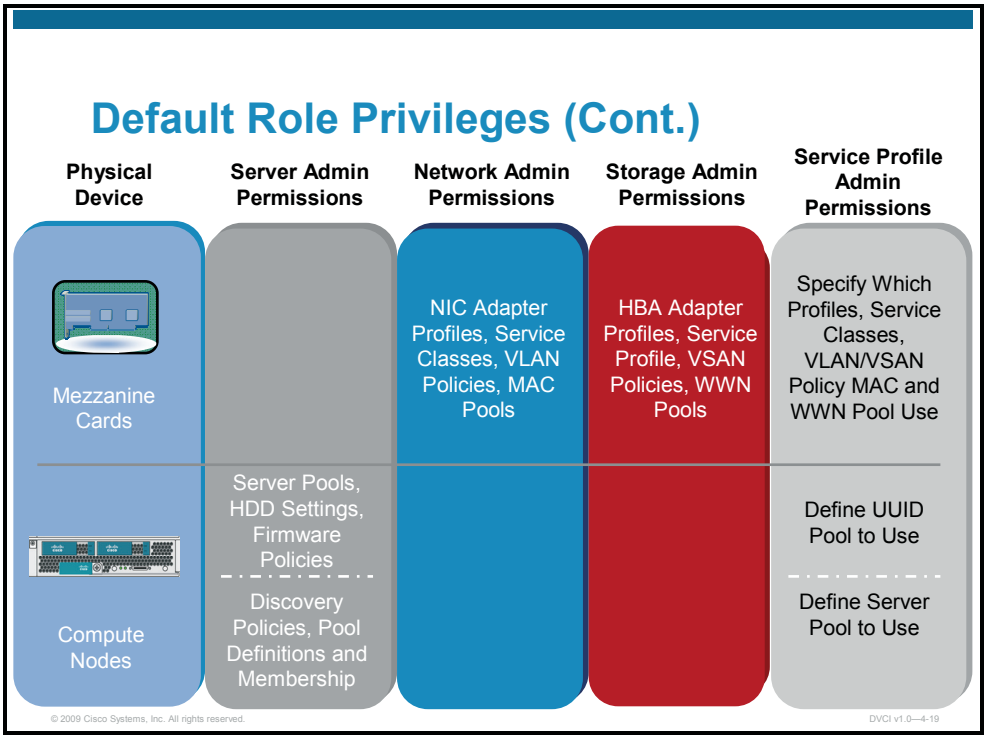

The figure illustrates the physical and logical resources that could be configured by an administrator assigned to one of the pre-defined roles (continued).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

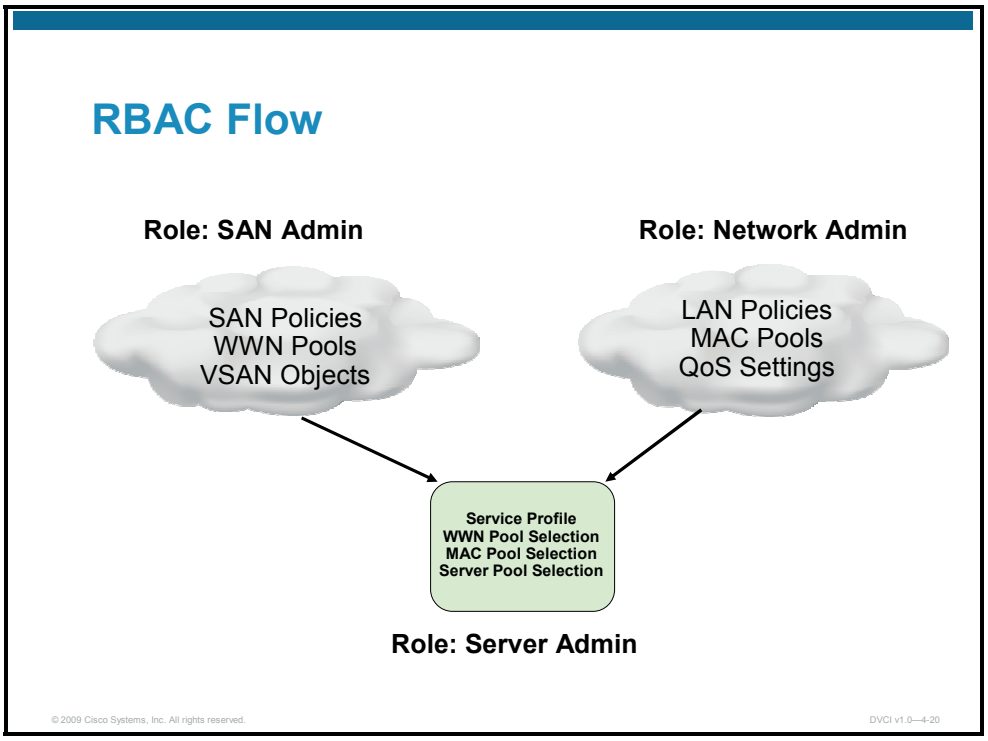

RBAC provides a powerful method of segregating responsibilities within Cisco UCS. This figure demonstrates how administrators with different roles collectively provide the structure necessary for the operation of a logical server. In this example, a SAN Administrator creates policies and objects related to the SAN. A Network Administrator creates policies and objects for Ethernet access. Finally, a Server Administrator can take the policies and objects created by the other two administrators and combines them into a Service Profile. Each administrator is responsible for their areas of expertise without having to administer the entire architecture.

https://t.me/learningnets

RBAC Configuration

This topic explains RBAC configuration.

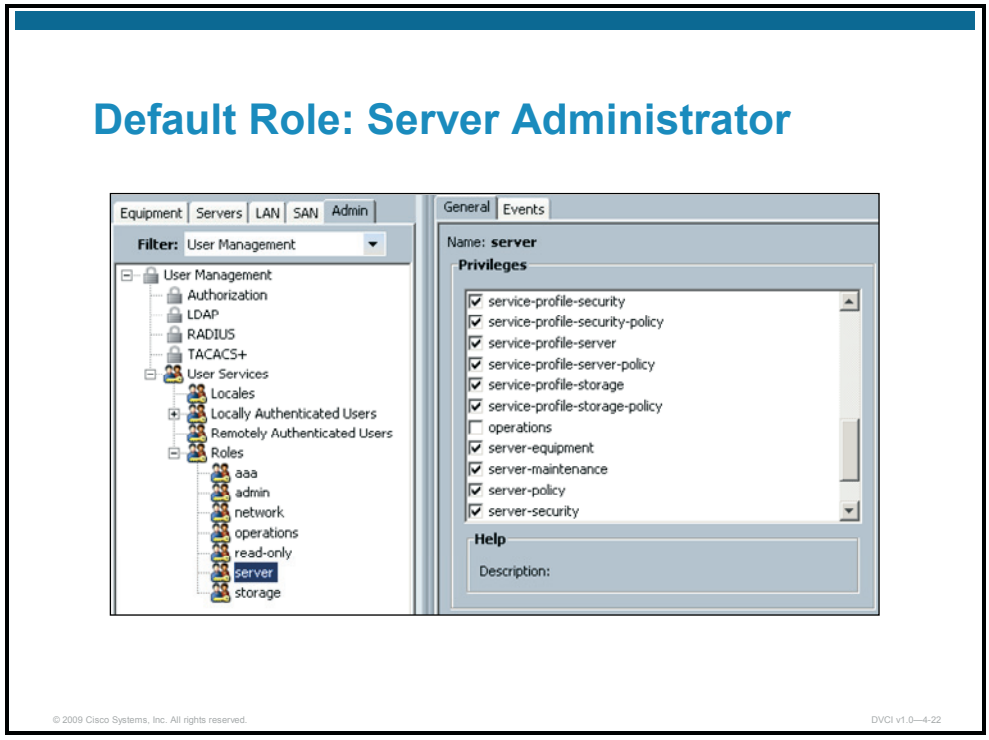

The figure shows the set of privileges for the Server Administrator role. These include all of the privileges that begin with “service-profile”, administration of logical server profiles, and all policies that apply to them.

Privileges for the Server role are organization-specific. A user given this role will have these particular privileges only within the organizations that are part of the user’s assigned locales.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Default Role: Network Administrator

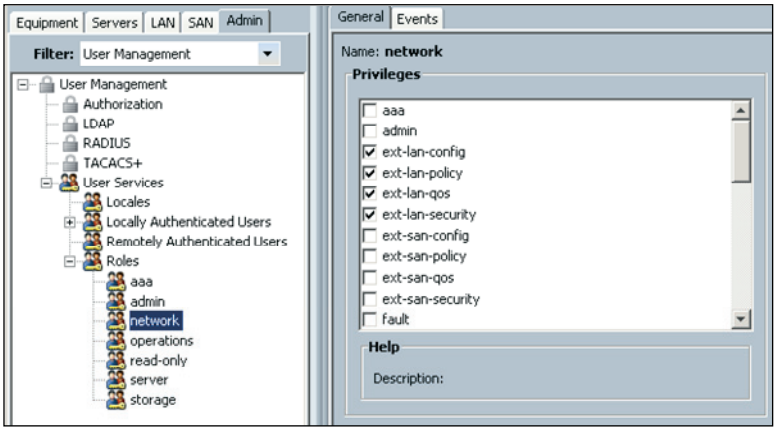

© 2009 Cisco Systems, Inc. All rights reserved.

The Network Administrator role defines a set of privileges that are not organization-specific. These privileges involve configuring the global LAN settings. For example, supported VLANs, and whether those VLANs are supported globally or per-switch.

https://t.me/learningnets

Default Role: Storage Administrator

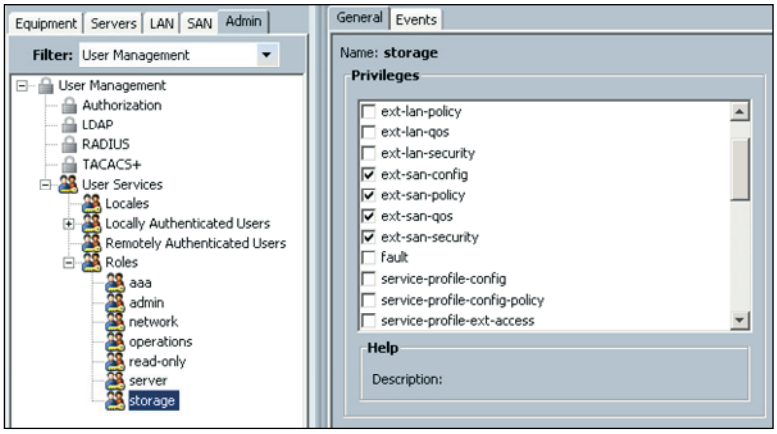

© 2009 Cisco Systems, Inc. All rights reserved.

The Storage Administrator role defines a set of privileges that allow configuration of the SAN network within Cisco UCS. These include border interfaces or SAN uplinks and internal VSAN configuration. Like the Network Administrator role, these privileges are not organization-specific.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

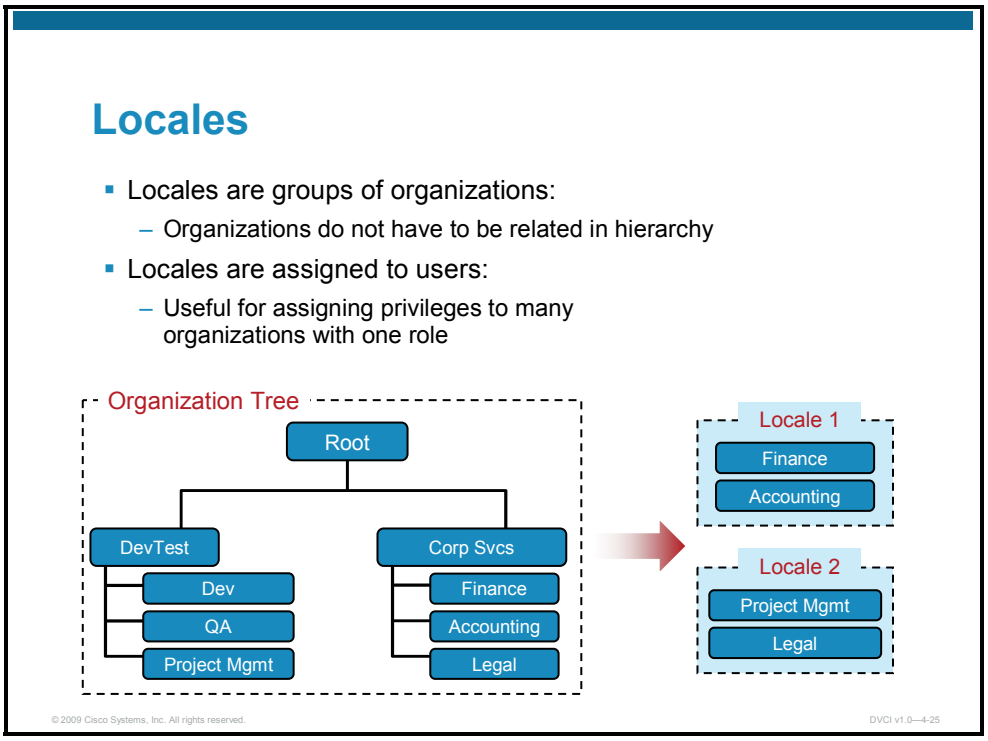

Locales provide a mechanism to group multiple organizations together for purposes of granting roles to users. A user can be assigned one or more locales. Each locale defines one or more organizations to which the user is allowed access.

Locales can be used without external authentication systems for local users, but are required for external AAA servers, such as RADIUS, TACACS+, or Lightweight Directory Access Protocol (LDAP).

https://t.me/learningnets

Creating Locales

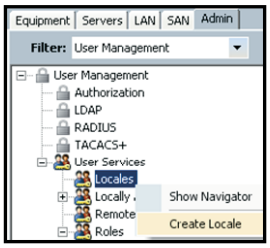

- Right click **Locales**
- Select **Create Locale**
- Assign **Organizations** to the **Locale**

© 2009 Cisco Systems, Inc. All rights reserved.

Locales are sets of organizations used only within the RBAC scheme. They allow subsets of organizations to be assigned to users quickly.

Sets of organizations can be dragged into the locale creation pane for faster locale configuration.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Users

- A user object contains Authentication and Authorization credentials as well as role and locale assignment
- Users have assigned roles in configured locales

New User:

User Name:	**jdoe**
Password:	********
Locale(s):	**studentLocale**
Role(s):	**studentRole1**
	studentRole2

© 2009 Cisco Systems, Inc. All rights reserved.

User objects are the final building blocks of the RBAC scheme. User objects contain definitions of the login credentials, locales, and the roles that an administrator can use.

https://t.me/learningnets

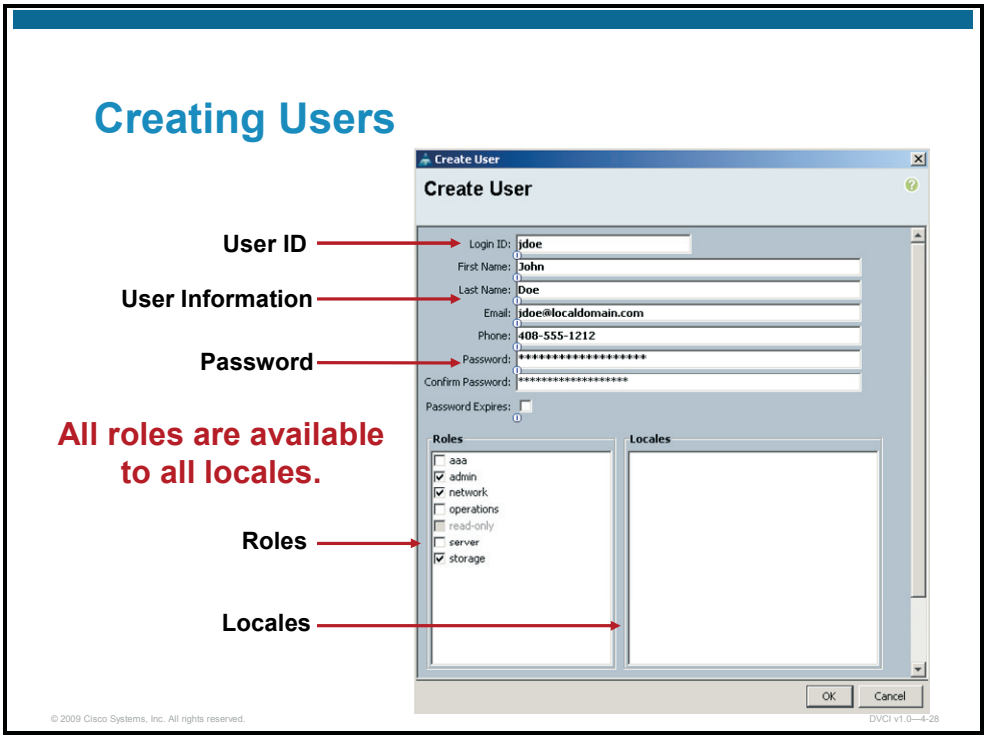

The figure shows the creation of Users through the Cisco UCS Manager GUI.

https://t.me/learningnets

Remote RBAC Configuration

This topic explains remote RBAC configuration.

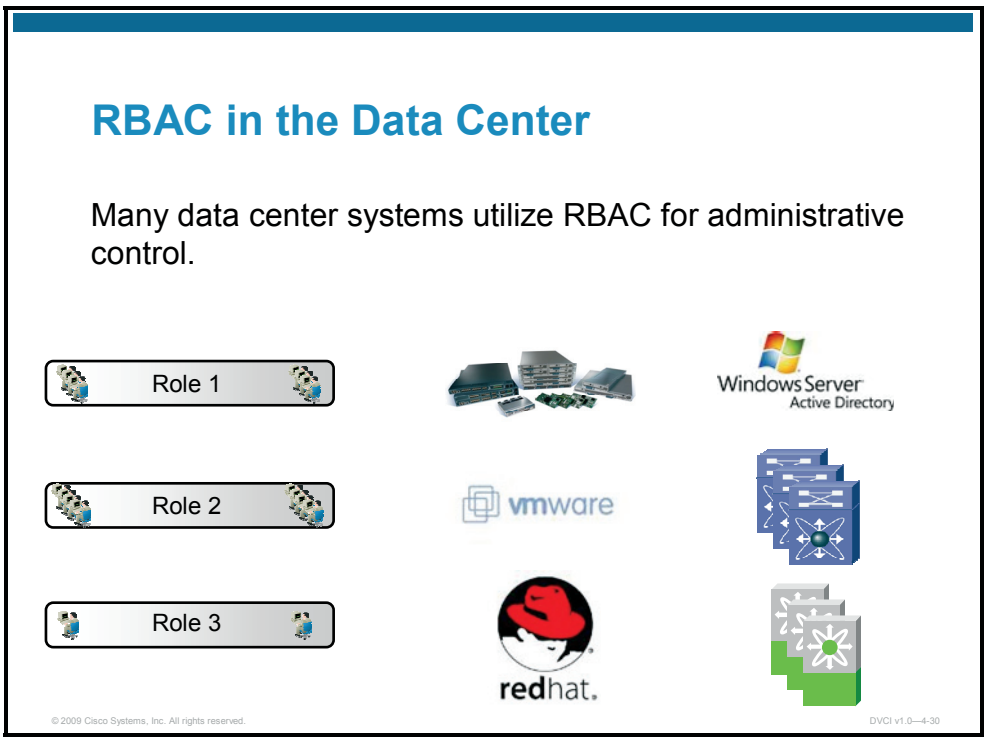

RBAC is used in many systems to provide granular access control. These systems rely on internal users and roles.

https://t.me/learningnets

Centralized Authentication

In order to simplify administration of multiple RBAC schemes, various systems exist to centralize Authentication.

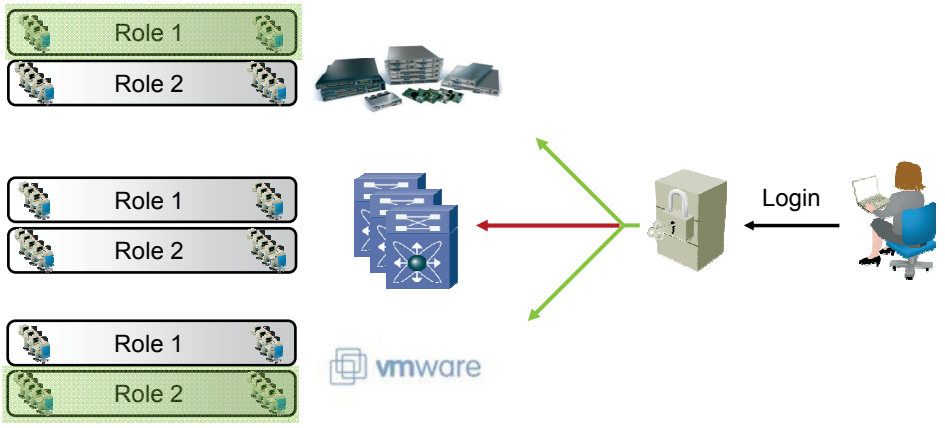

© 2009 Cisco Systems, Inc. All rights reserved.

RBAC is used in many systems to provide granular access control. These systems rely on internal users and roles.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

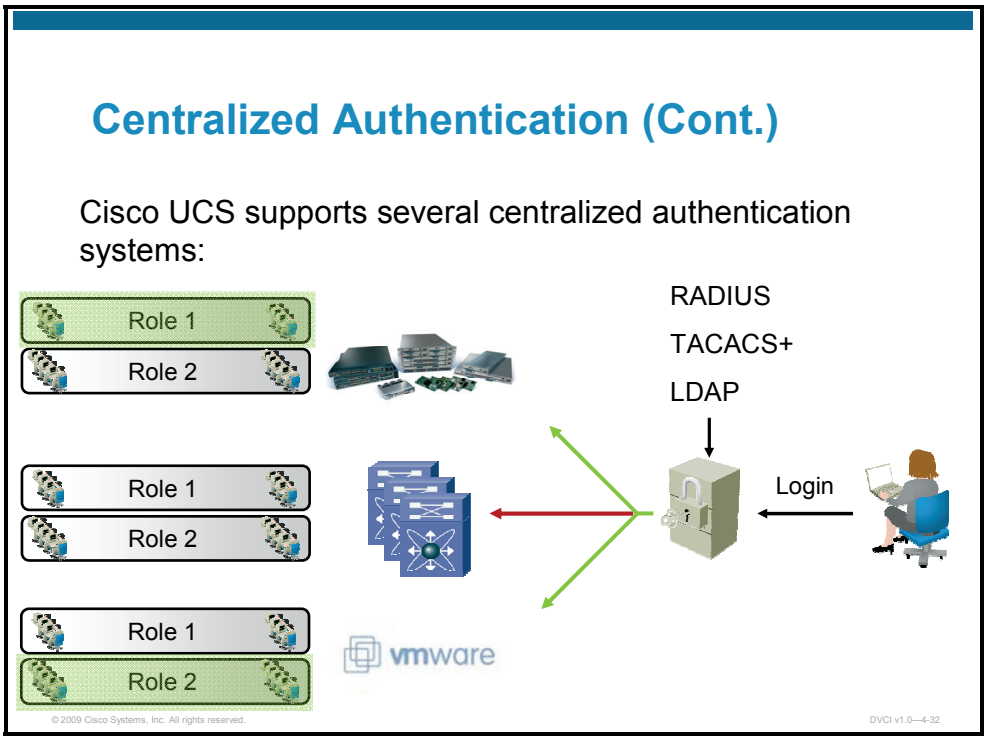

Cisco UCS supports RADIUS, TACACS+, and LDAP for centralized RBAC authentication. These are common authentication systems in today's data centers.

https://t.me/learningnets

Centralized Authentication Operation

Steps in centralized authentication:

- User logs into the device (Cisco UCS in this example)
- Device passes credentials to authentication server for approval
- Authentication server grants or denies based on credentials:
  - Passes approved role(s) to requesting device
- Requesting device grants or denies access based on server reply

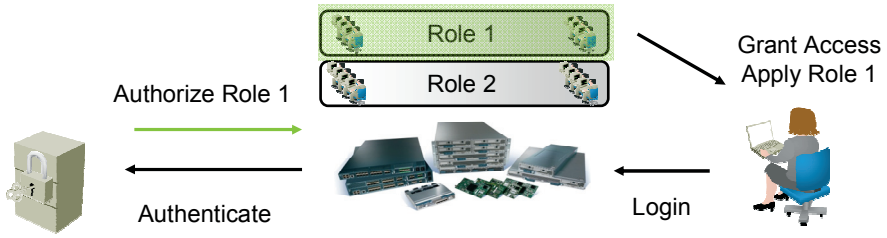

© 2009 Cisco Systems, Inc. All rights reserved.

Using centralized authentication for RBAC requires that the device be configured to pass login requests to a central server or server group. These credentials are checked against a user database on the authentication server and permissions are granted based on matches.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

AAA Configuration

This topic describes the configuration of AAA features and options.

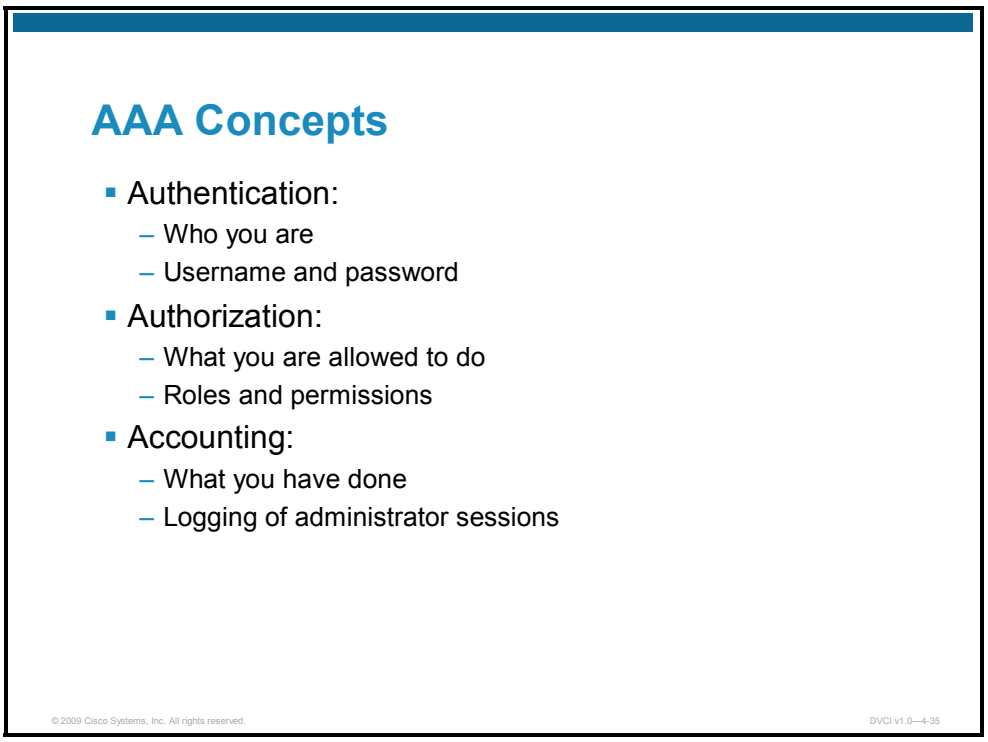

Authentication

Authentication is the process of verifying the identity of the person or device accessing the Cisco UCS Manager . Verification is based on the user ID and password combination.

Authorization

Authorization is the process of associating an authenticated user to a locale and one or more roles. You can configure role-based authorization by assigning user roles locally, or by using remote AAA servers.

Accounting

The accounting function tracks and maintains a log of every management session used to access the Cisco UCS Manager .

https://t.me/learningnets

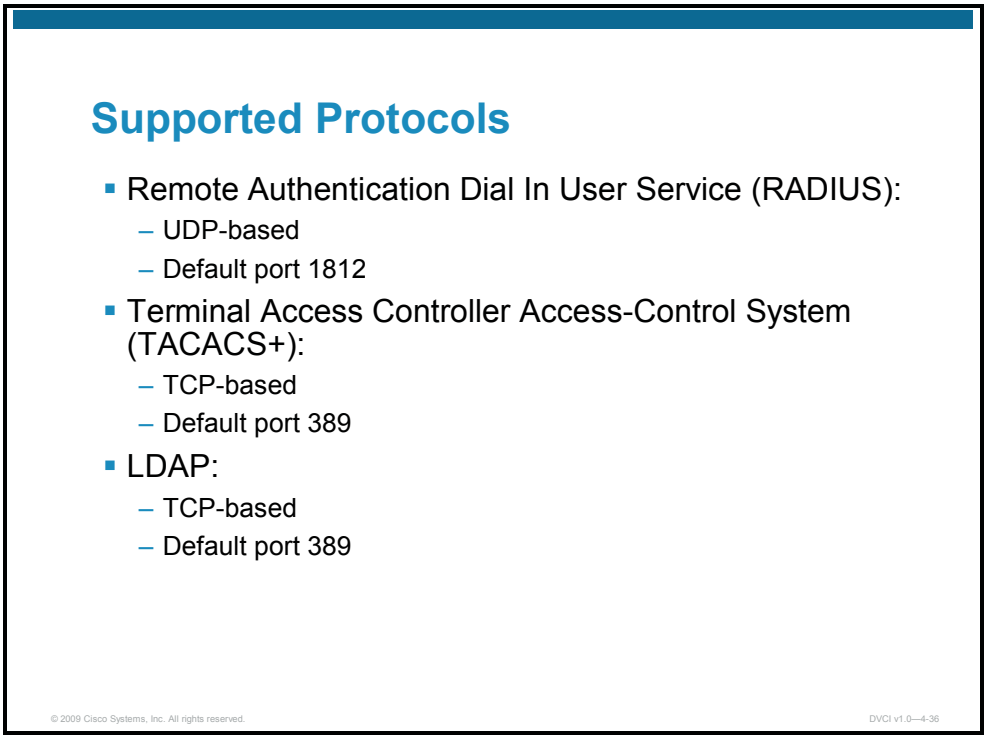

TACACS+ servers are based on RFC 1492. TACACS+ software was developed at Cisco and is freely available. TACACS+ was defined to provide AAA functionality for Cisco routers, network access servers, and other networked devices with one or more centralized servers. TACACS+ predates RADIUS and evolved from two earlier versions (TACACS and XTACACS) to provide a flexible, secure, and powerful modular AAA framework. With TACACS+, each AAA service can be tied in to its own database or performed independently by different centralized servers, allowing authorization to be performed without authentication (a feature missing in RADIUS). TACACS+ provides the option to encrypt the whole payload in a security exchange and thereby provide a higher level of data confidentiality than the shared secret password offered by RADIUS.

RADIUS and TACACS+ both provide AAA solutions. Each has its own advantages. RADIUS advantages include the following:

- It is a more commonly deployed AAA service that uses a fully open protocol. RADIUS source code is freely available and can be customized for a given use.
- RADIUS uses less memory and fewer CPU cycles than similarly configured TACACS+ servers. RADIUS uses efficient User Datagram Protocol (UDP) packets as a transport.
- RADIUS is a distributed client/server system. In the Cisco implementation, RADIUS clients run on Cisco MDS switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

TACACS+ advantages include the following:

- TACACS+ is designed to provide independent, modular AAA services.
- TACACS+ encrypts the entire payload and thus provides better network security than RADIUS.
- TACACS+ uses connection oriented TCP packets as a transport. TCP based TACACS provides reliable and acknowledged transport for AAA messages.

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

AAA Configuration

- Create a provider object for each protocol needed
- Change authentication method for “default” or “console”
- Console authentication affects only serial console
- Default authentication affects SSH, Telnet, GUI access
- Requires external database schema modifications
  - Cisco-specific attribute added to user account:
    - LDAP: CiscoAVPair
    - RADIUS/TACACS+: cisco-av-pair

© 2009 Cisco Systems, Inc. All rights reserved.

For each external authentication server used, a provider object is created within the Cisco UCS Manager , containing values specific to that protocol. Once an authentication provider is configured, the Fabric Interconnect can be configured to use that provider for login authentication and authorization. Two types of authentication are supported; default and console. The console method is used only for direct console logins through the serial console port on the Fabric Interconnect. The default method is used for all other login types, including the Cisco UCS Manager GUI, SSH, and Telnet access to the Fabric Interconnect.

https://t.me/learningnets

RADIUS Configuration

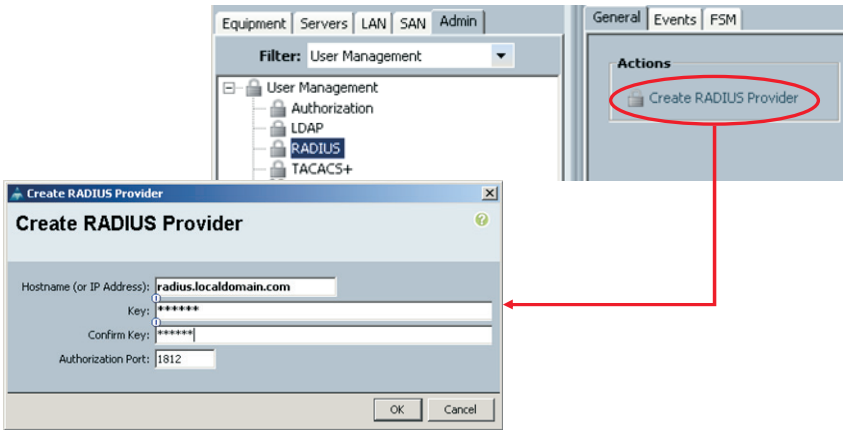

© 2009 Cisco Systems, Inc. All rights reserved.

To create a RADIUS provider, click the **Create RADIUS Provider** link in the content pane of the RADIUS configuration pane. Supply a hostname or IP address for the RADIUS server, as well as the shared secret key that was previously configured in the RADIUS server. If the port used is non-standard (not port 1812), it can be changed as well.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

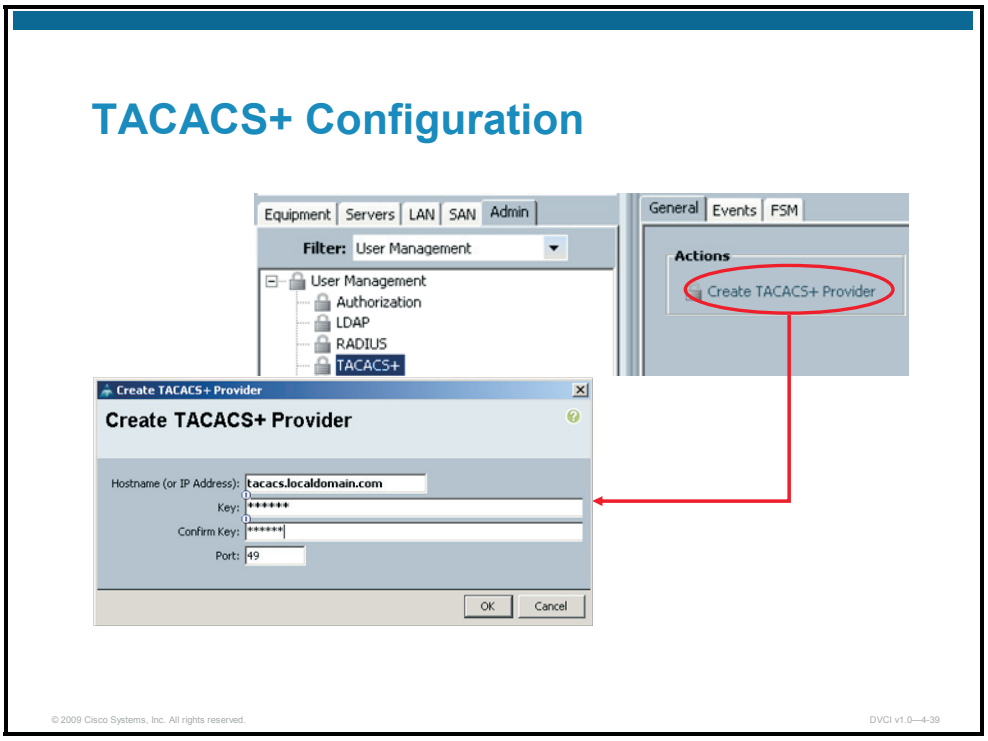

To create a TACACS+ provider, click the **Create TACACS+ Provider** link in the content pane of the TACACS+ configuration pane. Supply a hostname or IP address for the TACACS+ server, as well as the shared secret key that was previously configured in the TACACS+ server. If the port used is non-standard (not port 49), it can be changed as well.

https://t.me/learningnets

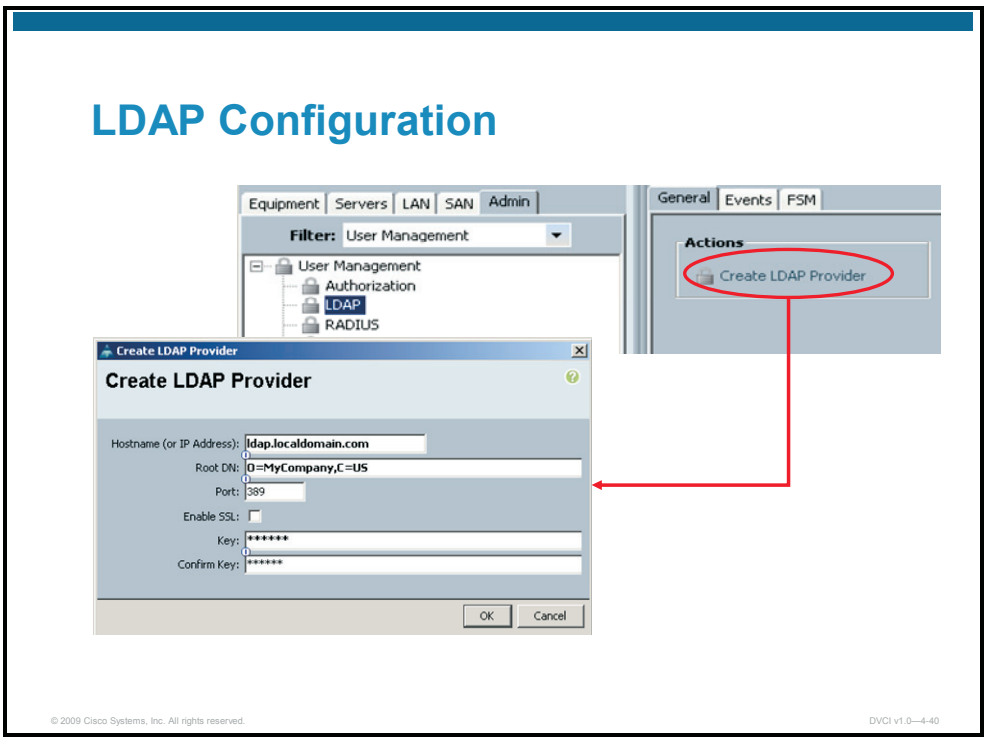

To create a LDAP provider, click the **Create LDAP Provider** link in the content pane of the LDAP configuration pane. Supply a hostname or IP address for the LDAP server, as well as the shared secret key that was previously configured in the LDAP server. A Root DN must also be specified, as required by the particular LDAP server in use. If the port used is non-standard (not port 389), it can be changed as well.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

LDAP Configuration (Cont.)

- **Base DN** sets the portion of the organizational tree that is searched for user accounts
- **Attribute** specifies the user attribute that is checked for authentication values

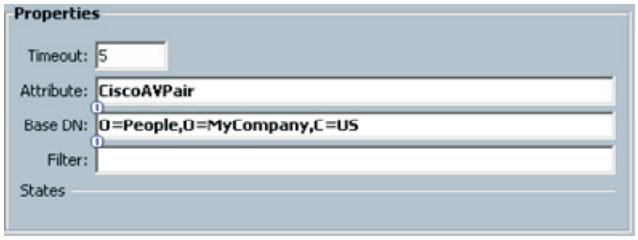

© 2009 Cisco Systems, Inc. All rights reserved.

Additionally, the LDAP provider object can be configured with additional options. The Attribute option sets the LDAP attribute checked in the user object for the role and locale information. A Base DN can be configured to identify the portion of the LDAP tree that is searched for users. Finally, a Filter can be defined to further restrict which users will be returned during searches.

https://t.me/learningnets

Authorization Configuration

- After creating desired provider configuration, set **Authorization** to use protocol
- Will fail to local authentication if remote protocol fails

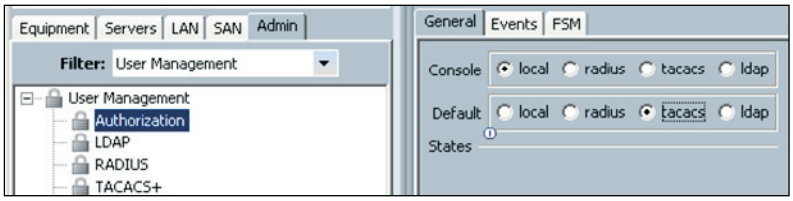

© 2009 Cisco Systems, Inc. All rights reserved.

Once the provider object(s) are created, the Cisco UCS Manager must be configured to use the provider. Select the Authorization section in the navigation pane and select the desired protocol for each of the Console and Default authentication methods. If the external authentication provider is unavailable, the Cisco UCS Manager will attempt to authenticate the user with the local user database.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Multi-Tenancy Features and Options

This topic describes multi-tenancy features and options.

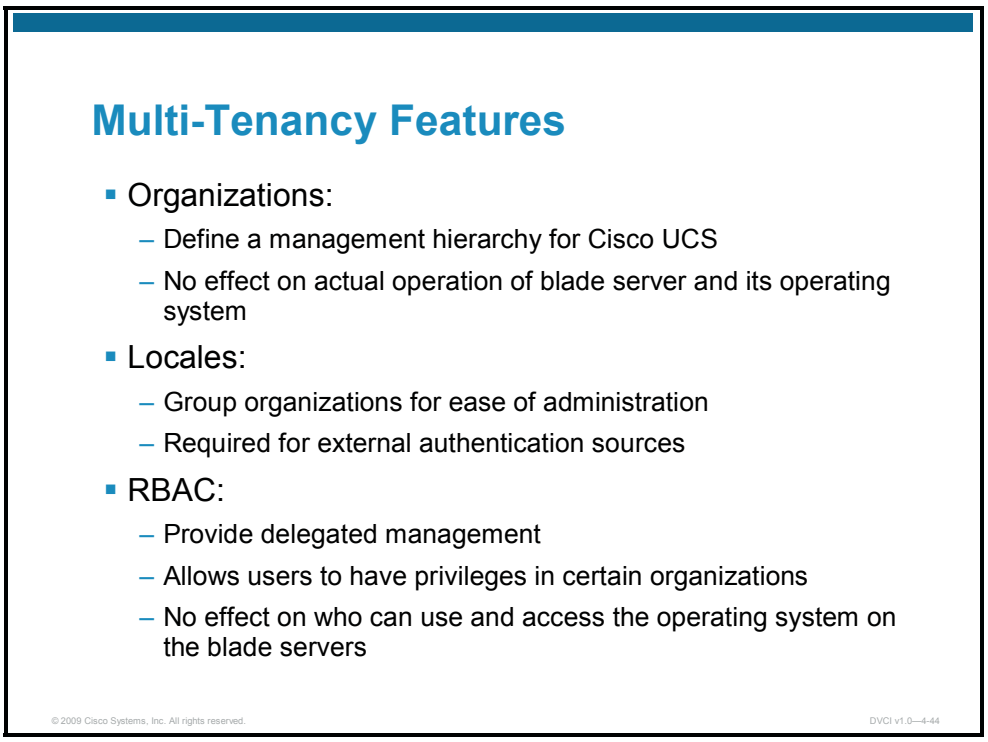

The organization structure defines a management hierarchy within the Cisco UCS implementation. The hierarchy is strictly for management of resources and logical objects inside the system; it has no effect on the actual operation of a blade server and its operating system.

The organization structure can be used together with the RBAC delegated management features. The RBAC features not only allow fine-grained privileges to be assigned to different management users, but also allow restriction of a user's privileges to a specific set of organizations.

https://t.me/learningnets

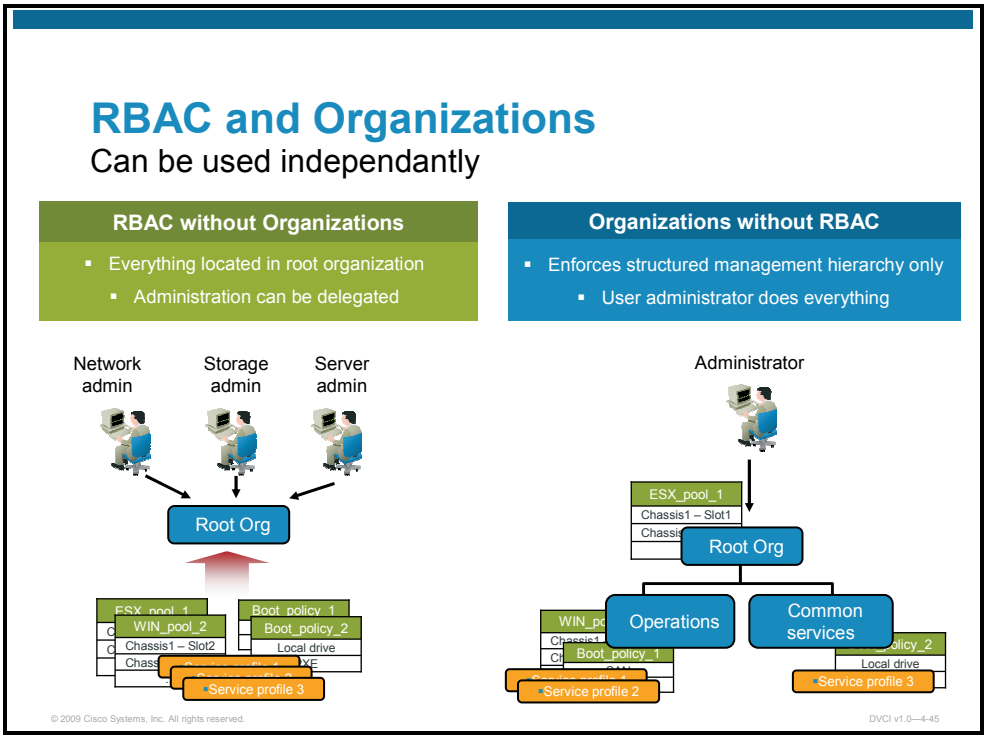

RBAC and organizations can be used independently. It is not necessary to use RBAC if you are managing a hierarchical set of organizations, or vice versa.

Using RBAC without Organizations

In case where RBAC is used without organizations – i.e. everything is in the root organization – the RBAC style of delegated management advantages 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 case where organizations are used without RBAC the organizations 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.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Compute Node Ownership

- Compute nodes are independent of the organizations
- Same compute node can be in many server pools, in many orgs
- Compute node can be associated with logical Service Profile in any org

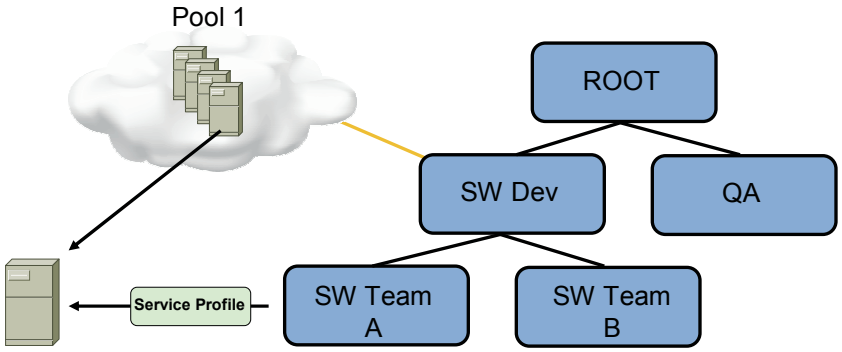

© 2009 Cisco Systems, Inc. All rights reserved.

Remember that compute nodes are never part of an organization. Any compute node can appear in as many pools in as many different organizations as you like. Any compute node can be associated with a Service Profile in any organization.

While server pools are part of an organization, they are advisory only. There is nothing to prevent associating any individual compute node to any Service Profile in any organization (bypassing the pools), regardless of how many pools the compute node might reside in.

Note how the rule is basically “first come, first served”. The first profile that successfully selects a compute node for association (through pools, or directly without pools) is the owner of that blade server until it is made available again through disassociation.

https://t.me/learningnets

Summary

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

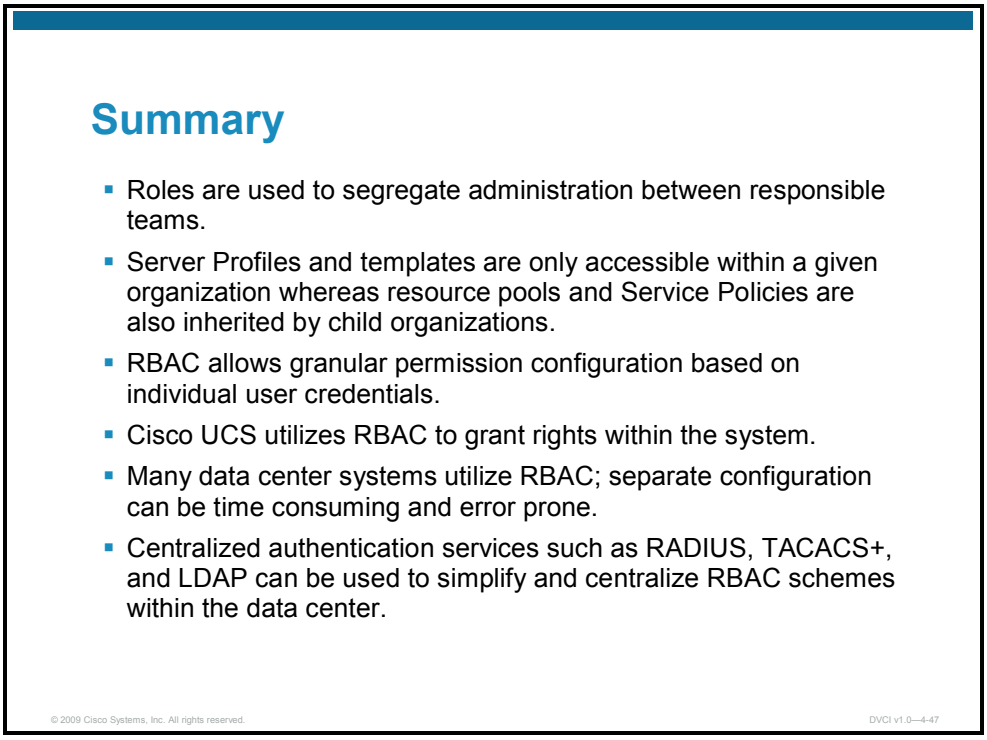

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

fModule Summary

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

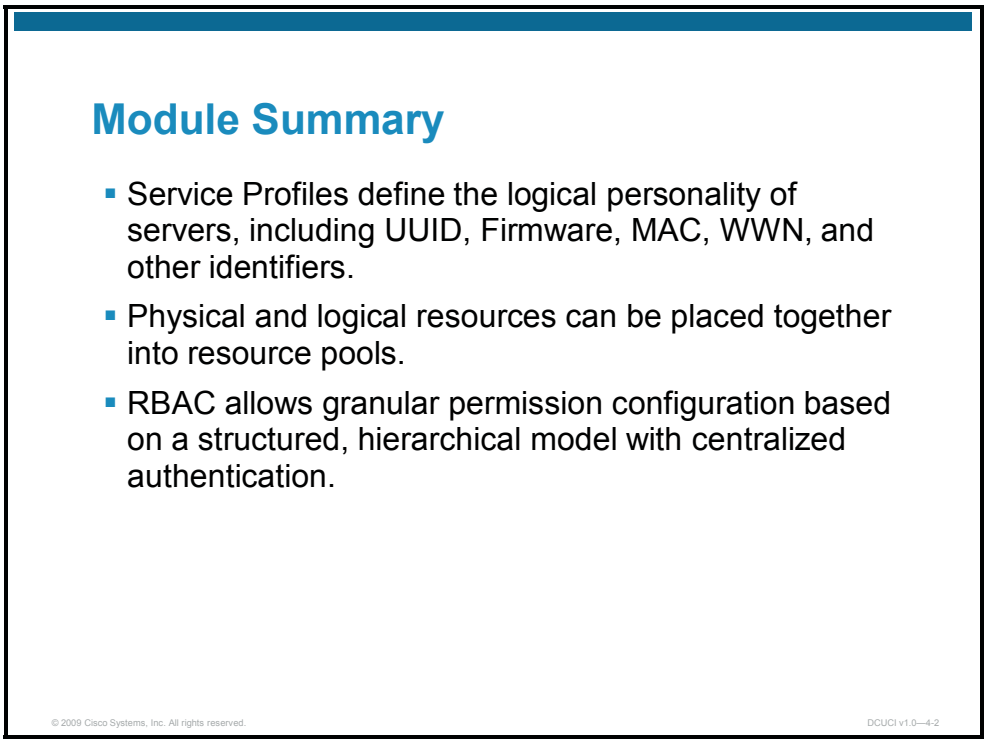

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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) What Cisco UCS feature enables mobility of server workloads? (Source: UCS Manager Fundamentals)

A) Fibre Channel over Ethernet (FCoE)
B) Hot-swap hard drive bays
C) Service Profiles
D) VSANs

Q2) A locale is a collection of what? (Source: Role-Based Access Control)

A) Roles
B) Privileges
C) Users
D) Organizations

https://t.me/learningnets

Module Self-Check Answer Key

Q1) C

Q2) D

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 5

Establishing Compute Node Connectivity

Overview

This module describes the configuration and options for establishing compute node network connectivity.

Module Objectives

Upon completing this module, you will be able to describe the requirements and capabilities of compute node connectivity. This includes being able to meet these objectives:

- Describe the configuration requirements and options for compute node LAN connectivity in Cisco UCS
- Describe the configuration requirements and options for compute node SAN connectivity in Cisco UCS

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Compute Node LAN Connectivity

Overview

In this lesson you will learn how to configure Cisco UCS for connectivity between compute nodes and Ethernet networks through the Fabric Interconnect.

Objectives

Upon completing this lesson, you will be able to describe the configuration requirements and options for compute node LAN connectivity in Cisco UCS. This includes being able to meet these objectives:

- Describe server and uplink Ethernet ports in Cisco UCS
- Describe Ethernet Host Virtualizer mode
- Describe Ethernet switching mode
- Describe the configuration and use of VLANs in the Cisco UCS environment
- Describe Virtual Network Interface Cards
- Describe port channel configuration
- Describe how Pin Groups function
- Describe LAN tools

https://t.me/learningnets

Server and Uplink Ports

This topic describes server and uplink Ethernet ports in Cisco UCS.

Uplink Ports

- Uplink ports connect the Fabric Interconnect to the upstream Ethernet cloud
- Can be configured in Ethernet Host Virtualizer mode or switching mode
- In EHV mode, uplink ports appear to be a host with multiple MAC addresses, similar to a VMware Virtual Switch
- In switching mode, uplink ports appear to be an Ethernet switch

© 2009 Cisco Systems, Inc. All rights reserved.

Uplink port configuration will be described later in this lesson. The Fabric Interconnect can be configured in EHV mode or switching mode.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Server Ports

- Server ports are not directly configured
- Server connectivity is configured exclusively through Service Profiles
- Server VLAN and traffic properties are assigned dynamically based on service profile associations
- Attributes include VLAN configuration, trunking characteristics, and redundancy settings

© 2009 Cisco Systems, Inc. All rights reserved.

Server ports are not directly configured; instead, they are dynamically adjusted depending on the Service Profiles deployed to a particular chassis. This has the advantage of reducing broadcast traffic reaching a given chassis–only those VLANs actually in use within the chassis will be allowed.

https://t.me/learningnets

Ethernet Host Virtualizer Mode

This topic describes Ethernet Host Virtualizer (EHV) mode.

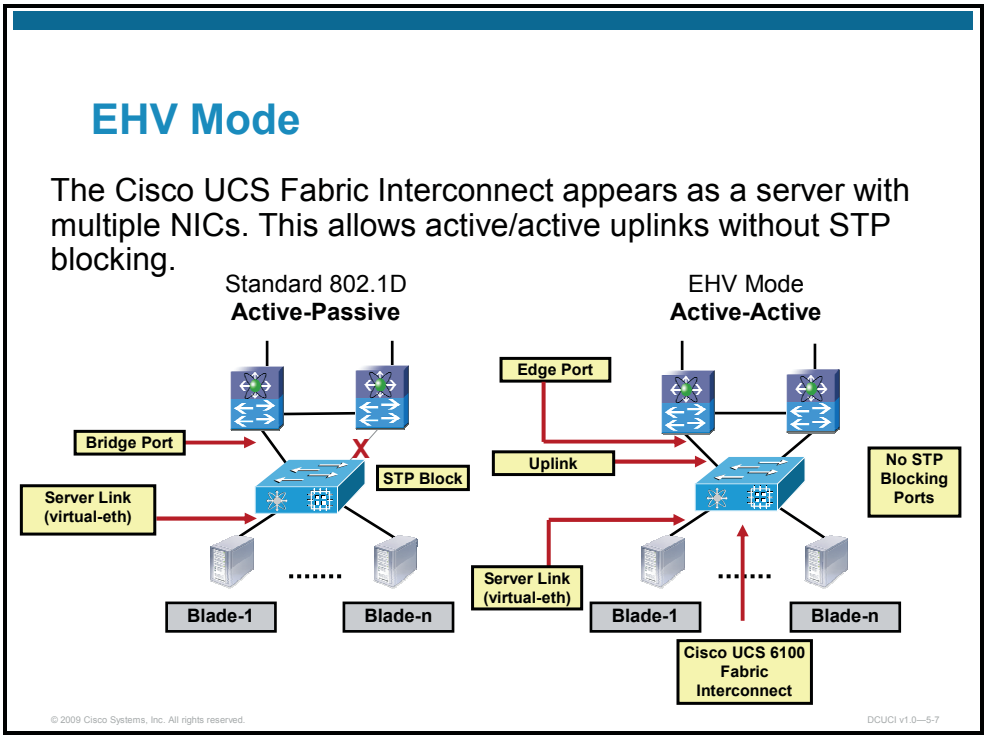

EHV appears to the external LAN as a host with many MAC addresses:

- EH-node has two types of links—server links (virtual-eth) and uplinks.
- Each uplink can be an aggregated or a single link.
- EH-node switch does not talk Spanning Tree Protocol (STP) over border links.
- Traffic forwarding among border links is prevented.
- All uplinks must connect to the Layer 2 network.
- Advantage—reduces scale of STP control plane.
- Advantage—active-active use of redundant links toward the Layer 2 network.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

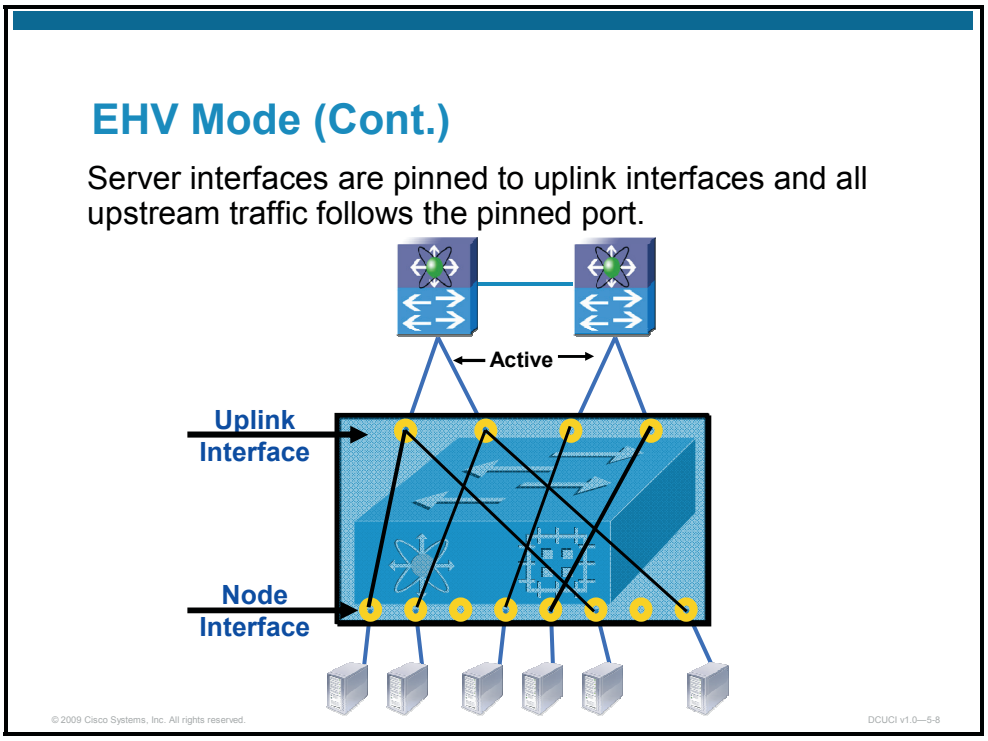

Each server link is pinned to exactly one uplink. Pinning logic load-balances blade interfaces to various uplinks:

- Server-to-server traffic is locally switched.
- Server-to-network traffic goes out on pinned uplink.
- Network-to-server unicast traffic is forwarded to server only if it arrives on pinned uplink (a Reverse Path Forwarding [RPF] check).
- Server traffic received on any uplink except pinned border link is dropped (a déjà vu check).

https://t.me/learningnets

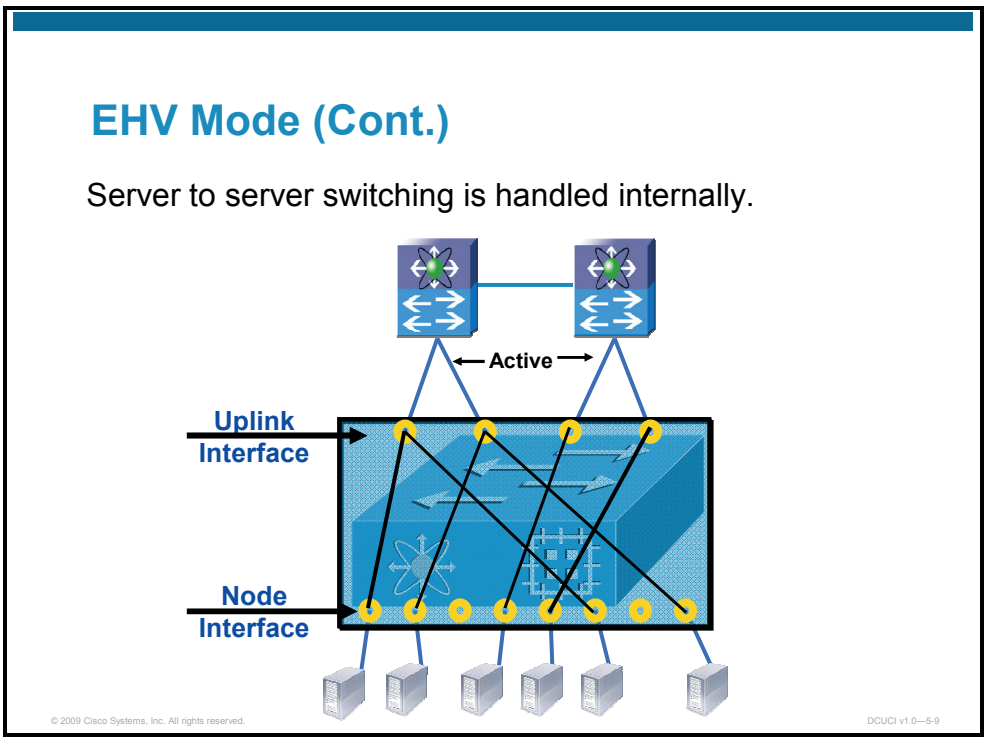

Fabric Interconnects perform local switching for server to server traffic. They learn MAC Addresses from inside the chassis only.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

EHV Forwarding Rules

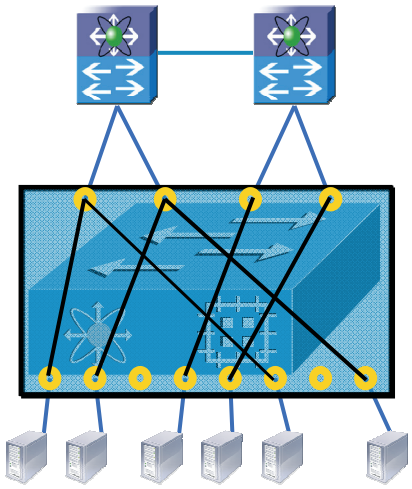

Unlearned unicast messages:

- Received on uplink interface are dropped
- Received on server interface are forwarded to the pinned port and all active server ports

© 2009 Cisco Systems, Inc. All rights reserved.

DCUCI v1.0—5-10

Traffic Type	Server Link	Border Link
Unlearned unicast	Flooded to server links and pinned uplink	Dropped
Learned unicast learned server link	Send to learned server link	Send to
Known multicast IGMP snooping	As per IGMP snooping	As per
Broadcast/unknown multicast server links	Flooded to server links and pinned uplink	Flooded to

https://t.me/learningnets

EHV Forwarding Rules (Cont.)

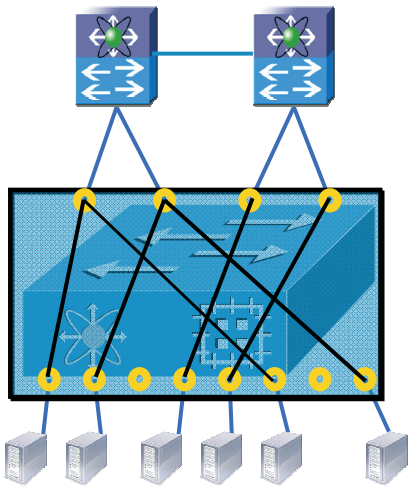

Broadcast messages:

- Received on border interface forwarded to server interfaces
- Received on server interface are flooded to the pinned port and all active server ports

© 2009 Cisco Systems, Inc. All rights reserved.

Traffic Type	Server Link	Border Link
Unlearned unicast	Flooded to server links and pinned border link	Dropped
Learned unicast learned server link	Send to learned server link	Send to
Known multicast IGMP snooping	As per IGMP snooping	As per
Broadcast/unknown multicast server links	Flooded to server links and pinned border link	Flooded to

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

EHV Learning Rules

- Learning is disabled on uplinks:
  - 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 server link goes down or is deleted:
  - Server MACs can move (in the event of re-pinning)
- Server MAC address can be statically configured.

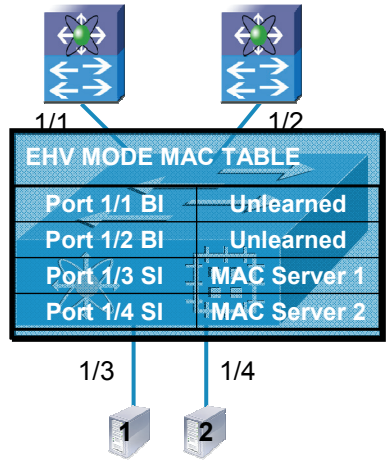

© 2009 Cisco Systems, Inc. All rights reserved.

The figure summarizes the MAC learning rules for EHV mode. In this mode, learning of MAC addresses is disabled on uplinks. On server links, learning is enabled to allow Cisco UCS to forward traffic between servers. Learned MAC addresses never age unless a server link goes down, or is deleted, or in the event that servers are re-pinned. Alternately, server MAC addresses can be statically configured.

https://t.me/learningnets

Ethernet Switching Mode

This topic describes Ethernet switching mode.

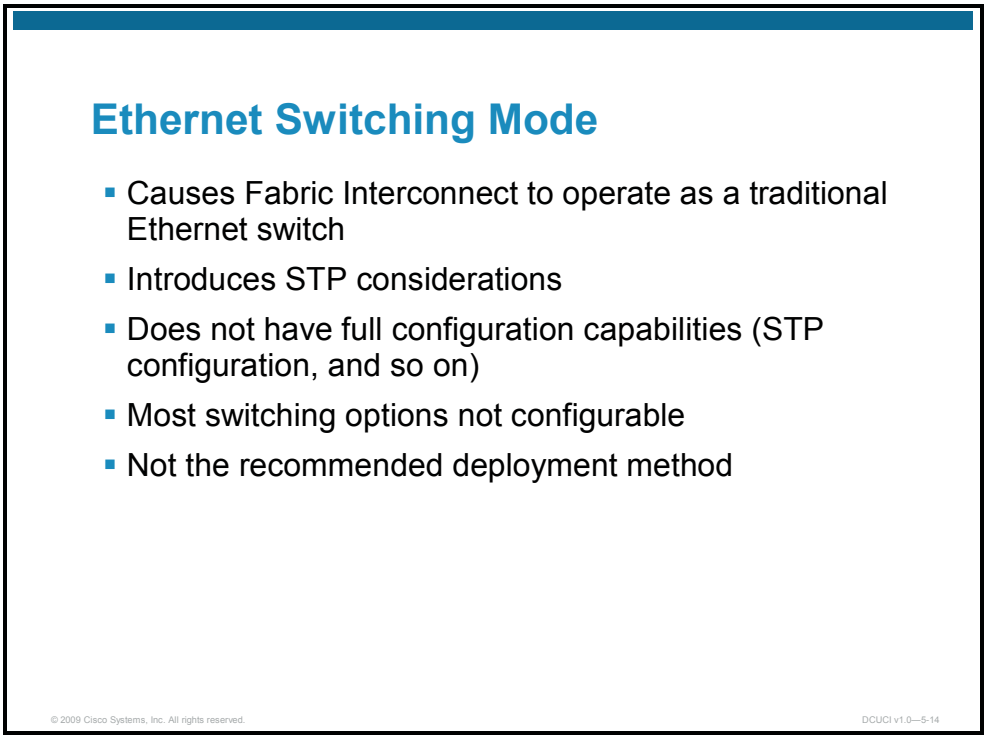

If desired, the Fabric Interconnects can be set to operate in traditional Ethernet switching mode. This introduces the need for STP to avoid broadcast loops. Despite operating in this manner, not all typical Cisco switch options are available. For this reason, this deployment method is not typically recommended.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Enabling Switching Mode

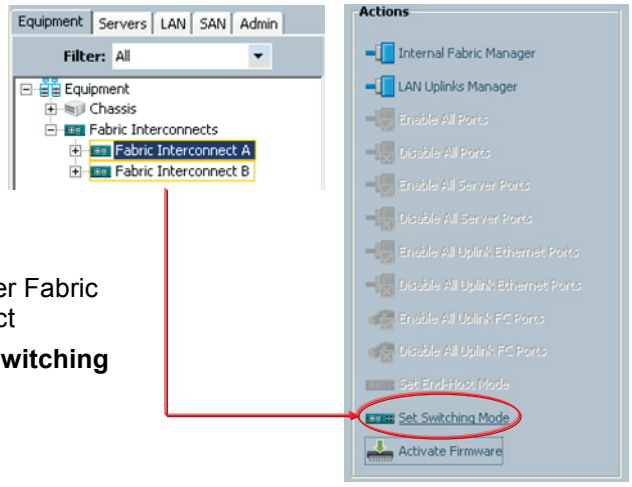

- Select either Fabric Interconnect
- Click **Set Switching Mode**

© 2009 Cisco Systems, Inc. All rights reserved.

To enable switching mode, click the **Set Switching Mode** link for either Fabric Interconnect.

https://t.me/learningnets

Enabling Switching Mode (Cont.)

- Enabling Switching Mode is disruptive
- Both cluster nodes are changed and rebooted

© 2009 Cisco Systems, Inc. All rights reserved.

Changing the Fabric Interconnect to Switching Mode is a disruptive operation. Both Fabric Interconnects will be changed and rebooted.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Describe the use of VLANs

This topic describes VLANs, as used within the Cisco UCS environment.

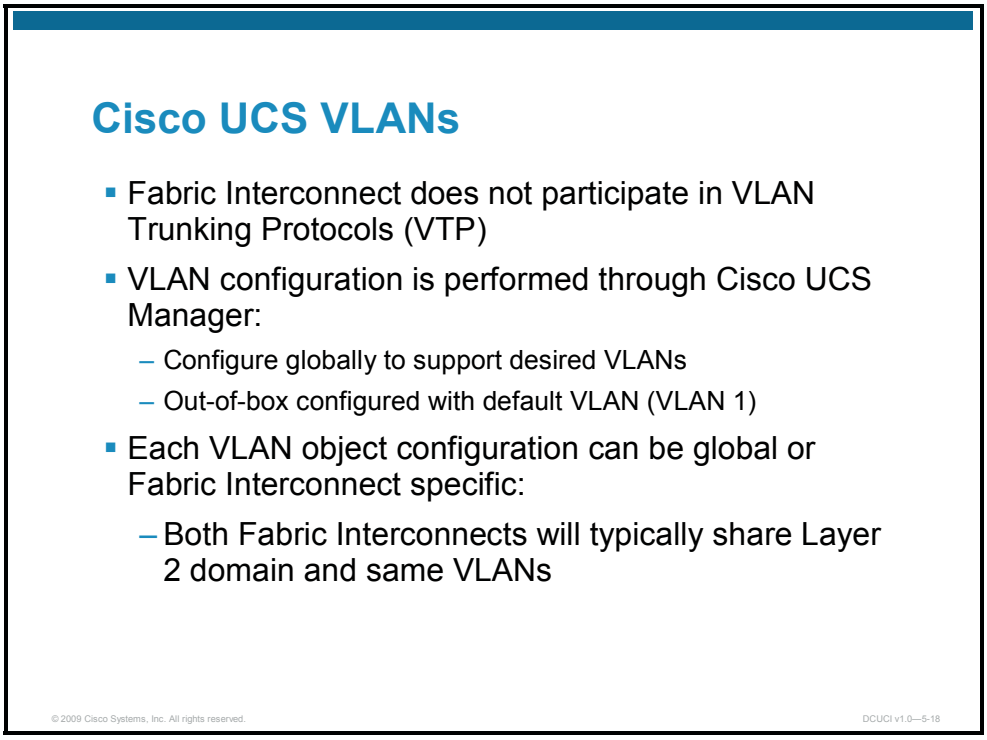

All VLAN configuration must be performed through the Cisco UCS Manager. Fabric Interconnects do not inherit any configuration through VTP. Creation of a VLAN object in the Cisco UCS Manager makes the VLAN available for use by compute nodes, depending on Service Profile configuration.

https://t.me/learningnets

Upstream Switch Configuration

- Upstream switch port should use 802.1Q VLAN trunking
- Trunk all VLANs that will be used within Cisco UCS
- Cisco UCS will automatically adjust trunking allowed lists on Fabric Interconnect ports based on VLAN objects defined

© 2009 Cisco Systems, Inc. All rights reserved.

The upstream switch should supply an 802.1Q VLAN trunk with any VLANs used by Cisco UCS compute nodes. The trunking allowed list on the uplink ports is dynamically adjusted depending on the VLAN objects defined within the Cisco UCS Manager.

When operating in EHV mode, the upstream switch ports should be configured for “Portfast” mode, since the Fabric Interconnect is not participating in STP.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Global VLAN Configuration in Cisco UCS

- Create VLAN under “LAN Cloud”:
  - Will work across entire Cisco UCS
- Optionally create Fabric Interconnect (A-specific or B-specific) configuration

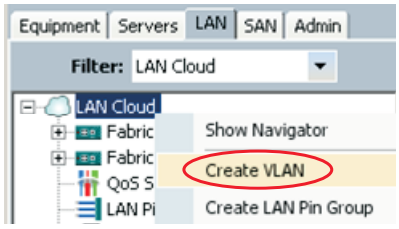

© 2009 Cisco Systems, Inc. All rights reserved.

For any VLAN that you want to support on any compute node, a VLAN object must be created in the global Cisco UCS configuration. Use the **LAN** tab on the navigation panel to create the VLAN object, as shown in the figure.

You can also create VLANs that are associated with only Fabric Interconnect A or only Fabric Interconnect B.

https://t.me/learningnets

Creating Global VLANs

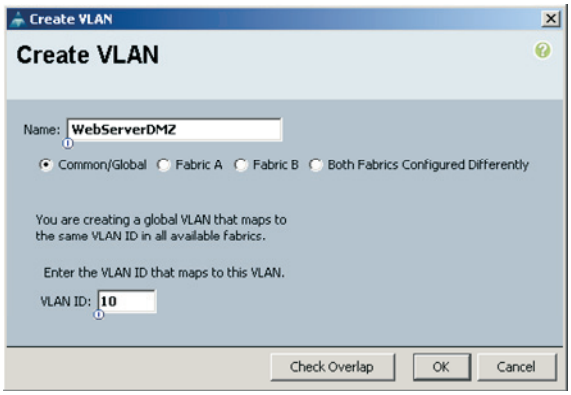

Typical implementation will use Common/Global VLANs with both Fabric Interconnects participating in the same VLANs.

© 2009 Cisco Systems, Inc. All rights reserved.

This example shows the creation of global VLAN support for a particular VLAN (VLAN 10) inside Cisco UCS. The name is simply a logical name inside the Cisco UCS Manager. By abstracting the name of the VLAN from the VLAN number, changes can be made to the VLAN architecture in one place instead of potentially many Service Profiles.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Individual VLANs

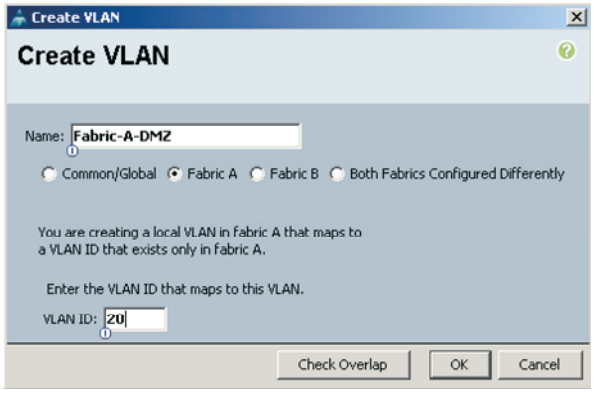

VLANs can be created that exist only on a single Fabric Interconnect.

© 2009 Cisco Systems, Inc. All rights reserved.

This example shows creation of local VLAN support for a particular VLAN (VLAN 20) on Fabric Interconnect A.

https://t.me/learningnets

Creating Differently Configured VLANs

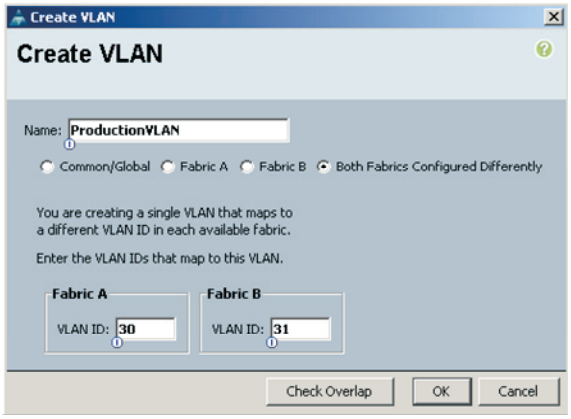

If necessary, the same VLAN label can apply to different VLANs in each Fabric Interconnect.

© 2009 Cisco Systems, Inc. All rights reserved.

In some cases, it may be necessary to have the same VLAN label apply to different VLANs, depending on the Fabric Interconnect. For example, if the same network is accessed through two different VLAN numbers in different upstream switches, this feature allows the same label to map to the proper VLAN.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Virtual NICs

This topic describes Virtual Network Interface Cards (vNICs).

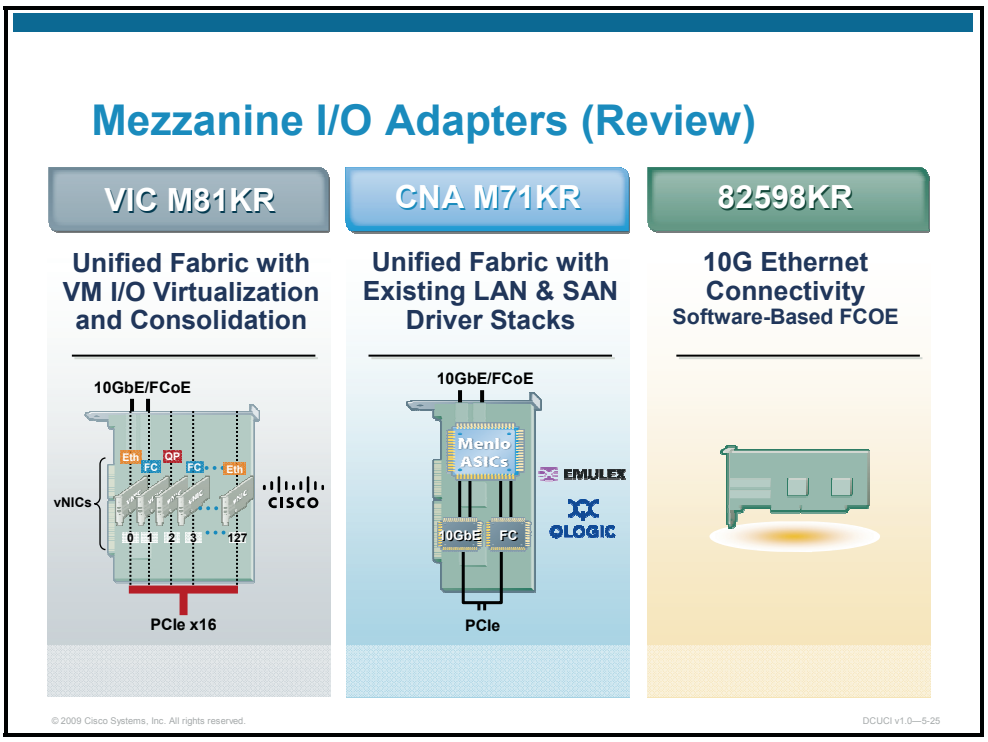

Mezzanine Cards

The types of adapters, or mezzanine cards, available in a Cisco UCS blade server include:

- Cisco UCS VIC M81KR: A virtual interface card
- Cisco UCS CNA M71KR-E: An Emulex converged network adapter
- Cisco UCS CNA M71KR-Q: A QLogic converged network adapter
- Cisco UCS 82598KR-CI: A 10-GE network adapter

Cisco UCS VIC M81KR

The Cisco UCS VIC M81KR supports multiple Ethernet virtual network interface cards (vNICs) or Fibre Channel Host Bus Adapters (vHBAs), running at 500 KIOPs in both initiator and target mode. This card supports kernel/hypervisor bypass. The Cisco UCS VIC M81KR is a converged network adapter (CNA) with dual 10-GE ports and dual Fibre Channel ports to the backplane. The card provides failover between the redundant Ethernet links. No multipathing software is required on the host operating system for this functionality.

Cisco UCS CNA M71KR

The Cisco UCS CNA M71KR is a converged network adapter (CNA) with two host-side 10-GE ports and two Fibre Channel ports to the backplane. The two network ports can run either native Ethernet or FCoE protocols and can be configured for failover. This failover is performed by the FCoE ASIC and does not require multipathing software on the host operating system:

https://t.me/learningnets

The Cisco UCS CNA M71KR ASIC is a Cisco-designed multiplexor and FCoE protocol offload engine with a 350 MHz 24K MIPS processor. There are two versions of this card (Cisco UCS CNA M71KR-E Emulex Converged Network Adapter, and Cisco UCS CNA M71KR-Q QLogic Converged Network Adapter), thereby supporting existing proven driver stacks to the customer. Both cards use an Intel 82598 10GE Ethernet chipset.

Cisco UCS 82598KR-CI

The Cisco UCS 82598KR-CI Network Adapter is made by Intel. It has two 10-GE ports to the backplane which run native Ethernet. operating system-implemented FCoE may be run over the Cisco UCS 82598KR-CI to provide “free SAN access” to the host. The card does not provide failover functionality.

Comparison

- The Cisco UCS VIC M81KR provides enhanced virtualization. This adapter supports mulitple vNICs presented to the ESX server which can then be used as physical NICs for virtual machines (VMs).
- The Cisco UCS CNA M71KR adapter provides increased scalability. The FCoE protocol encapsulation is offloaded from the host and performed in hardware on the Cisco UCS CNA M71KR.
- The Cisco UCS 82598KR-CI provides standard 10-GE connectivity to servers not needing Fibre Channel protocol support, or servers providing Fibre Channel protocol and FCoE encapsulation in software.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

vNIC

- For physical mezzanine adapter (CNA or 10GE):
  - Create one vNIC in profile for each port that you want to make accessible
  - Maximum of two vNIC per adapter card; one to Fabric Interconnect A, and one to Interconnect B
  - Failover is only supported on CNA cards

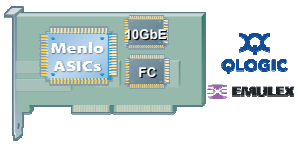

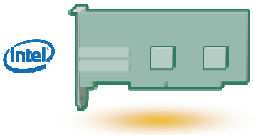

© 2009 Cisco Systems, Inc. All rights reserved.

For adapters with only physical NICs (Cisco UCS CNA M71KR, Cisco UCS 82598KR-CI), you must create a vNIC for each network interface card (NIC) you want to make usable on the network within Cisco UCS. The vNIC has a switch setting and a failover flag. For Cisco UCS 82598KR-CI, you must match the physical setting (first adapter goes to Switch A, second to Switch B), and you can not choose failover. For Cisco UCS CNA M71KR, each vNIC is associated with a particular switch, but you can enable failover.

https://t.me/learningnets

vNIC with Virtual Interface Card

- For Virtual Interface Card (VIC) mezzanine adapter:
  - Create multiple vNICs if desired
  - Number of vNICs supported depends on cabling topology
  - Each is presented as physical NIC to the operating system on the blade server
  - Choose Fabric Interconnect A or B with failover
  - Special feature to bypass VMware vSphere virtual switch and present (N-1) of these directly to VMware guest OSes

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS VIC M81KR supports NIC virtualization for either a single operating system or for VMware vSphere. Inside a single operating system, each Cisco UCS VIC M81KR vNIC is presented as a physical adapter. For VMware, a special feature allows Cisco UCS VIC M81KR vNICs to be presented directly to the guest operating system, bypassing the VMware virtual switch layer. This feature is also called "Hypervisor Bypass". This approach improves efficiency and also allows Cisco UCS to reconfigure network infrastructure if virtual machines migrate between vSphere instances on different compute nodes.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Virtual Interface Card Limitations

- Number of possible adapters depends on uplinks between I/O Module and Fabric Interconnect
- 128 virtual adapters is the architectural limit of the adapter chipset
- Each Fabric Interconnect port supports 128 VNTags (the mechanism used to support virtual interfaces)
- Some VNTags are consumed by chassis management functions

With 1 I/O Module, 8 Blades		With 2 I/O Modules, 8 Blades	
Uplinks to Fabric Interconnect	Virtual Interfaces per Adapter	Uplinks to Fabric Interconnect	Virtual Interfaces per Adapter
1	13	1	26
2	28	2	56
4	58	4	116

© 2009 Cisco Systems, Inc. All rights reserved.

The number of virtual interfaces supported on an adapter depends on the number of uplinks between the I/O Module (IOM) and the Fabric Interconnect, as well as the number of interfaces in use on other adapters sharing the same uplinks. The numbers shown in the table reflect the worst case scenario, where all nodes are attempting to use virtual interfaces.

The Cisco Virtual Interface Card mezzanine adapter is an example of an interface card that can support the creation of multiple virtualized hardware interface ports.

These limits apply to any combination of vNICs and vHBAs.

https://t.me/learningnets

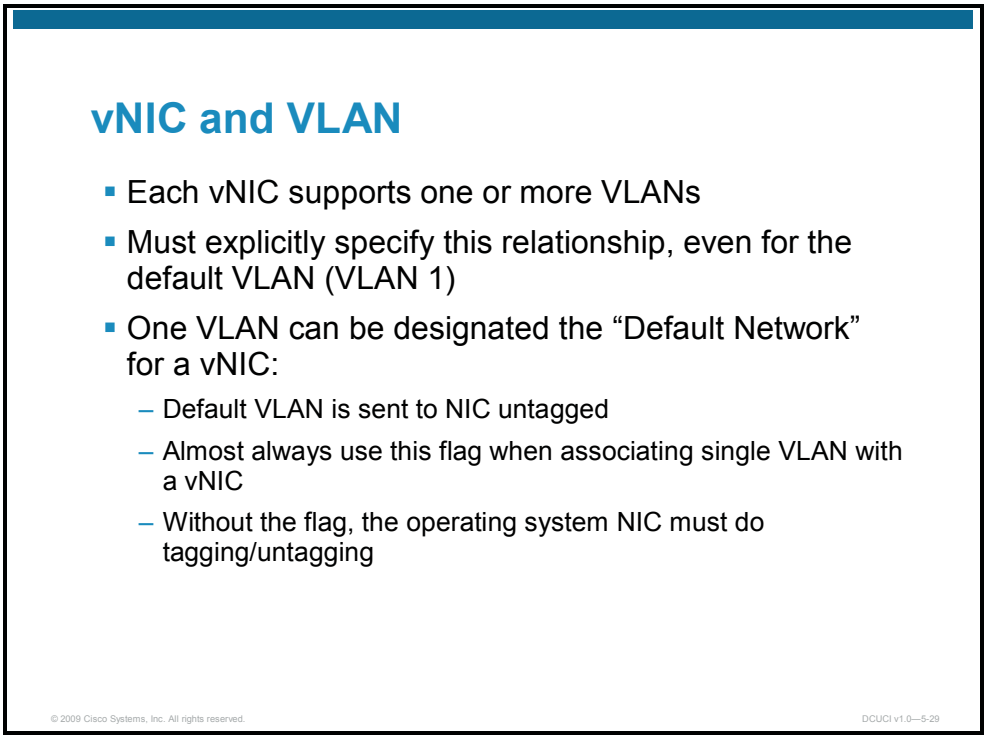

Each vNIC must be associated with one or more VLANs. Each VLAN must be configured before it is configured into a vNIC. The default VLAN (VLAN 1) must still be specifically associated with a vNIC if any default network traffic is to reach the adapter on the blade server that is associated with the profile containing the vNIC.

A typical simple configuration is a vNIC that supports only the default VLAN.

A flag that is most often associated with the default VLAN can be used to designate a VLAN as the “Default Network” for a particular vNIC. This flag indicates that traffic on that VLAN will come through to the NIC untagged; in other words, the NIC in the operating system can remain “VLAN-unaware”.

If a vNIC supports VLANs that are not the “Default Network” for that particular vNIC, traffic for those VLANs will come through to the NIC with VLAN tagging intact. This NIC must then be configured in its operating system as VLAN-aware.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Port Channels

This topic describes the link aggregation methods used by the Fabric Interconnect on uplink Ethernet ports.

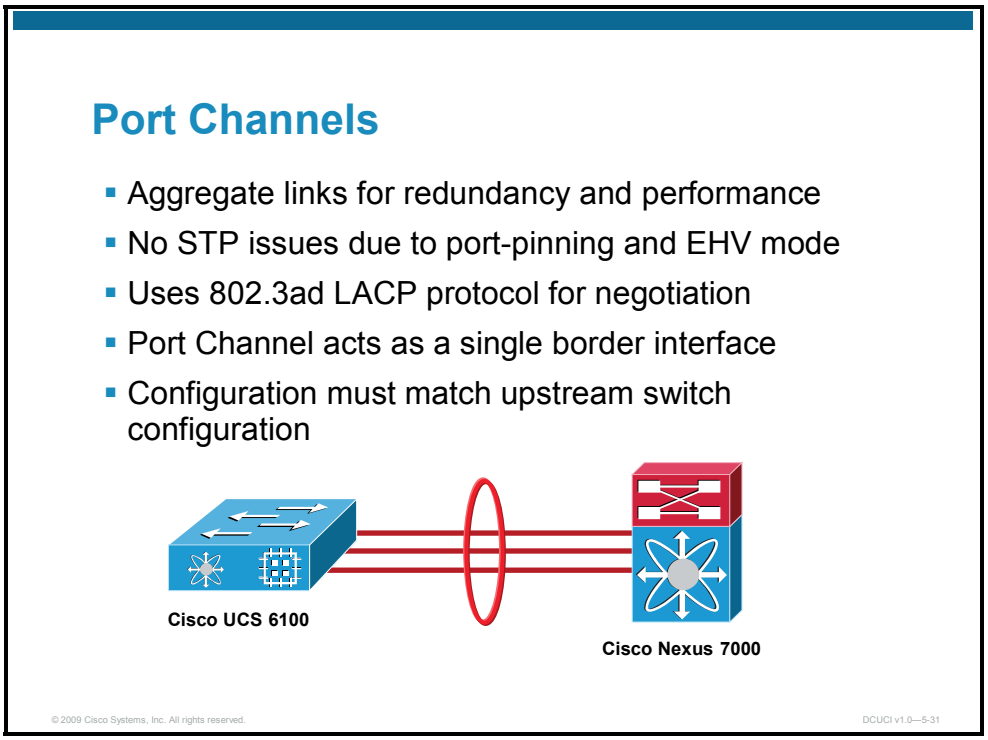

You can configure port channels on the uplink ports from the Cisco UCS Fabric Interconnect to the upstream Ethernet switch.

Only 802.3ad LACP is supported for negotiating port channels with the upstream switch. The port channel will appear to the Fabric Interconnect as a single border interface, and vNICs can be pinned to the port channel instead of individual links. The upstream switch must have LACP enabled on the port channel links to negotiate the aggregate.

https://t.me/learningnets

Creating Port Channels

- Always Fabric Interconnect specific:

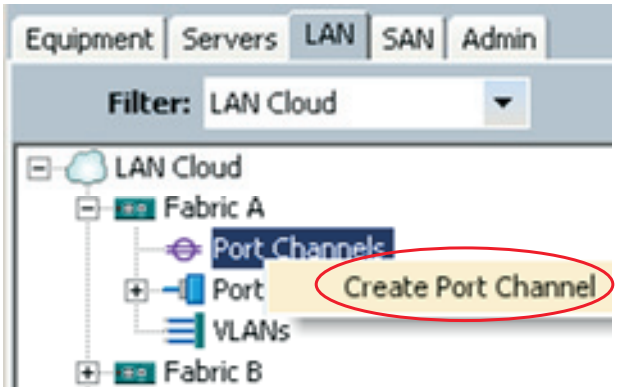

© 2009 Cisco Systems, Inc. All rights reserved.

Port Channel creation is always specific to a particular Fabric Interconnect.

https://t.me/learningnets

Creating Port Channels (Cont.)

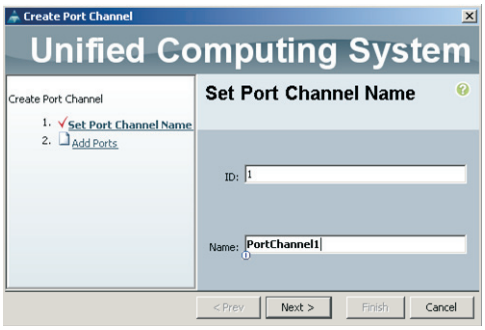

- Note that all ports are listed, even ports configured as server ports
- Selecting a server port will cause a warning message

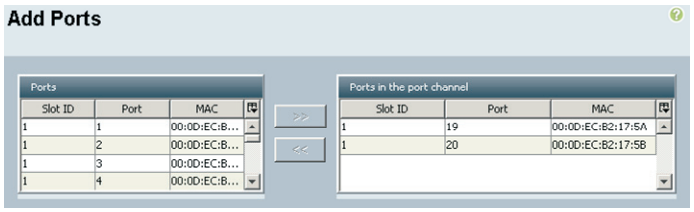

© 2009 Cisco Systems, Inc. All rights reserved.

In this example, a port channel has been created with ID 1. A matching configuration must be configured at the upstream switch that is hosting the other side of the uplinks.

https://t.me/learningnets

Pin Groups

This topic describes how Pin Groups are used to organize or segregate Ethernet traffic on uplink Ethernet links.

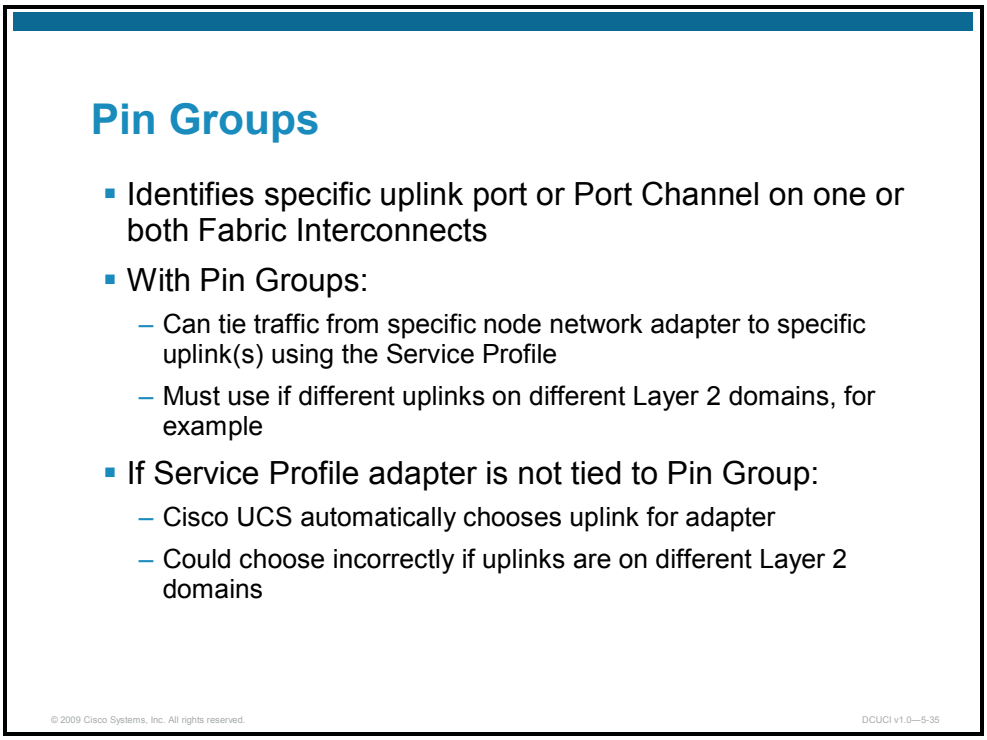

Pin Groups are a mechanism for pinning traffic from a particular blade server adapter (through the specification in its Service Profile) to a particular individual port or Port Channel on each switch.

Without Pin Groups (or even if you create them but do not use them in a Service Profile), the Cisco UCS automatically chooses an uplink port or Port Channel for the adapter on the node associated with each profile.

If your environment uses some concept of a “correct” uplink for a particular node (for example, if the uplinks are corrected to different Layer 2 domains), then you must create Pin Groups and tie them to your Service Profiles.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Pin Groups

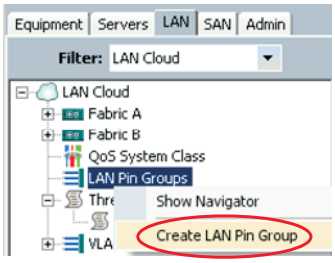

- Pin Groups point to a single link or port channel
- Each Fabric Interconnect can use different target types

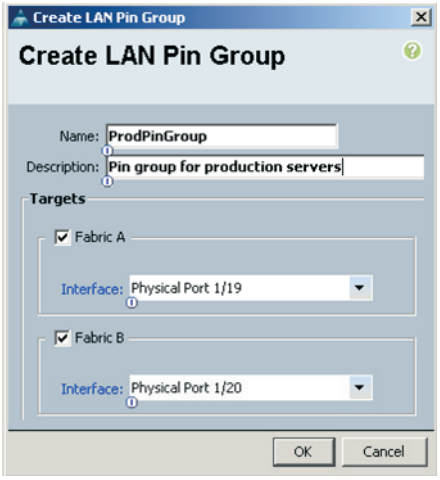

© 2009 Cisco Systems, Inc. All rights reserved.

Pin Groups are defined at the LAN Cloud level. A Pin Group is a named object in the Cisco UCS Manager that refers to a single port or Port Channel on Fabric Interconnect A, and/or a single port or Port Channel on Fabric Interconnect B.

By itself, the Pin Group does nothing. However, when it is set as an adapter property inside a Service Profile, all traffic for the adapter on the blade server that is associated with the profile will be physically tied to the Pin Group.

https://t.me/learningnets

Selecting Pin Groups

- Pin Group selection is a function of the vNIC configuration
- While creating vNIC using expert Service Profile Wizard, Pin Groups can be selected or created

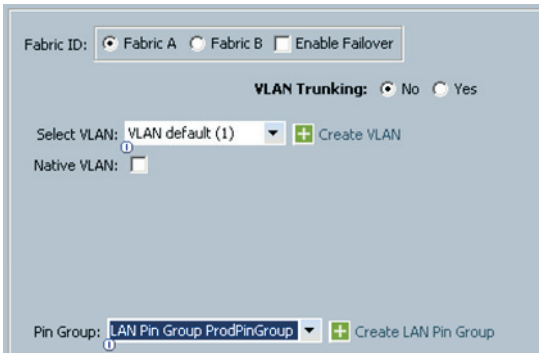

© 2009 Cisco Systems, Inc. All rights reserved.

During the creation of a vNIC in the expert Service Profile Wizard, the desired Pin Group can be selected. Alternately, the properties of the vNIC can be edited later to select the Pin Group.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

LAN Tools

This topic describes additional LAN management tools available in the Cisco UCS Manager.

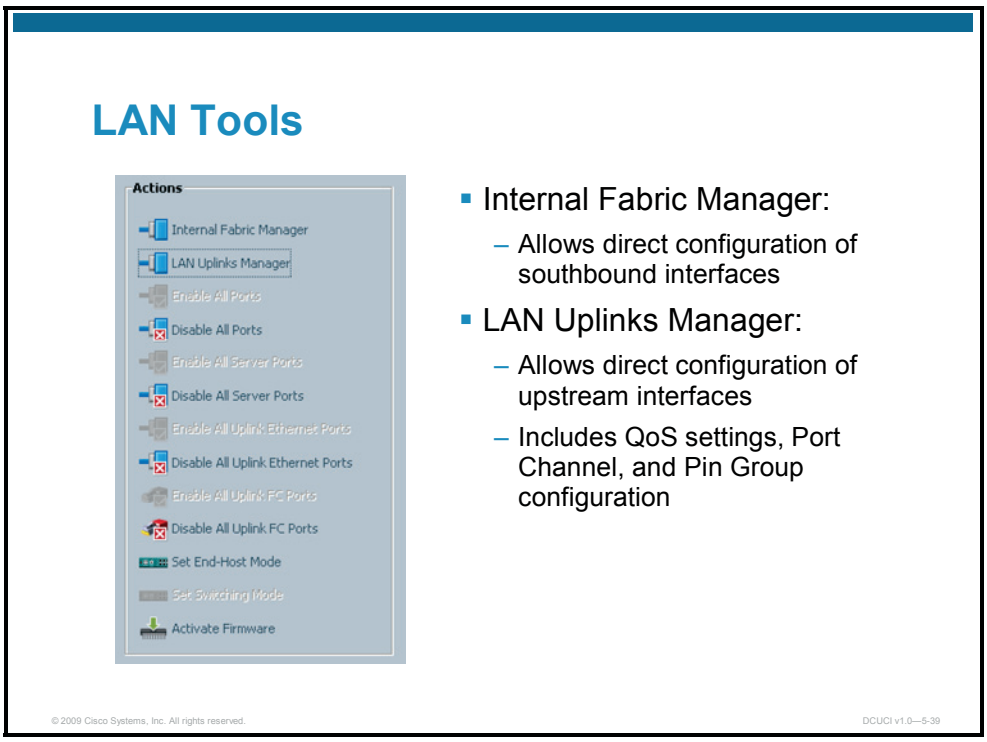

The Cisco UCS Manager provides several methods to manage the LAN connections within the Fabric Interconnect. Each port can be managed individually, or the Internal Fabric Manager and LAN Uplinks Manager tools can be started to modify multiple ports simultaneously.

https://t.me/learningnets

Internal Fabric Manager

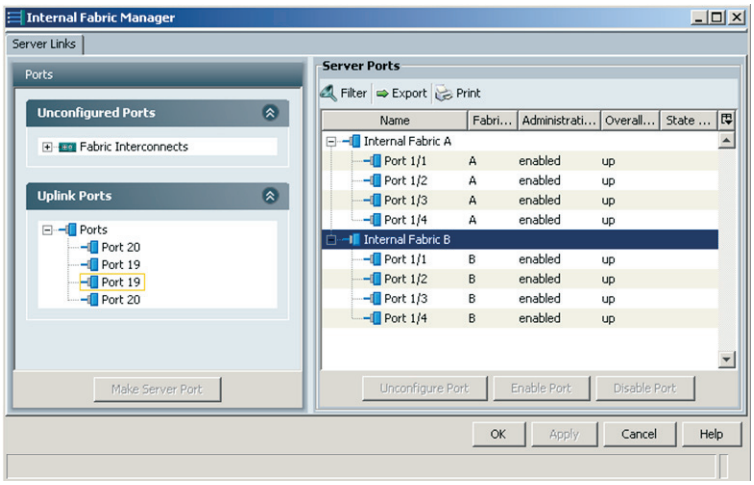

© 2009 Cisco Systems, Inc. All rights reserved.

The Internal Fabric Manager allows configuration of Server Ports connected to the chassis I/O Modules. Un-configured ports can be configured as Server Ports if desired.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

LAN Uplinks Manager

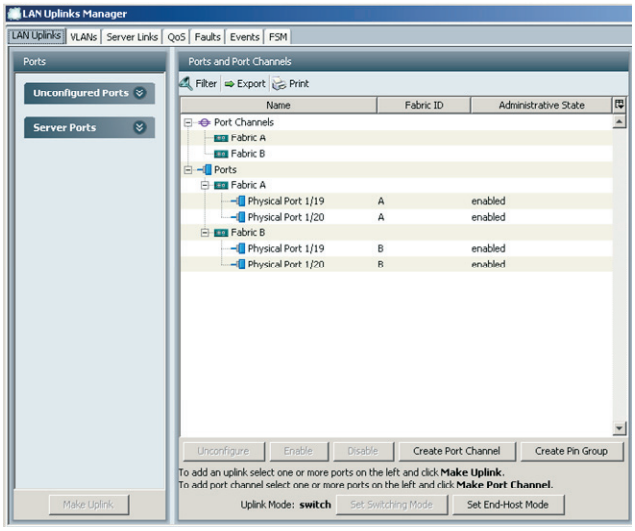

© 2009 Cisco Systems, Inc. All rights reserved.

The LAN Uplinks Manager provides similar functionality for the upstream connections from the Fabric Interconnects. Multiple tabs contain additional configuration options, such as VLAN and QoS configuration, as well as tabs containing fault and status information about the uplinks.

https://t.me/learningnets

LAN Uplinks Manager QoS

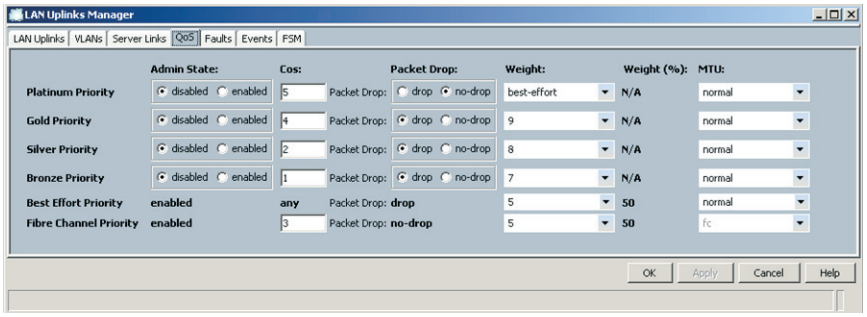

- QoS settings are available in LAN Uplinks Manager

© 2009 Cisco Systems, Inc. All rights reserved.

The QoS tab in the LAN Uplinks Manager allows for QoS configuration and mapping of named priority values to specific CoS tags. The flow control behavior of the CoS values can also be configured in this screen. The values listed here are the default values.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

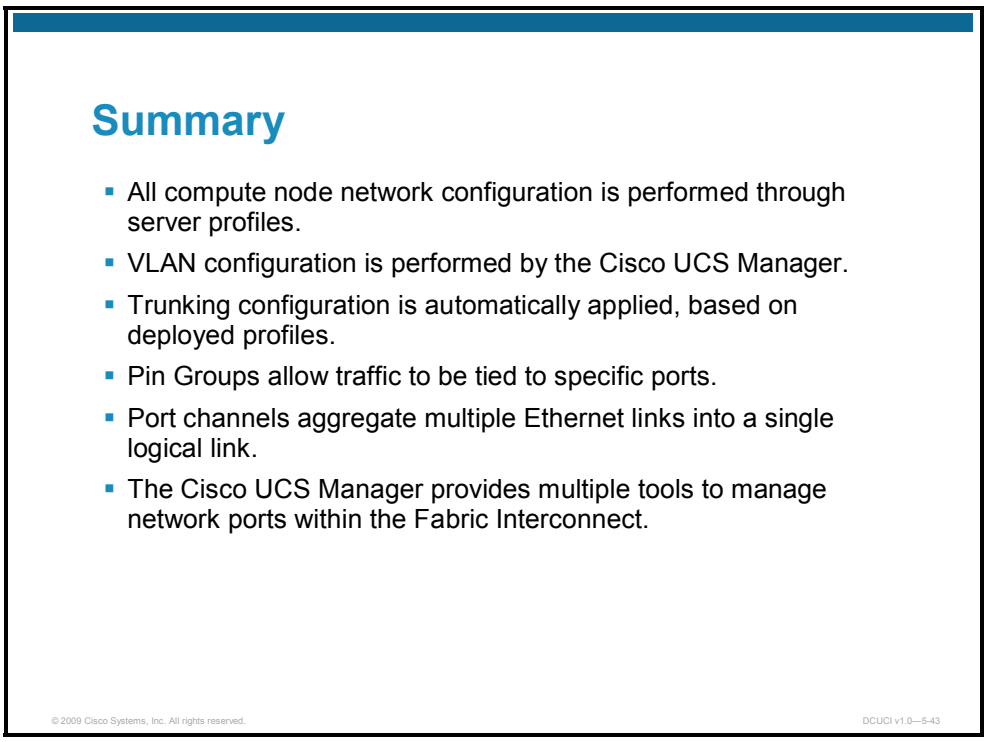

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Compute Node SAN Connectivity

Overview

In this lesson you will learn how to configure Cisco Unified Computing System (Cisco UCS) for connectivity between compute nodes and Fibre Channel networks through the Fabric Interconnect.

Objectives

Upon completing this lesson, you will be able to describe the configuration requirements and options for compute node SAN connectivity in Cisco UCS. This includes being able to meet these objectives:

- Describe how N port identifier virtualization provides a means to assign multiple FCIDs to a single N port
- Describe the configuration and use of VSANs
- Describe virtual host bus adapters
- Describe SAN uplink and server Fibre Channel ports
- Describe how Pin Groups function

https://t.me/learningnets

N Port Identifier Virtualization

This topic describes how N port identifier virtualization (NPIV) provides a means to assign multiple Fibre Channel IDs (FCIDs) to a single N_port.

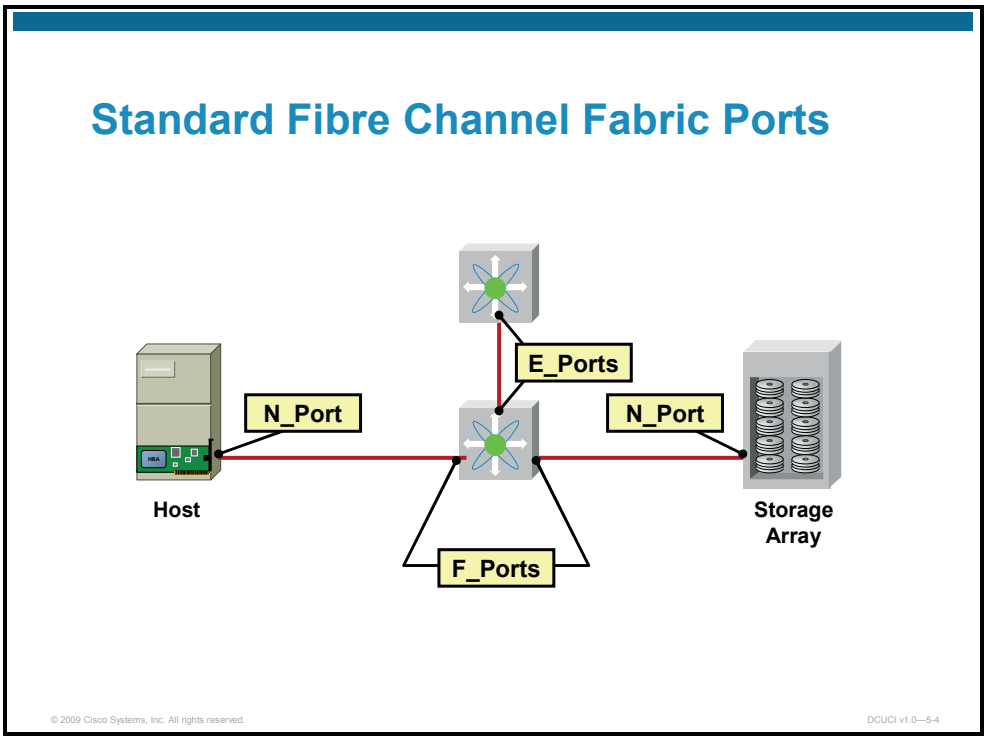

- An N_Port (Node Port) is a port on a node that connects to a fabric:
  - I/O adapters and array controllers contain one or more N_Ports
  - N_Ports can also directly connect two nodes in a point-to-point topology
- An F_Port (Fabric Port) is a port on a switch that connects to an N_Port.
- An E_Port (Expansion Port) is a port on a switch that connects to an E_Port on another switch.
- An FL_Port (Fabric Loop Port) is a port on a switch that connects to an arbitrated loop:
  - Logically, an FL_Port is considered part of both the fabric and the loop.
  - FL_Ports are always physically located on the switch.
- Note that Fibre Channel hubs, although they obviously have physical interfaces, do not contain Fibre Channel ports. Hubs are basically just passive signal splitters and amplifiers. They do not actively participate in the operation of the network. On an arbitrated loop, the node ports manage all Fibre Channel operations.
- Not all switches support FL_Port operation. For example, some McDATA switches do not support FL_Port operation.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

- An NL_Port (Node Loop Port) is a port on a node that connects to another port in an arbitrated loop topology. There are two types of NL_Ports:
  - Private NL_Ports can communicate only with other loop ports.
  - Public NL_Ports can communicate with other loop ports and with N_Ports on an attached fabric.
- Note that the term L_Port (Loop Port) is sometimes used to refer to any port on an arbitrated loop topology. L_Port can mean either FL_Port or NL_Port. In reality, there is no such thing as an L_Port.

https://t.me/learningnets

N_Port Virtualizer Mode Network Topology

Fabric treats a Fabric Interconnect as an Fibre Channel node with multiple ports and Fibre Channel IDs:

- Cisco UCS Fabric Interconnect functions in N_Port Virtualizer (NPV) mode
- Uplink ports function as N_Port proxies (**NP_Ports**)
- Blade server interfaces function as **F_Ports**

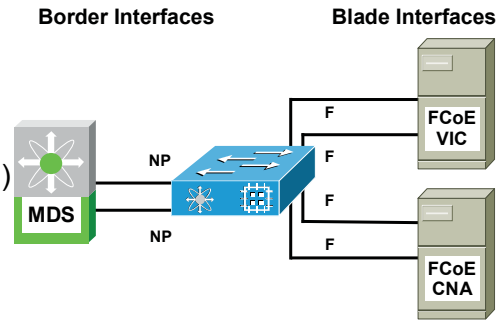

© 2009 Cisco Systems, Inc. All rights reserved.

There are two types of interfaces in NPIV topologies.

Server Interfaces

Server-facing interfaces are either physical Fibre Channels, or virtual Fibre Channel interfaces with F_Port modes:

- No local switching
- All packets are forwarded
- FLOGI-related processing is relayed in software (FLOGI, FDISC, and corresponding LS_ACC, LS_RJT, and so on) to the same uplink interface
- NPIV supported on F_Ports

Border Interfaces

Border interfaces are network facing and will always be port type NP:

- Internal FLOGI is sent to the core Fibre Channel switches
- Register with name servers on the successful internal FLOGI
- Every uplink can be connected to different Fibre Channel switches and virtual storage area networks (VSANs)

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

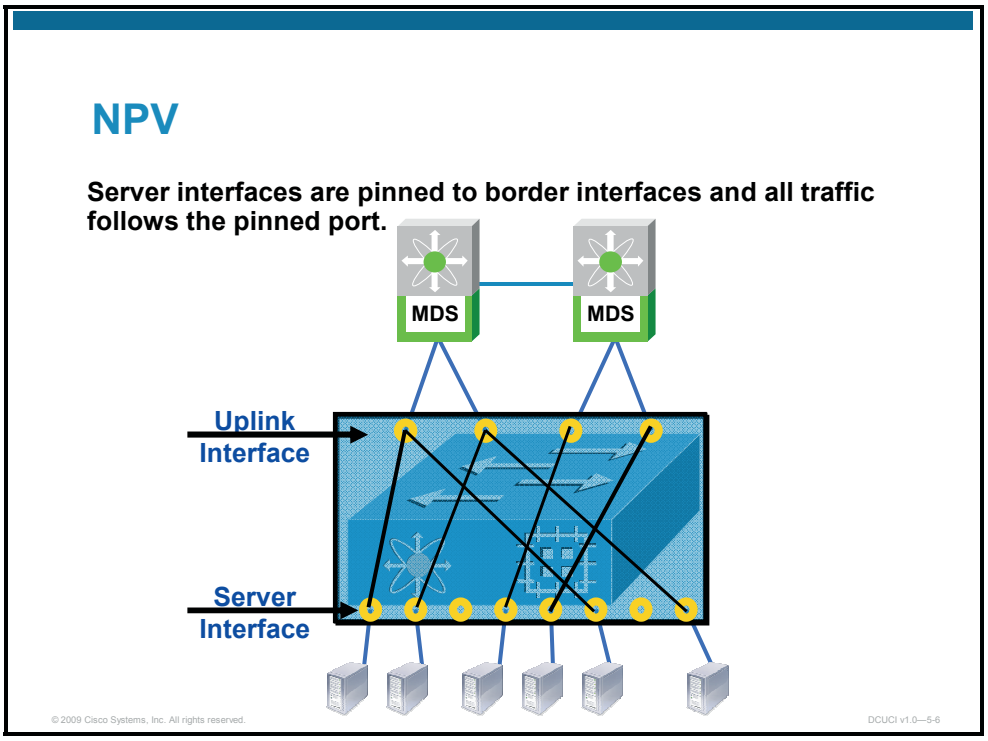

Each server link is pinned to exactly one uplink:

- Pinning logic distributes server links to various uplinks.
- All traffic is passed up-stream for switching.

https://t.me/learningnets

NPV in Detail

- Each Server Interface is pinned to one Border Interfaces
- All traffic follows the pinned port
- Pinning is based on a round robin mechanism
- Supports NPIV over Server Interfaces (nested NPIV)
- Relies on NPIV (NPIV must be enabled on the core switch)
- Retries all failed login requests from a Border Interface on a new Server Interface
- Handles error events by generating proxy LOGOs

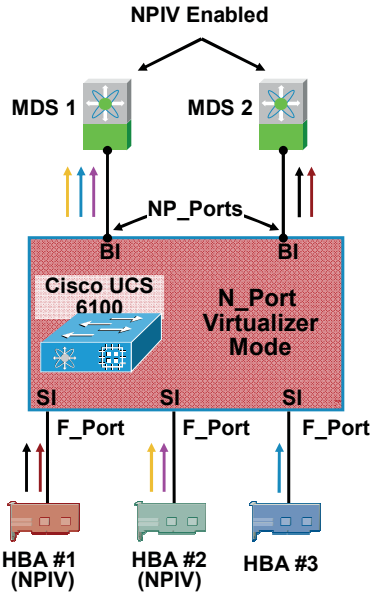

© 2009 Cisco Systems, Inc. All rights reserved.

NPIV proxy modules in the Cisco Nexus Operating System (NX-OS) provide the proxy function of distributing fabric login (FLOGI) requests from servers over the available Border Interfaces (BIs). The Fibre Channel Host Bus Adapters (HBAs) in servers and Fibre Channel switches assume they are directly connected to each other by using a physical cable.

Note	The NPIV proxy function allows the NPIV to be used between the Fabric Interconnect and the Fibre Channel switch, even if some or all HBAs implement only the basics of N_Port functionality.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Pinning and Load Balancing

Pinning and load balancing:

- Server interfaces pinned to uplink interfaces (not FLOGI)

Data packets on the server interface:

- Packets are sent to the pinned interface without forwarding lookup (NPV switches do not participate in FSPF)
- Binding checks are performed to make sure the server IDs are on the right server interface, prevents address spoofing.

Data packets on the external interface:

- A forwarding lookup is performed based on the domain ID frames that point to the server interface
- Packets are discarded on miss

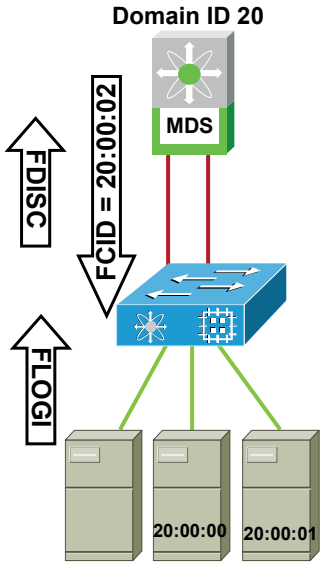

© 2009 Cisco Systems, Inc. All rights reserved.

Using NPV, each downstream device (server or blade server) will be pinned to an uplink port based on a round robin algorithm. The NPV mode switch will no longer handle FLOGI login requests or make routing decisions using Fabric Shortest Path First (FSPF). Instead, these operations are passed to the upstream switch, which is known as the NPV core switch. The NPV core switch will utilize NPIV in order to interpret multiple logins from the same port.

https://t.me/learningnets

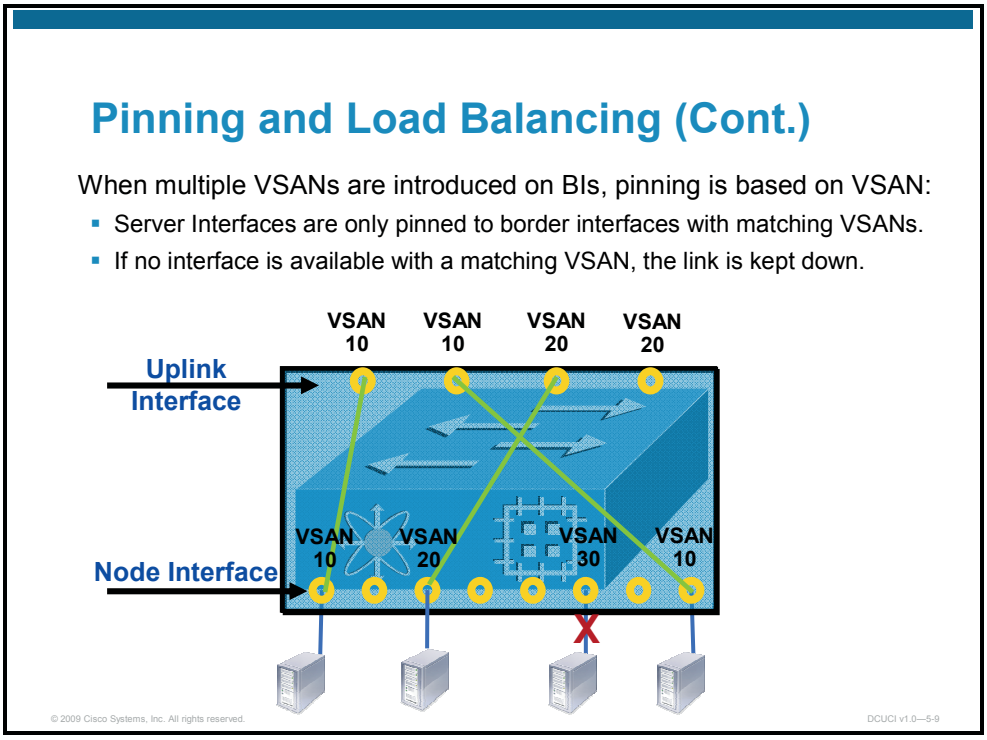

Uplink Interfaces and Node Interfaces can only be configured for a single VSAN. Node Interfaces will only be pinned to a port of the correct VSAN.

In the example:

- Two uplink interfaces are configured for VSAN 10 and the other two are configured for VSAN 20.
- Two blade ports are configured for VSAN 10
- One blade port is configured for VSAN 20
- One server port is configured for VSAN 30
- The blades configured for VSAN 10 will be pinned on one of the VSAN 10 BIs
- The blade configured for VSAN 20 will be pinned to one of the VSAN 20 BIs
- The blade interface configured for VSAN 30 will be kept down because there is no matching uplink interface

© 2009 Cisco Systems, Inc.
https://t.me/learningnets

VSANs

This topic describes VSANs as used within the Cisco UCS environment.

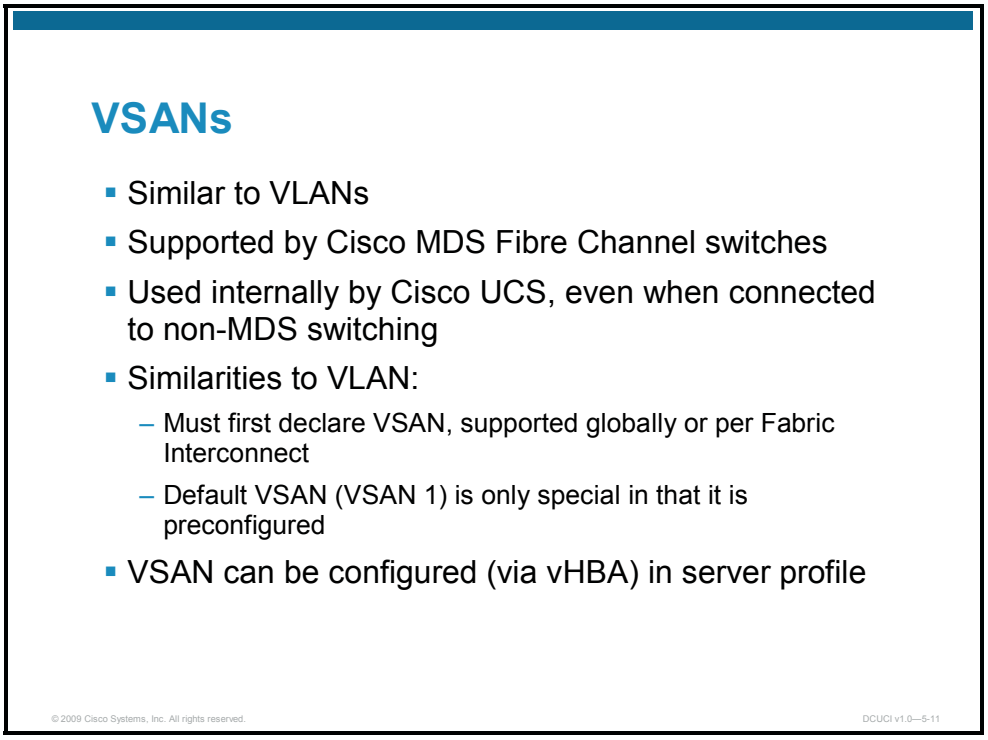

The VSAN/virtualized host bus adapter (vHBA) logic is similar to the VLAN/vNIC logic discussed in the previous lesson.

You must have a Cisco UCS Manager vHBA object in a profile to specify connectivity for a Fibre Channel adapter on a compute node to which the profile will be associated.

To support any VSAN, that VSAN must be configured globally into Cisco UCS Manager, and then 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.

https://t.me/learningnets

VSANs and Non-MDS Switches

- VSANs are still used internally by Cisco UCS
- Each uplink connects to one physical fabric; mapped to a VSAN number in Cisco UCS
- No VSAN information traverses the uplink port
- Cisco UCS uses NPIV protocol to connect to upstream Fibre Channel switch
- Use host-based Fibre Channel failover for redundancy
  - While vNICs can be automatically failed over (CNA, VIC), vHBAs are never failed over

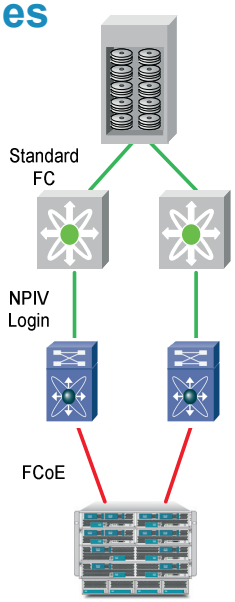

© 2009 Cisco Systems, Inc. All rights reserved.

While VSANs are used by Cisco MDS switches, other vendors have not yet implemented VSAN technology within their switching architectures. The 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 through the use of VSANs and no inter-VSAN routing is possible.

https://t.me/learningnets

FCoE VLANs

- Must specify a FCoE VLAN for each VSAN
- All server Fibre Channel traffic is carried via FCoE in dedicated VLANs
- FCoE VLANs must not conflict with Cisco UCS Manager VLAN objects
- Select unused range of VLANs and dedicate that range to FCoE

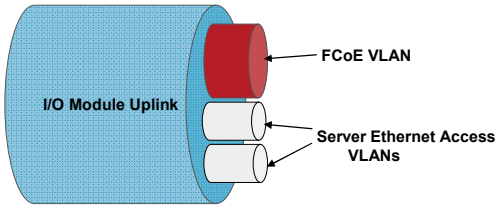

© 2009 Cisco Systems, Inc. All rights reserved.

Due to the internal Fibre Channel over Ethernet (FCoE) architecture of Cisco UCS, each VSAN supported within the architecture requires a dedicated VLAN to carry the FCoE traffic. FCoE VLANs are designated during VSAN configuration and are not created like Ethernet VLANs, as shown in the previous lesson. The FCoE VLANs must not conflict with Ethernet VLAN objects. It may be wise to designate an unused range of VLANs and dedicate that range to FCoE traffic; the specific range will depend on the VLAN usage and requirements within the organization.

https://t.me/learningnets

Creating VSANs

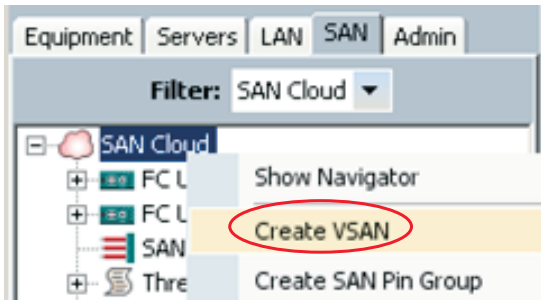

VSANs can be configured globally, per Fabric Interconnect, or with the same label, but configured differently on each Fabric.

© 2009 Cisco Systems, Inc. All rights reserved.

To create a VSAN object with Cisco UCS, right click the **SAN Cloud** object in the SAN navigation tab.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Configuring Global VSANs

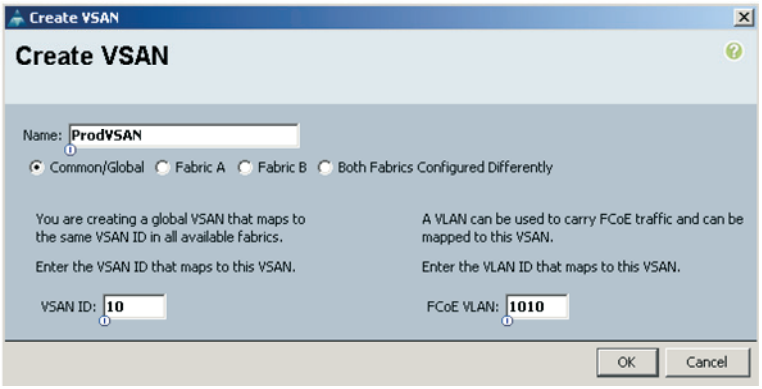

- Common/Global VSANs will be available on both Fabric Interconnects
- Requires uplinks from both Fabric Interconnects to the same physical upstream fabric

© 2009 Cisco Systems, Inc. All rights reserved.

Common/Global VSANs are created on both Fabric Interconnects and will use the same VSAN ID and FCoE VLAN. This type of configuration requires that both Fabric Interconnects be connected to the same physical Fibre Channel fabric.

https://t.me/learningnets

Configuring Local VSANs

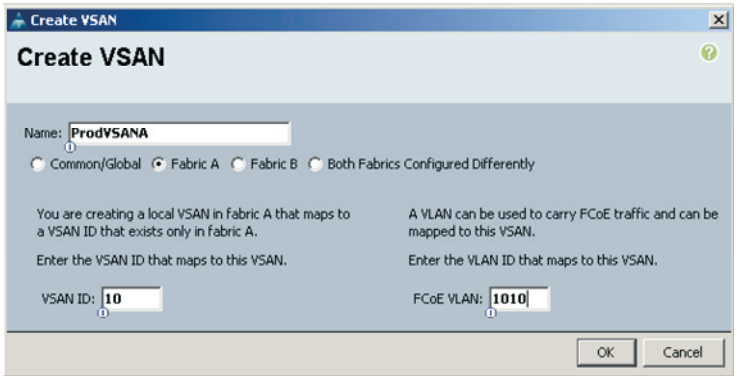

- Locally configured VSANs will only be available on the Fabric Interconnect designated
- Closely maps most traditional SAN deployments with physical separation of fabrics

© 2009 Cisco Systems, Inc. All rights reserved.

Locally configured VSANs are only available through the Fabric Interconnect on which they are configured. In this configuration, each Fabric Interconnect would likely connect to different upstream Fibre Channel fabrics; this closely resembles traditional deployments and maintains the physical separation of the fabrics. Only the mezzanine cards on the compute nodes would be attached logically or physically to the same physical fabric.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Configuring VSANs Differently

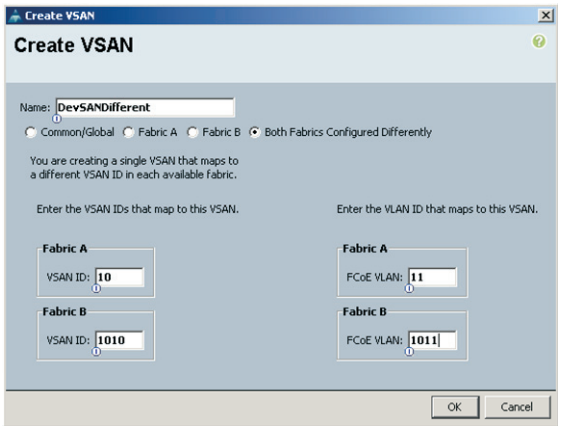

- Differently configured VSANs allow the same label to map to different VSAN and FCoE VLANs on each Fabric Interconnect

© 2009 Cisco Systems, Inc. All rights reserved.

When using the “Both Fabrics Configured Differently” option, the same VSAN label is applied to two separate VSANs and FCoE VLANs. Note that for this to behave as expected, the same external resources need to be accessible through each of the upstream fabrics.

https://t.me/learningnets

Virtual HBAs

This topic describes virtual host bus adapters (vHBAs).

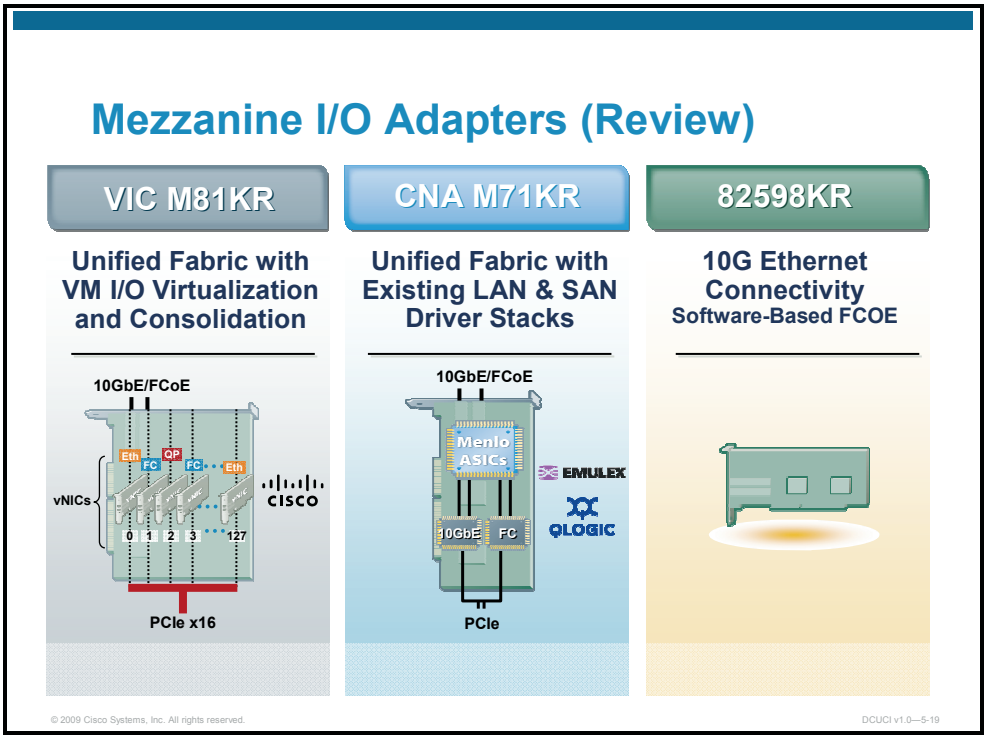

Mezzanine Cards

The types of adapters, or mezzanine cards, available in a Cisco UCS blade server include:

- Cisco UCS VIC M81KR: A virtual interface card
- Cisco UCS CNA M71KR-E: An Emulex converged network adapter
- Cisco UCS CNA M71KR-Q: A QLogic converged network adapter
- Cisco UCS 82598KR-CI: A 10-GE network adapter

Cisco UCS VIC M81KR

The Cisco UCS VIC M81KR supports multiple Ethernet virtual network interface cards (vNICs) or Fibre Channel Host Bus Adapters (vHBAs), running at 500 KIOPs in both initiator and target mode. This card supports kernel/hypervisor bypass. The Cisco UCS VIC M81KR is a converged network adapter (CNA) with dual 10-GE ports and dual Fibre Channel ports to the backplane. The card provides failover between the redundant Ethernet links. No multipathing software is required on the host operating system for this functionality.

Cisco UCS CNA M71KR

The Cisco UCS CNA M71KR is a converged network adapter (CNA) with two host-side 10-GE ports and two Fibre Channel ports to the backplane. The two network ports can run either native Ethernet or FCoE protocols and can be configured for failover. This failover is performed by the FCoE ASIC and does not require multipathing software on the host operating system:

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

The Cisco UCS CNA M71KR ASIC is a Cisco-designed multiplexor and FCoE protocol offload engine with a 350 MHz 24K MIPS processor. There are two versions of this card (Cisco UCS CNA M71KR-E Emulex Converged Network Adapter, and Cisco UCS CNA M71KR-Q QLogic Converged Network Adapter), thereby supporting existing proven driver stacks to the customer. Both cards use an Intel 82598 10GE Ethernet chipset.

Cisco UCS 82598KR-CI

The Cisco UCS 82598KR-CI Network Adapter is made by Intel. It has two 10-GE ports to the backplane which run native Ethernet. operating system-implemented FCoE may be run over the Cisco UCS 82598KR-CI to provide “free SAN access” to the host. The card does not provide failover functionality.

Comparison

- The Cisco UCS VIC M81KR provides enhanced virtualization. This adapter supports mulitple vNICs presented to the ESX server which can then be used as physical NICs for virtual machines (VMs).
- The Cisco UCS CNA M71KR adapter provides increased scalability. The FCoE protocol encapsulation is offloaded from the host and performed in hardware on the Cisco UCS CNA M71KR.
- The Cisco UCS 82598KR-CI provides standard 10-GE connectivity to servers not needing Fibre Channel protocol support, or servers providing Fibre Channel protocol and FCoE encapsulation in software.

https://t.me/learningnets

vHBAs

- For CNA adapter on blade:
  - Create one vHBA in the profile for each port that is to be accessible
  - Maximum of two vHBA per adapter card; one to Fabric Interconnect A, and one to B
  - No hardware failover is supported
  - Use host-based multi-pathing for redundancy

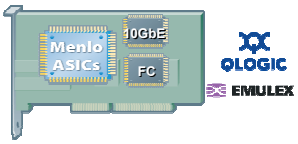

© 2009 Cisco Systems, Inc. All rights reserved.

Each Menlo-based mezzanine adapter can support two vHBAs. One vHBA will connect to Fabric A, the other to Fabric B. No hardware failover is supported, unlike the vNIC behavior. Any redundancy of the fabrics will be managed by multipathing software or drivers within the operating system of the compute node.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

vHBA with Virtual Interface Card

- For virtual interface card mezzanine adapter:
  - Create multiple vHBAs if desired
  - Number of vHBAs supported depends on cabling topology
  - Each is presented as physical HBA to the operating system on the blade server
  - Choose Fabric Interconnect A or B
  - No hardware based failover of vHBAs

© 2009 Cisco Systems, Inc. All rights reserved.

The virtual interface card can support many vHBA adapters. Each vHBA is presented to the host operating system as a physical adapter. The number of possible adapters is reviewed on the following slide.

https://t.me/learningnets

Virtual Interface Card Limitations

- Number of possible adapters depends on uplinks between I/O Module and Fabric Interconnect
- 128 virtual adapters is the architectural limit of the adapter chipset
- Each Fabric Interconnect port supports 128 VNTags (the mechanism used to support virtual interfaces)
- Some VNTags are consumed by chassis management functions

With 1 I/O Module, 8 Blades		With 2 I/O Modules, 8 Blades	
Uplinks to Fabric Interconnect	Virtual Interfaces per Adapter	Uplinks to Fabric Interconnect	Virtual Interfaces per Adapter
1	13	1	26
2	28	2	56
4	58	4	116

© 2009 Cisco Systems, Inc. All rights reserved.

The number of virtual interfaces supported on an adapter depends on the number of uplinks between the I/O Module (IOM) and the Fabric Interconnect, as well as the number of interfaces in use on other adapters sharing the same uplinks. The numbers shown in the table reflect the worst case scenario, where all nodes are attempting to use virtual interfaces.

The Cisco Virtual Interface Card mezzanine adapter is an example of an interface card that can support the creation of multiple virtualized hardware interface ports.

These limits apply to any combination of vNICs and vHBAs.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

vHBA and VSAN

- Each vHBA supports one VSAN
- Association of vHBA to VSAN determines possible uplinks
- Will support NPIV from host operating system

© 2009 Cisco Systems, Inc. All rights reserved.

vHBAs can currently support one VSAN. No VSAN tagging of the Fibre Channel frame occurs. Only uplinks mapped to the same VSAN will be used, maintaining the separation of multiple physical fabrics, if they are in use.

https://t.me/learningnets

SAN Port Connectivity

This topic describes Fibre Channel configuration options for uplink and server ports.

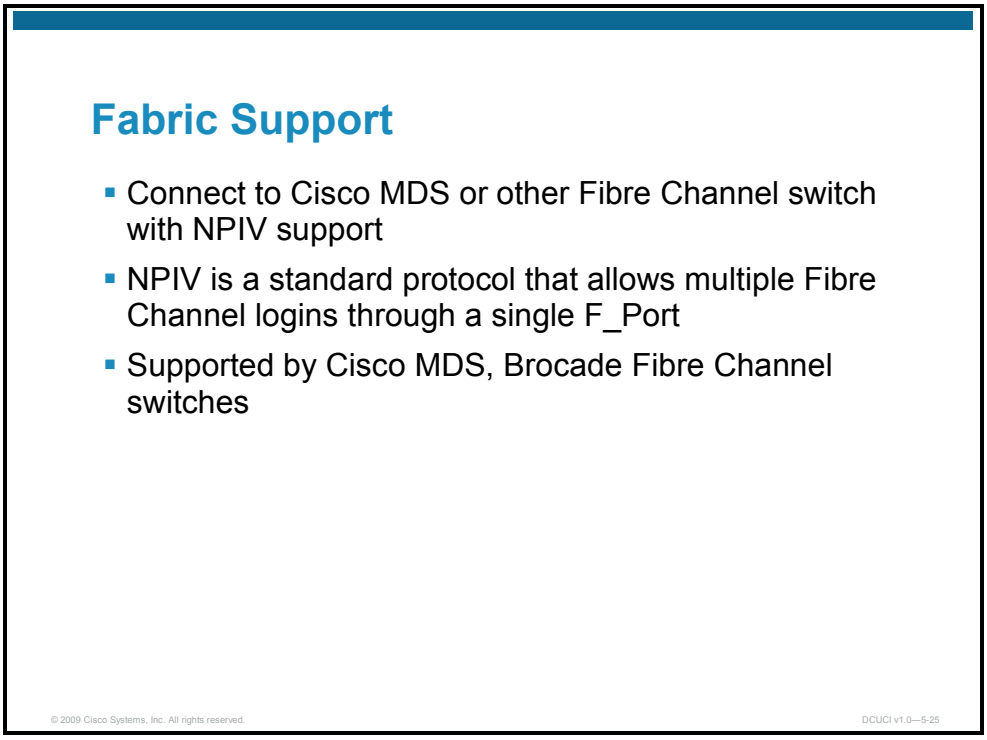

Cisco UCS supports connection to Cisco MDS Fibre Channel switches, or to any Fibre Channel switch supporting the NPIV standard.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Configuring Fibre Uplinks on VSANs

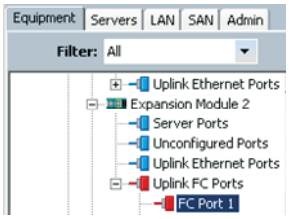

- VSAN assignment is by named VSAN object
- Each uplink supports one VSAN

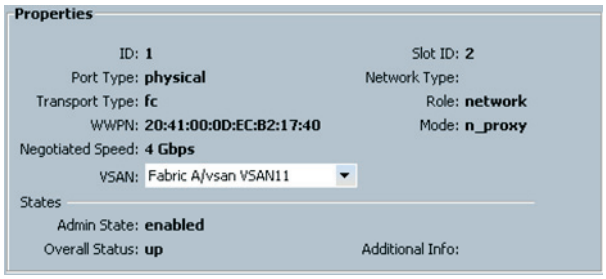

© 2009 Cisco Systems, Inc. All rights reserved.

Each fibre uplink (on the equipment panel) supports exactly one VSAN. You must choose an uplink that was globally configured into the system in the SAN panel.

The uplink port will not be configured correctly (the operational state will not be “up”) unless it is configured for the correct VSAN when connecting to a Cisco MDS switch.

If connecting to a non-MDS switch, the VSAN assignment will only determine which Cisco UCS-internal VSAN is mapped to the non-MDS physical fabric. Through the use of internal VSANs, Cisco UCS supports connection to multiple physical Fibre Channel fabrics without merging them.

https://t.me/learningnets

Server Ports

- Server ports are not directly configured
- Server VSAN and traffic properties are assigned dynamically, based on service profile associations
- Each vHBA belongs to a single VSAN

© 2009 Cisco Systems, Inc. All rights reserved.

Like server Ethernet connections, all vHBAs are configured through the use of server profiles. No direct configuration of the vHBA is possible or permitted. Each server vHBA can support one VSAN mapping per port.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Pin Groups

This topic describes how Pin Groups are used to organize or segregate Fibre Channel traffic on uplink Fibre Channel links.

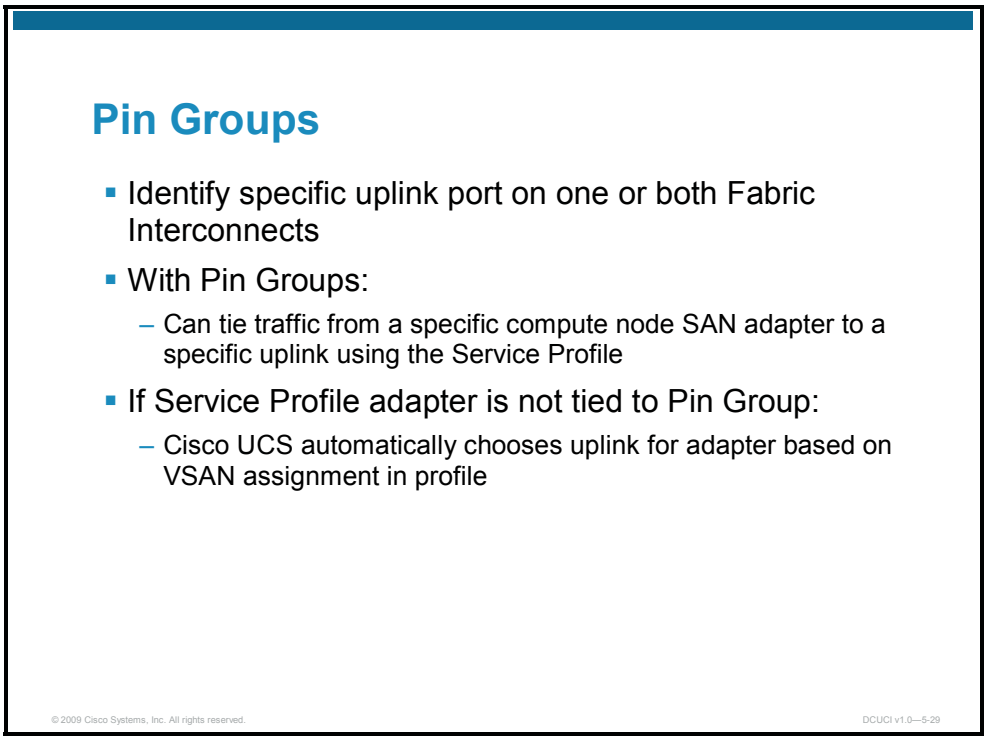

Pin Groups are a mechanism for pinning traffic from a specific compute node mezzanine adapter to a specific individual port on each Fabric Interconnect.

Without Pin Groups, the Cisco UCS automatically chooses an uplink port for the adapter on the blade associated with each profile.

https://t.me/learningnets

Summary

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

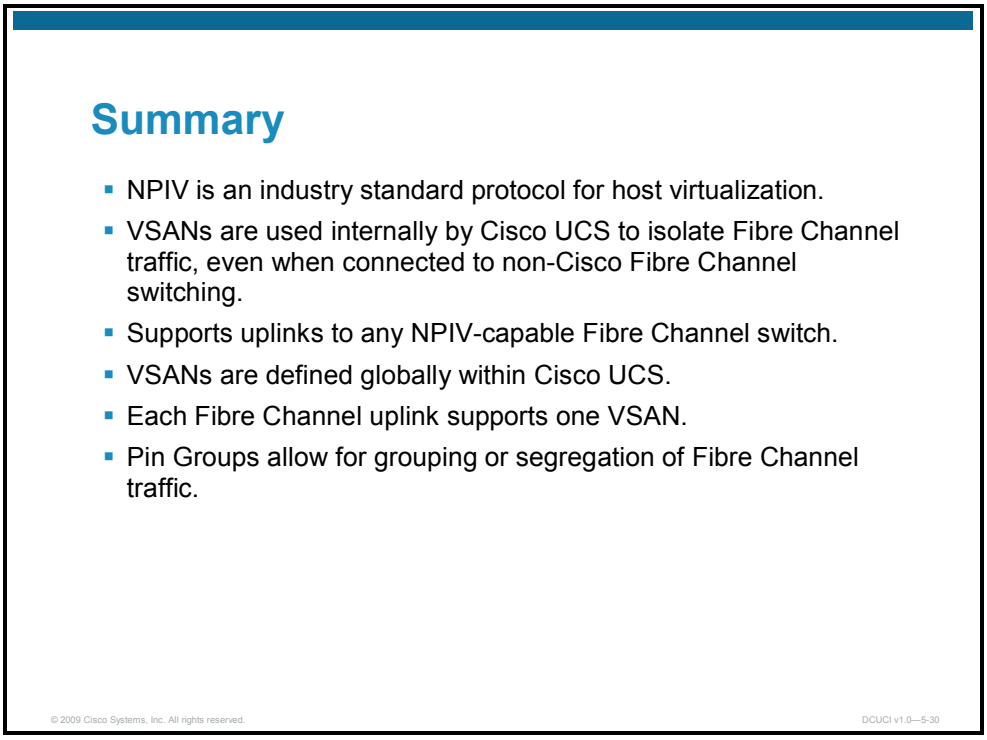

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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

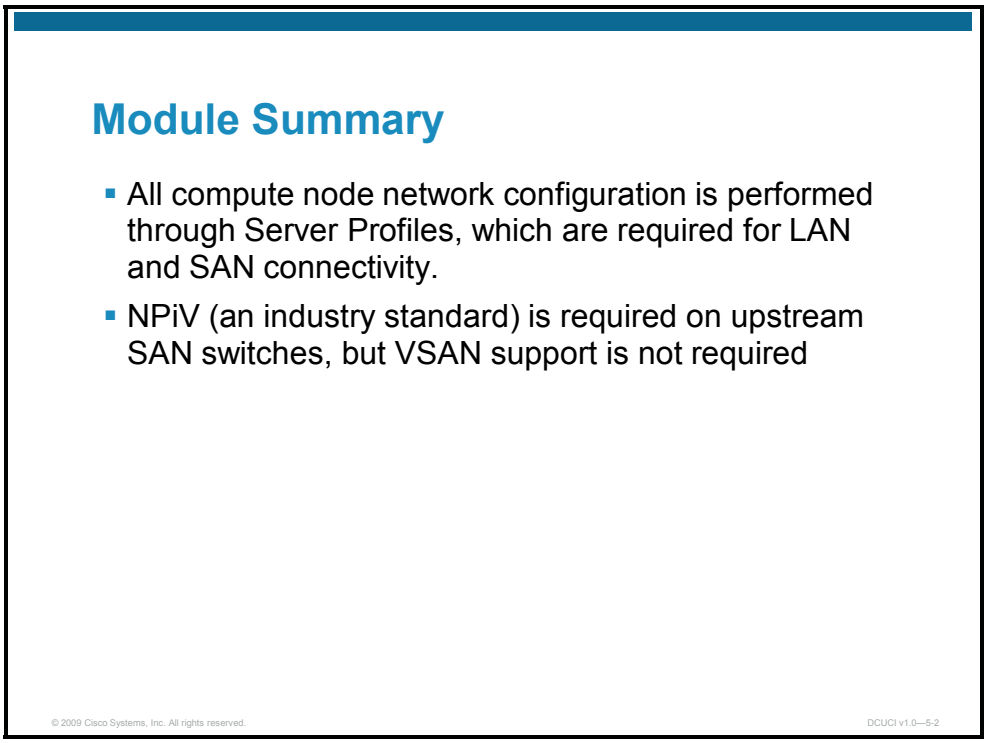

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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) What is the purpose of the Ethernet Host Virtualizer (EHV) feature? (Source: Compute Node LAN Connectivity)

A) Allows administrators to transparently migrate virtual machines
B) Allows Cisco UCS to virtualize server MAC addresses
C) Allows active/active uplinks from the Fabric Interconnect to the LAN
D) Allows virtual machines to connect transparently to the LAN

Q2) What feature must be enabled on any Fibre Channel switch to which the Fabric Interconnects are connected? (Source: Compute Node SAN Connectivity)

A) N-Port ID Virtualization (NPiV)
B) N-Port Virtualizer (NPV)
C) Fabric Shortest Path First (FSPF)
D) Fibre Channel Name Services (FCNS)

https://t.me/learningnets

Module Self-Check Answer Key

Q1) C

Q2) A

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 6

Implementation of Server Resources

Overview

This module describes the process of assigning physical resources in a Cisco Unified Computing System (Cisco UCS) environment.

Module Objectives

Upon completing this module, you will be able to describe the process of assigning physical resources through server profiles. This includes being able to meet these objectives:

- Describe how to create identity and resource pools
- Describe how to create Service Profiles on the Cisco Unified Computing System (Cisco UCS) platform
- Describe how to manage Service Profiles

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Identity and Resource Pools

Overview

In this lesson you will learn how to create identity and resource pools in Cisco UCS Manager.

Objectives

Upon completing this lesson, you will be able to describe how to create identity and resource pools. This includes being able to meet these objectives:

- Describe the purpose and use of Mobile Service Profiles
- Describe MAC addresses and pools
- Describe the format and use of world wide names
- Describe the purpose and use of universally unique identifiers
- Describe server pools and how they are used

https://t.me/learningnets

Mobile Service Profiles

This section describes the purpose and use of Mobile Service Profiles.

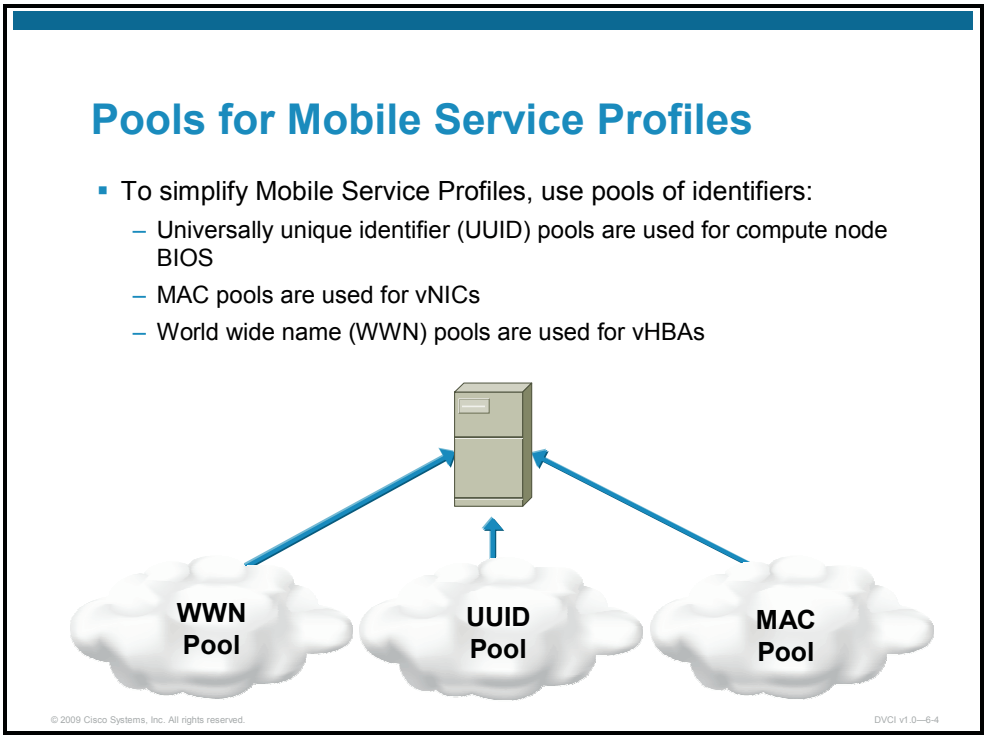

The figure describes pools for Mobile Service Profiles.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

OUIs and Hardware Identifiers

- Organizationally unique identifiers are used by hardware manufacturers in the generation of MAC and WWN addresses
- Assigned by IEEE
- In binary, first octet of IEEE OUI will always end in 00 or 01:
  - If first octet ends in 10 or 11, address is locally assigned

Binary	Hexadecimal
0000 00**1**0	02
0000 01**1**0	06
0000 10**1**0	0A

© 2009 Cisco Systems, Inc. All rights reserved.

Organizationally unique identifiers (OUIs) are assigned by the IEEE to allow hardware manufacturers to guarantee universal uniqueness of their addressing schemes. To provide a mechanism for addresses that are not assigned by the IEEE, and without the risk of conflicts, a method was devised to identify IEEE-assigned addresses from those locally-assigned. For the 3-byte field where an OUI would be used, the least and second least significant bits hold special meaning. The least significant bit (that is, the last digit as represented in binary) is set to 0 to indicate a unicast address, or 1 to indicate a multicast address. The second-least significant bit (that is, the second to last digit as represented in binary) is set to 1 to indicate a locally administered or assigned address. An IEEE OUI will always have a 0 as its second-least significant bit.

https://t.me/learningnets

MAC Addresses

This section describes the format of MAC addresses and the creation of MAC address pools.

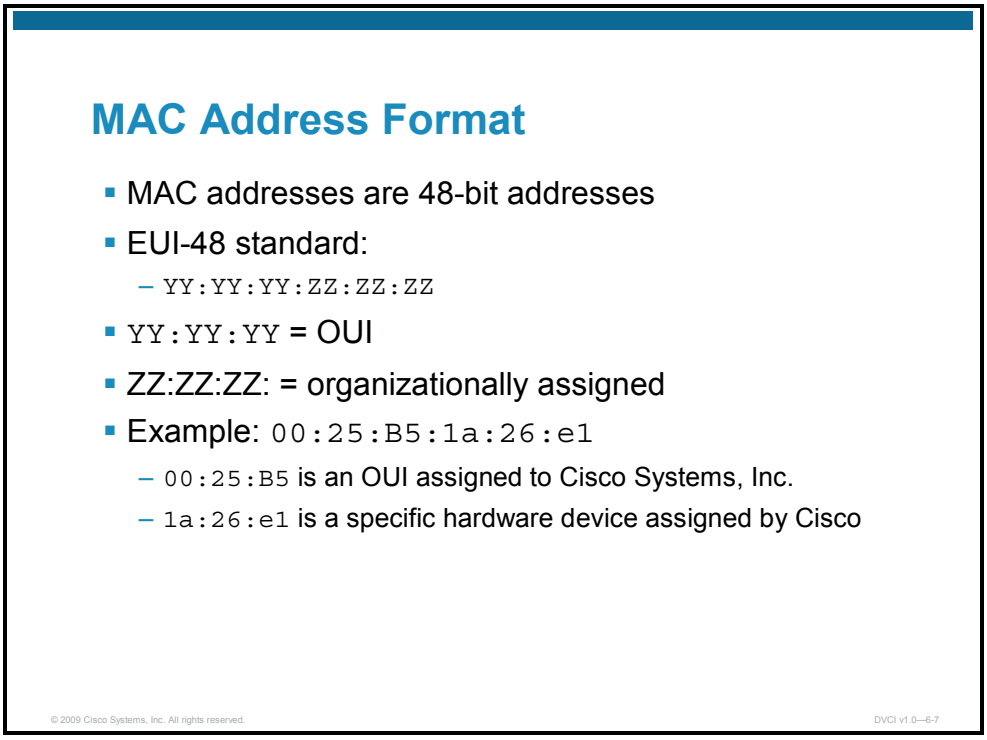

MAC addresses use the EUI-48 standard to define a 48-bit structure. The first 3 octets (24-bits) represent the OUI of the device manufacturer (or as previously discussed, a locally-administered OUI), while the final 3 octets are assigned by the organization itself. The example given uses a Cisco-assigned OUI (and specifically, the OUI used by Cisco UCS) and a randomly selected final 24-bits.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

MAC Address Considerations

- MAC addresses should be globally unique
- Cisco UCS Manager enforces a Cisco OUI of 00:25:B5
- Possibly encode meaning into fourth octet, such as Cisco UCS deployment, data center, etc.
  - By using 2 octets (16-bits), 32,767 possible MAC addresses
- Example: Data Center 1, Cisco UCS Deployment 4 might use 00:25:B5:14:00:00 to 00:25:B5:14:FF:FF
  - Easily distinguish location of MAC address, guaranteed uniqueness within organization

© 2009 Cisco Systems, Inc. All rights reserved.

MAC addresses should be globally unique, or at least unique within any Ethernet domains in which they might be used. Cisco UCS Manager enforces an OUI on all MAC pools of 00:25:B5, an IEEE-assigned OUI that is guaranteed not to conflict with products from other manufacturers. In the creation of a MAC address pool, it may be desirable to use the organizationally assigned field to include some meaning useful to the administrators of the system. This example suggests that the first 8 bits of the organizationally assigned range might use values to identify the Cisco UCS deployment and data center, to preserve uniqueness across multiple deployments or data centers within an organization. Optionally, the use of an external management entity (such as BladeLogic) can assist in keeping uniqueness among assigned values.

https://t.me/learningnets

Creating MAC Pools

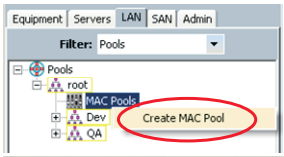

- Create a named MAC pool under the desired organization
- Pools will be available to all children organizations

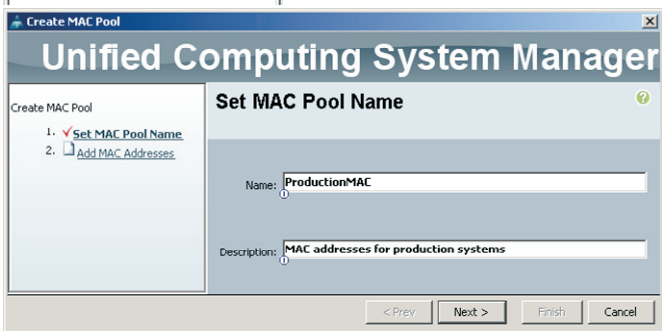

© 2009 Cisco Systems, Inc. All rights reserved.

A MAC pool is a range of MAC addresses. Create the MAC pool and assign a name. With MAC pools, Cisco UCS administration is made easier when scaling server deployment of Service Profiles, by prompting stake-holders to define a set of MAC addresses prior to actual deployment.

https://t.me/learningnets

Creating MAC Pools (Cont.)

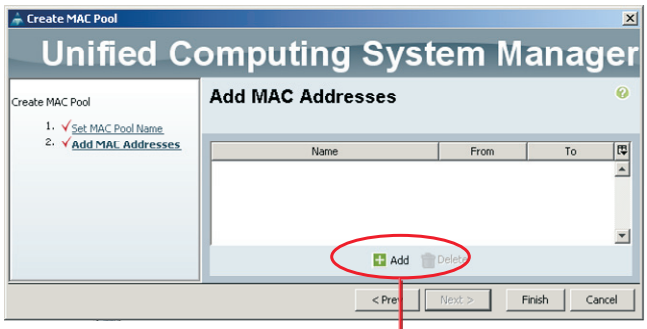

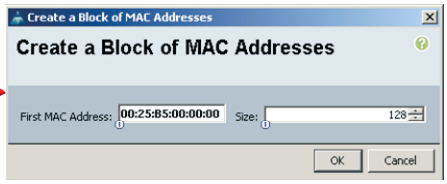

- Remember not to create more objects than necessary
- You can always add additional blocks later

© 2009 Cisco Systems, Inc. All rights reserved.

Provide the starting MAC address and the number of addresses to add to the block. Be aware that the Cisco UCS Manager creates an object in its database for each MAC address in the range. Creating an unnecessarily large block adds size and overhead to the Cisco UCS Manager database. It is recommended that MAC pools be created as part of the initial phase of a Cisco UCS deployment so that all stake-holders can have a chance to agree on the addressing that will be deployed when Service Profiles are later associated.

https://t.me/learningnets

World Wide Names

This section describes the format and use of world wide names (WWNs) and pools.

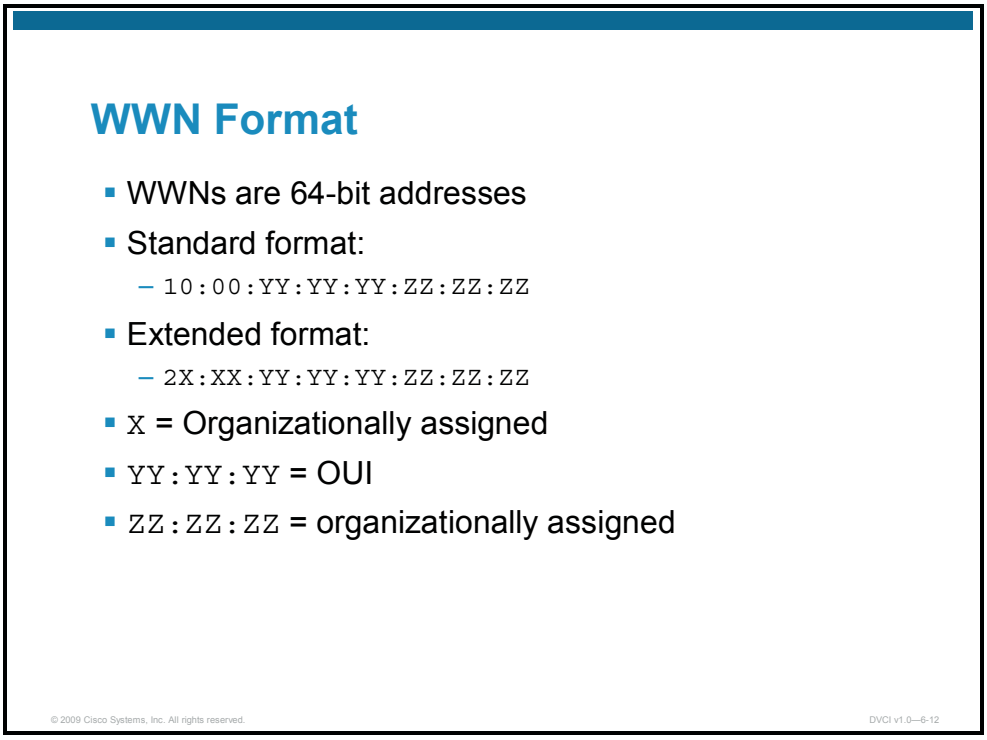

WWNs are 64-bit addresses that have many different possible formats. The information shown in the figure is for reference only, as Cisco UCS Manager enforces a specific format in all WWN pools.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

WWN Pool Considerations

- WWNs should be globally unique
- Recommend creating locally-administered OUI
- Possibly encode meaning into first octet of organizationally assigned field, such as with MAC addresses
- WWNN should be distinguishable from WWPN
  - WWNN might use `20:01:02:00:00:XX:XX:XX`
  - WWPN might use `20:00:02:00:00:XX:XX:XX`
- Pools may only use the 20:XX:XX:XX:XX:XX:XX:XX format

© 2009 Cisco Systems, Inc. All rights reserved.

Cisco UCS Manager enforces the use of world wide node names (WWNNs) that begin with "20:". All WWN pools must begin with that value, with any remaining values that the administrator wishes to create. In keeping with the standards set for WWNs globally, it is recommended that a locally-administered OUI be select (discussed earlier in this lesson) and used as the third through fifth octets. Additionally, it is useful if the WWNN and world wide pool name (WWPN) values be easily distinguishable, through any convention that the administrator desires. Because the second octet of a WWN can be organizationally assigned, that octet might be used to encode meaning for the WWNN or WWPN. This convention and address block should be agreed upon by all stake-holders in the initial implementation phase of a Cisco UCS deployment.

https://t.me/learningnets

WWNN Pool

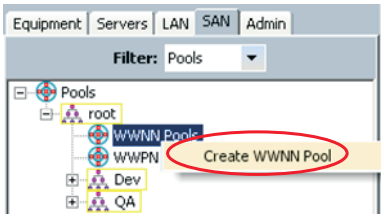

- Word wide node names are assigned on a per-Service Profile basis
- Create a named pool using process similar to that used to create MAC pools

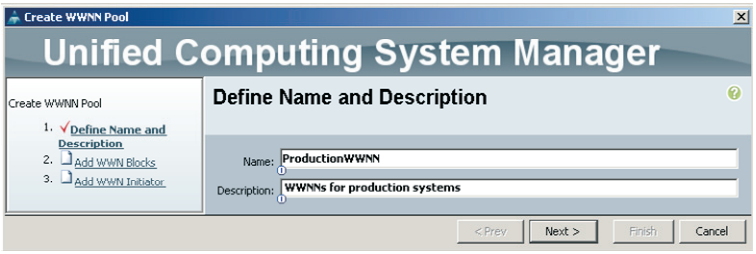

© 2009 Cisco Systems, Inc. All rights reserved.

WWN pools are similar to MAC pools.

You can use a WWN pool to choose a WWNN (associated with the entire Service Profile) and to choose a specific WWPN (associated with the specific vHBA). You must create separate pools for WWNNs and WWPNs. There is an available feature (discussed later in this module) that lets you associate SAN boot logical unit number (LUN) information with specific entries in the WWN pool. Much like an IP address scheme, WWNs and WWPNs should be carefully planned and determined as part of an initial phase of a Cisco UCS deployment.

https://t.me/learningnets

WWNN Pool

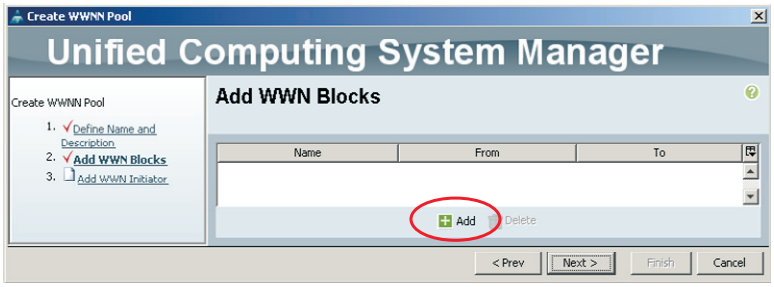

- Don't create too large a pool
- You can always add additional blocks later

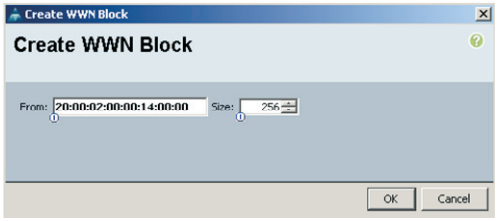

© 2009 Cisco Systems, Inc. All rights reserved.

Similar to MAC pools, an object is created in the Cisco UCS Manager database for each entry in the block. Creating an unnecessarily large pool adds size and overhead to the database. Additional blocks can be created later if necessary.

https://t.me/learningnets

WWPN Pool

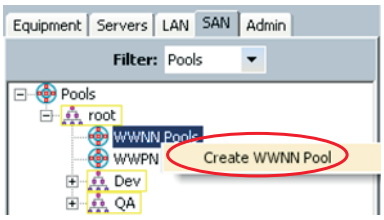

- Word wide port names are assigned on a per vHBA basis
- Create a named pool in a manner similar to the process for creating MAC pools

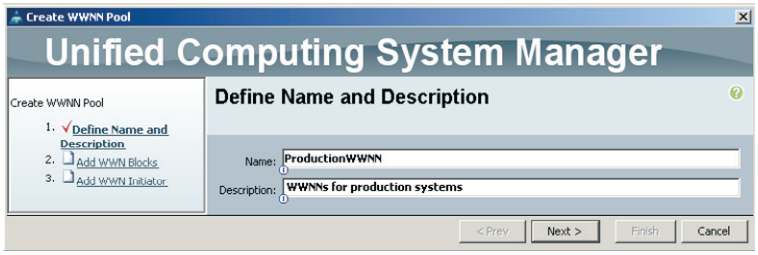

© 2009 Cisco Systems, Inc. All rights reserved.

Creating WWPN pools is identical to creating WWNN pools.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

WWPN Pool

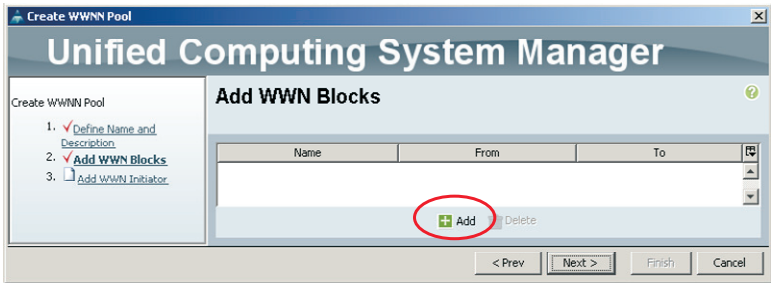

- Don't create too large a pool
- You can always add additional blocks later

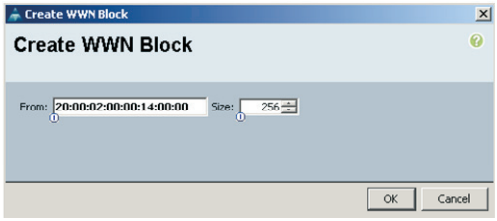

© 2009 Cisco Systems, Inc. All rights reserved.

This example shows the creation of a block of WWPNs.

https://t.me/learningnets

Adding Initiator Configuration

- Process is the same for WWNN or WWPN
- Allows for granular control of boot configuration based on assignment of WWN assigned:
  - If assigned to a particular WWN, use specific boot target
- Most environments will base storage configuration on WWPN

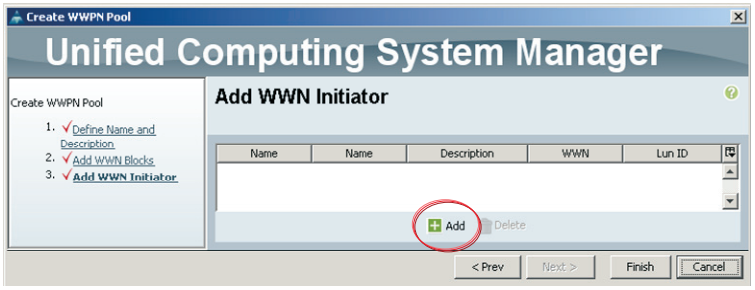

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes initiator configuration.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Adding WWN Initiator

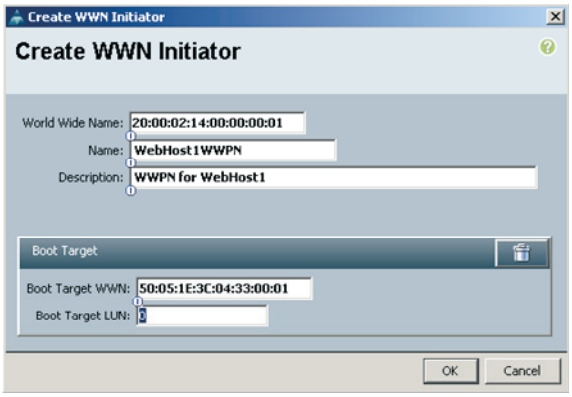

- Binds a specific WWN to boot configuration
- Useful if storage array requires specific boot targets, depending on source WWN

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes WWN initiator configuration.

https://t.me/learningnets

Universally Unique Identifiers

This section describes the purpose and use of universally unique identifiers (UUIDs).

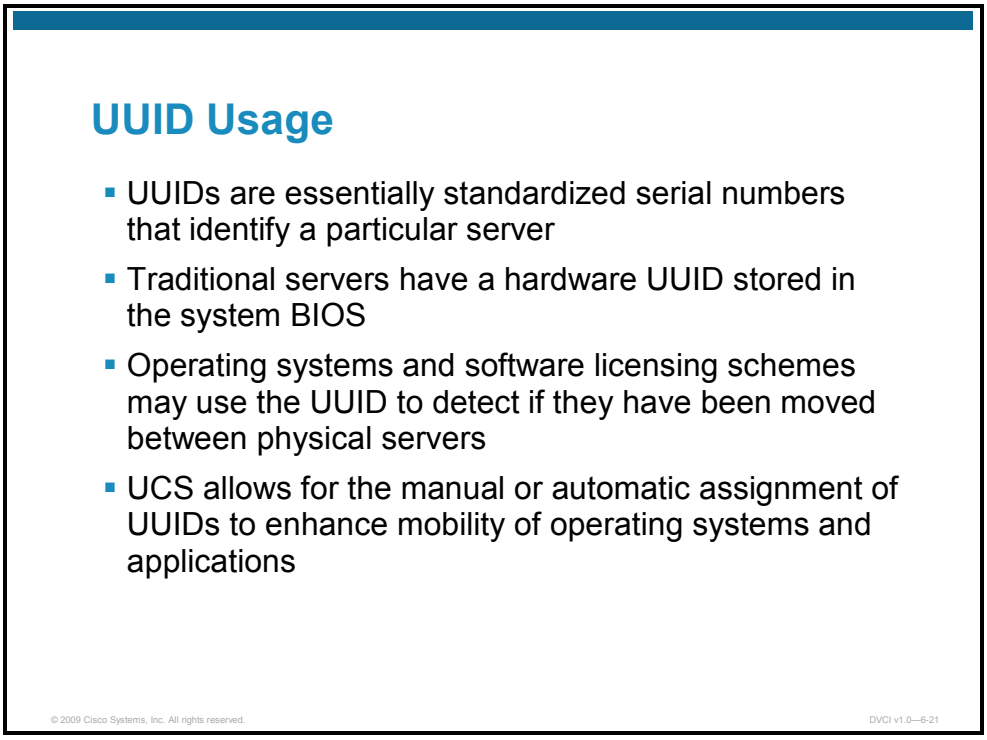

The figure describes the usage of UUIDs by operating systems and software licensing.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

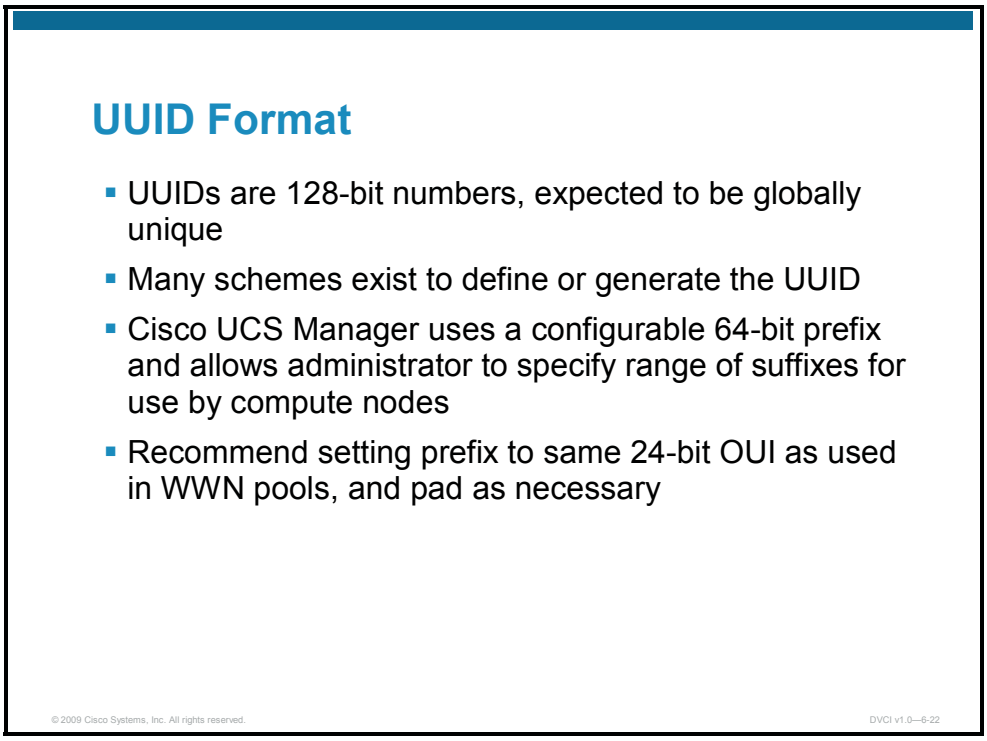

The figure describes the UUID format. UUIDs, much like other pools, should be discussed and planned out as part of an initial phase of a Cisco UCS deployment.

https://t.me/learningnets

Creating UUID Pools

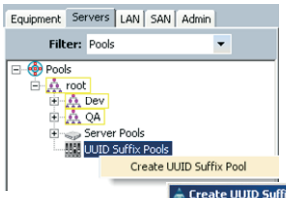

- Set a name, optional description, and prefix for the pool
- All UUIDs assigned from pool will begin with prefix

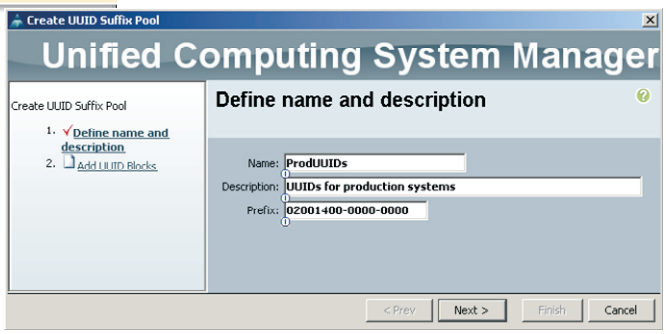

© 2009 Cisco Systems, Inc. All rights reserved.

To create a UUID pool, navigate to the **Servers** tab in the navigation pane.

Filter down to Pools and the organization in which the pool should be created. Right click **UUID Suffix Pools** and select **Create UUID Suffix Pool**.

Assign a name and optional description for the pool and click **Next**.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating UUID Pools (Cont.)

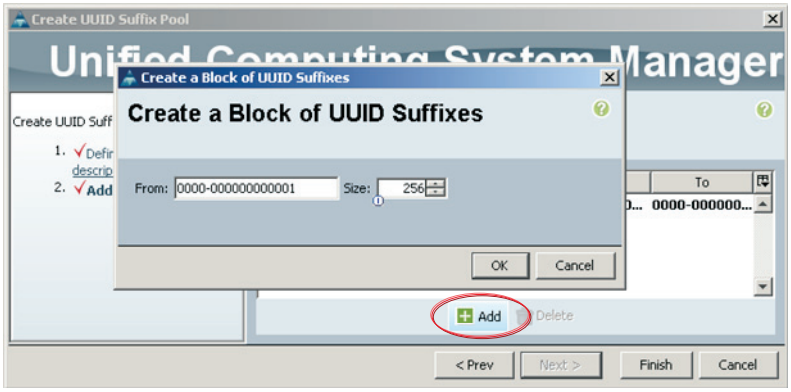

- Don’t create too large of a block
- Large blocks add unnecessary overhead and size to database

© 2009 Cisco Systems, Inc. All rights reserved.

Click **Add** and provide the starting UUID in the resulting dialog box. Select the number of UUIDs to include in the pool and click **OK**.

It is important not to create too large of a block of UUIDs. Each UUID creates an object in the Cisco UCS Manager database. Unnecessary UUIDs will add overhead and size to the database. If too small a block is created initially, the administrator can always add more blocks to the pool.

https://t.me/learningnets

Server Pools

This section describes server pools and how they are used in Cisco UCS.

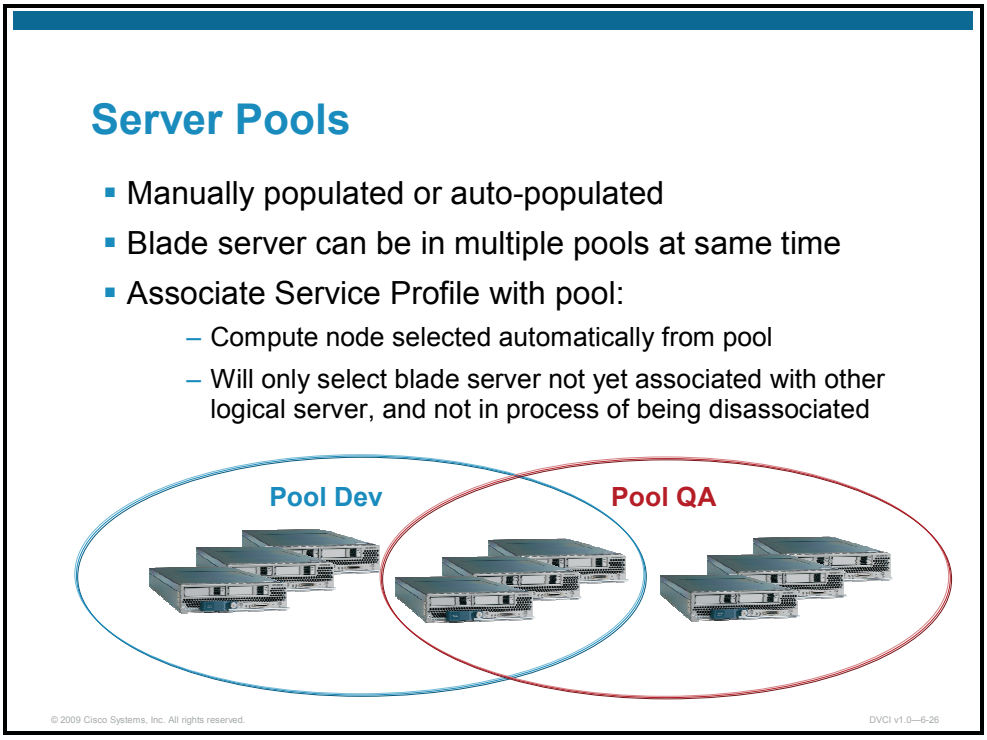

Pools can be manually populated or auto-populated using a Server Pool Policy, as demonstrated later in this lesson.

A compute node can be in multiple pools at the same time. The profile that is associated to a specific compute node owns the node, regardless of the number of pools the blade server resides in.

To use a server pool, associate the Service Profile with the pool. Cisco UCS Manager automatically selects an available compute node from the pool. An available blade server is one currently discovered, but not associated with any profile, and not in the process of being associated or disassociated.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Manual Server Pool

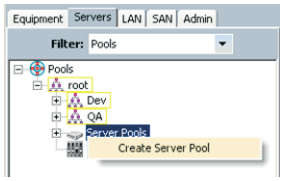

- Set a name and optional description for the pool
- Select compute nodes to be added to pool

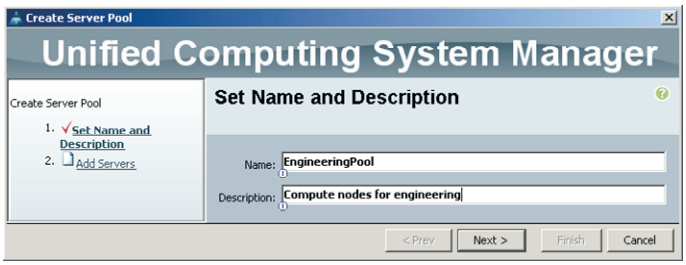

© 2009 Cisco Systems, Inc. All rights reserved.

Manually populating a server pool is the easiest way to create pools. Remember that a compute node can be in as many pools as desired.

https://t.me/learningnets

Creating Manual Server Pool (Cont.)

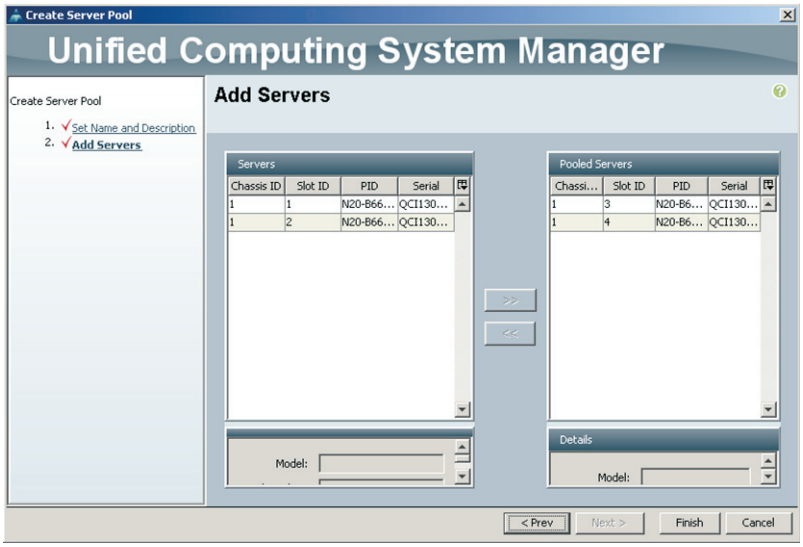

The figure shows how compute nodes are manually added to a pool.

https://t.me/learningnets

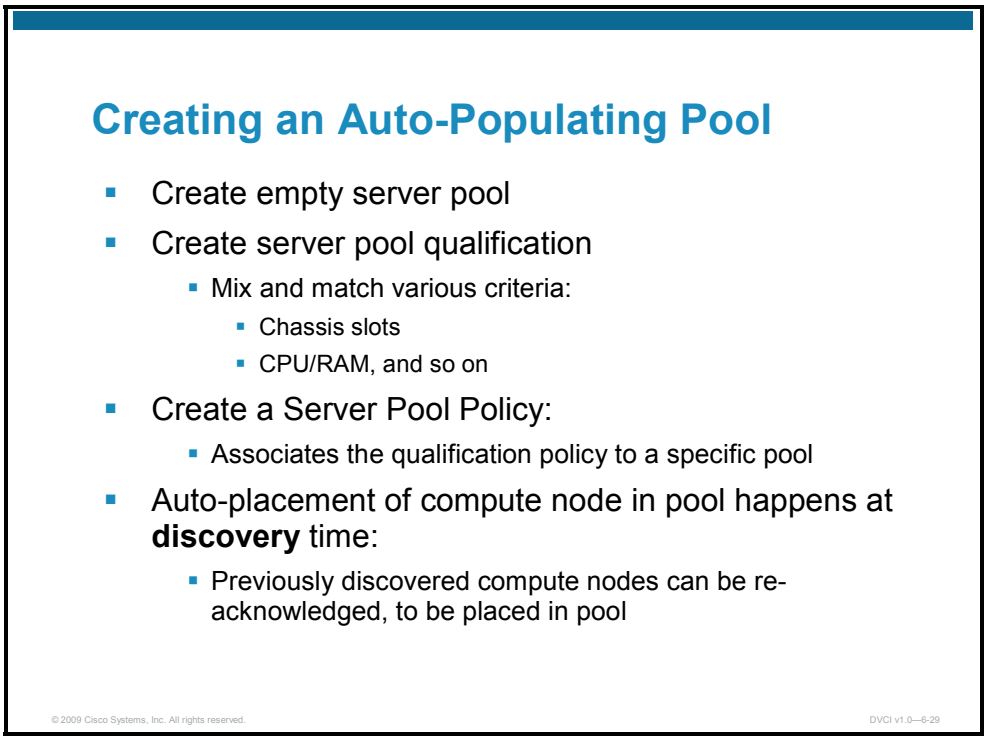

The auto-population feature lets you:

- Specify qualifications that will be used for matching specific blade servers
- Specify server pool policies—these will put every blade server that matches a particular qualification into a particular server pool

Note that server pool auto-population only happens as individual blade servers are discovered. If you want auto-population to apply to existing compute nodes servers, they need to be re-acknowledged.

https://t.me/learningnets

Creating Server Pool Qualification

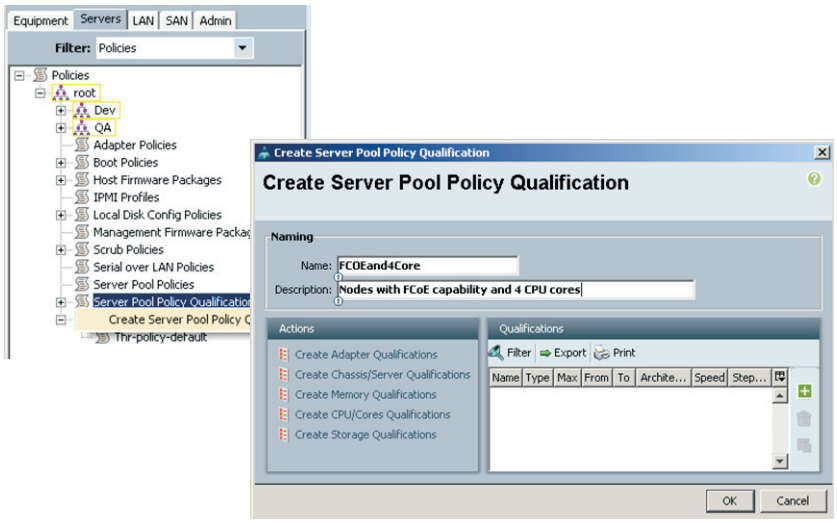

© 2009 Cisco Systems, Inc. All rights reserved.

A Server Pool Qualification is a named object that can combine (within one named qualification) various types of parameters (such as adapter card, memory, CPU, and storage). The figure shows the creation of a qualification named “FCOEand4Core” which will contain both adapter and CPU parameters.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Server Pool Qualification (Cont.)

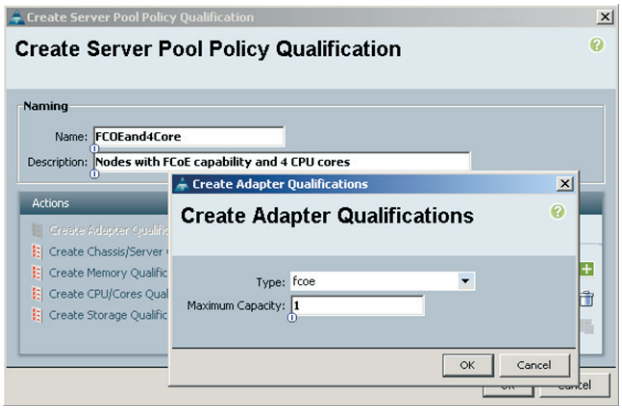

- This example requires an adapter capable of FCoE
- Multiple qualifications can be combined in the same policy

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows how to create a qualification for adapters. In this example, a qualification that is capable of FCoE is created.

https://t.me/learningnets

Creating a Server Pool Policy

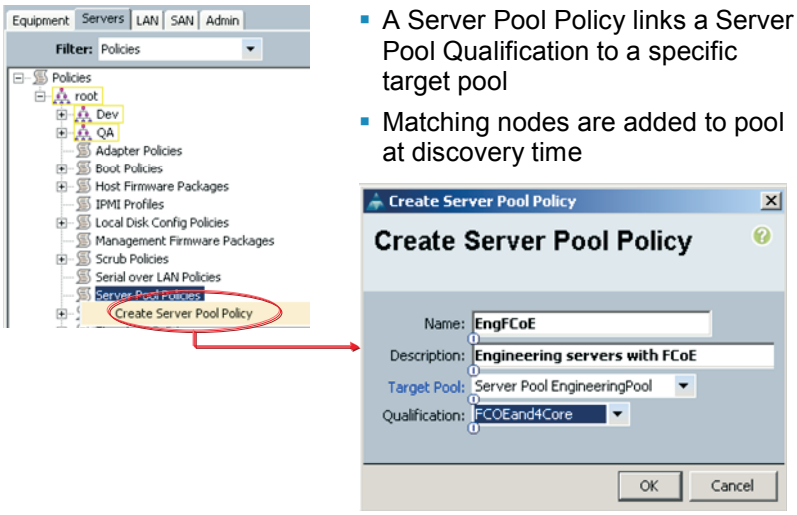

- A Server Pool Policy links a Server Pool Qualification to a specific target pool
- Matching nodes are added to pool at discovery time

© 2009 Cisco Systems, Inc. All rights reserved.

A Server Pool Policy ties together a named blade server qualification, such as “FCOEand4Core”, as shown, with an existing server pool. As compute nodes are discovered, those that meet the qualification will automatically be put in the pool.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

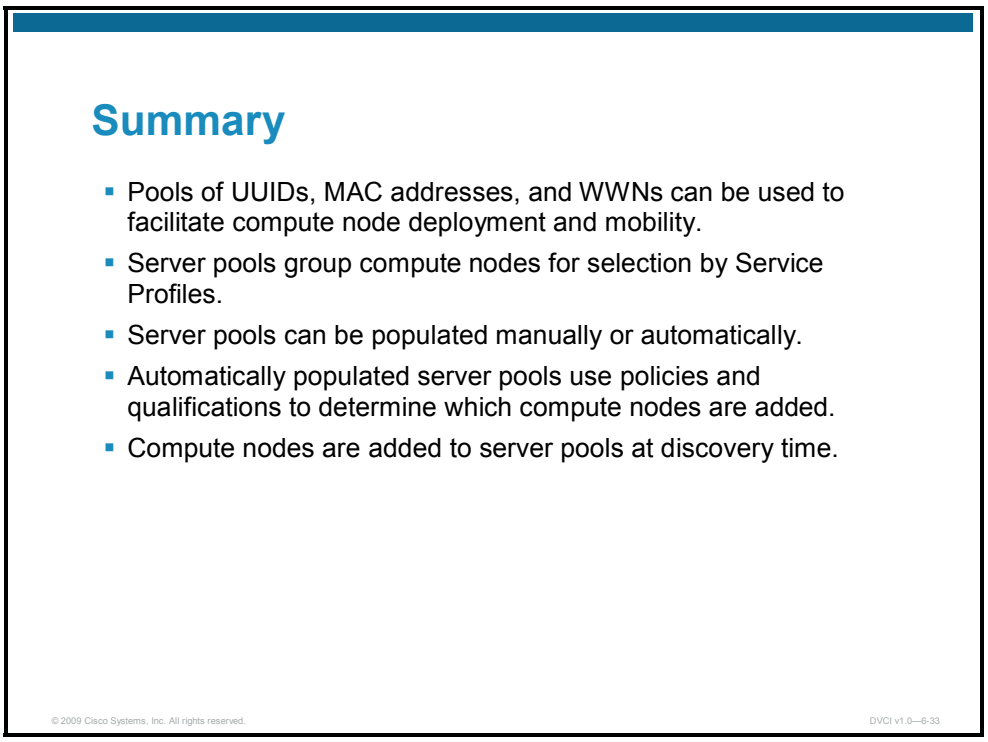

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Creating Service Profiles

Overview

In this lesson you will learn how to create Service Profiles on the Cisco UCS platform, including the basic opt-in model and the logical server model.

Objectives

Upon completing this lesson, you will be able to describe how to create Service Profiles on the Cisco Unified Computing System (Cisco UCS) platform. This includes being able to meet these objectives:

- Review the purpose and contents of a Service Profile
- Describe the two primary opt-in models of Cisco UCS
- Use the simple Service Profile wizard to create a Service Profile
- Use the expert wizard to create a Service Profile
- Describe the use of Service Profile templates
- Describe Scrub Policies and how they are used

https://t.me/learningnets

Service Profile Review

This topic reviews the contents and purpose of a Service Profile.

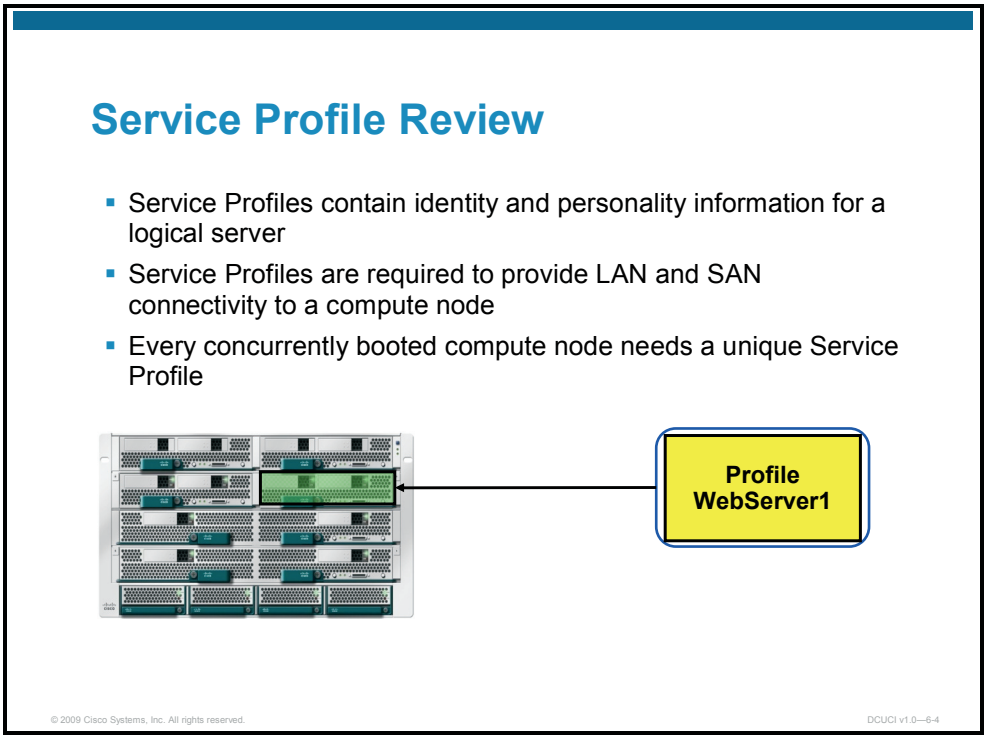

As described in a previous lesson, Service Profiles contain the identity and personality information for a logical server. Service Profiles abstract the identity and behavior of a server from the physical device on which it resides. In the Cisco UCS, every compute node that is running concurrently requires its own Service Profile. No LAN or SAN access is possible without a Service Profile.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

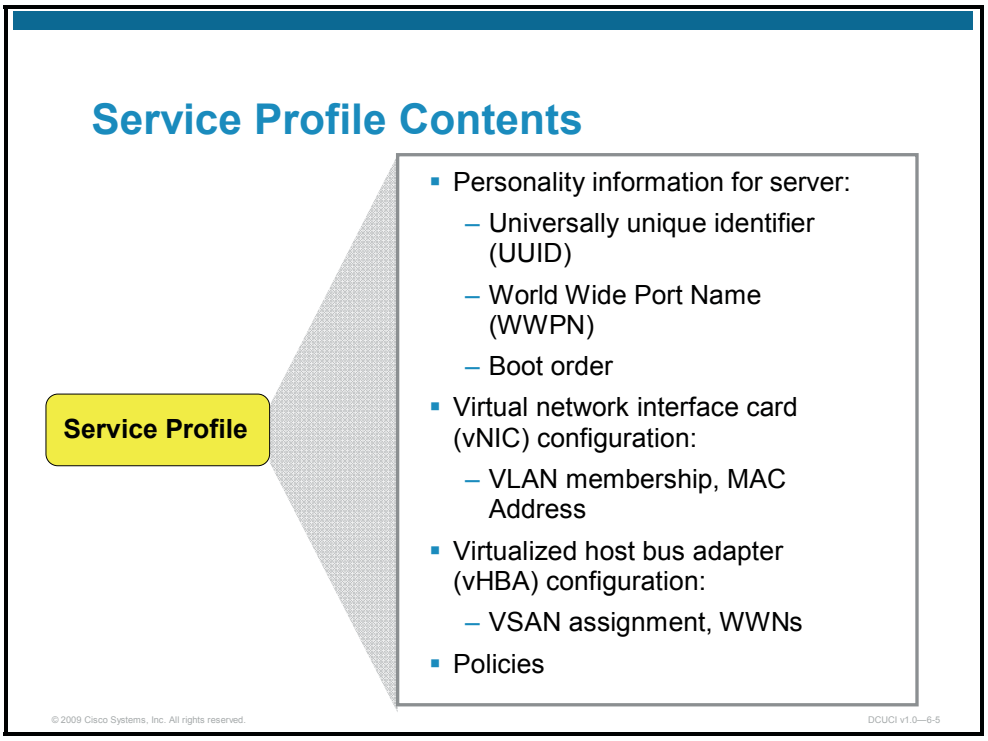

The Service Profile represents a logical view of a server without any ties to a specific physical device. The profile object contains the server personality. This identity contains all of the unique information for that server including, MAC, WWPN, UUID, boot order, and so on. Each profile can be associated with a single blade server at any given time.

Service Profiles facilitate the server mobility provided by UCS. Mobility is the ability to transfer server identity seamlessly between compute nodes.

In environments where blades will be managed as traditional individual servers, Service Profiles are still required. Service Profiles provide LAN and SAN connectivity configuration. Configuring Service Profiles in this fashion is similar to the need to configure individual LAN and SAN ports for traditional rack servers.

https://t.me/learningnets

Opt-In Models

This topic describes the two primary opt-in models of Cisco UCS.

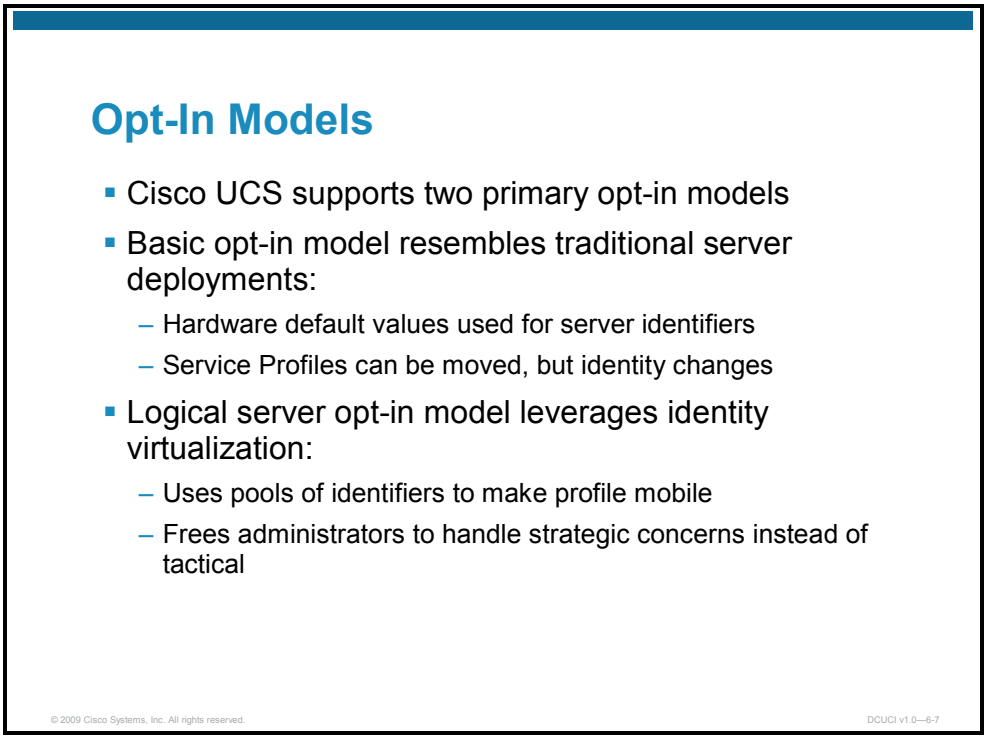

In the basic opt-in model, hardware values are used for identifiers. This mirrors the traditional method of deploying servers. In this model, Service Profiles can be moved, but the identifiers seen by the host operating system and network will change.

The logical server opt-in model uses virtualized identifiers, possibly from pools, to allow the Service Profile to be moved between physical compute nodes without the identifiers or personality changing.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Review of Mobile Service Profiles

- To simplify mobile Service Profiles, use pools of identifiers:
  - UUID pools used for compute node BIOS
  - MAC pools used for vNICs
  - World wide name (WWN) pools used for vHBAs

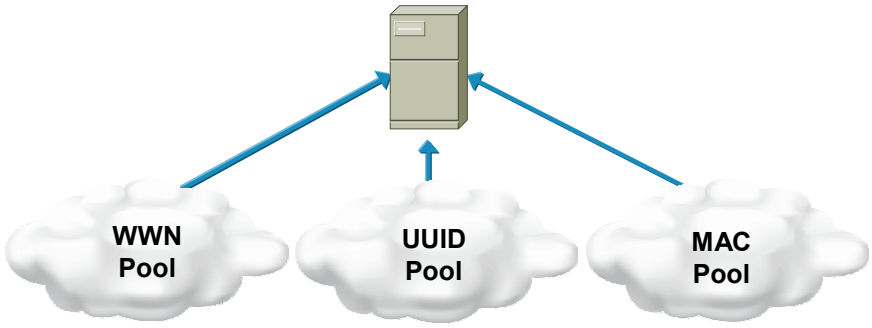

© 2009 Cisco Systems, Inc. All rights reserved.

As discussed in a previous lesson, the use of pools for identifiers makes the movement and mobility of Service Profiles easier. With virtualized identifiers, traditional problems with issues such as software licensing, network and SAN configuration, and so on, can be avoided.

https://t.me/learningnets

Service Profile Wizard

This topic describes how to use the wizard to create a Service Profile.

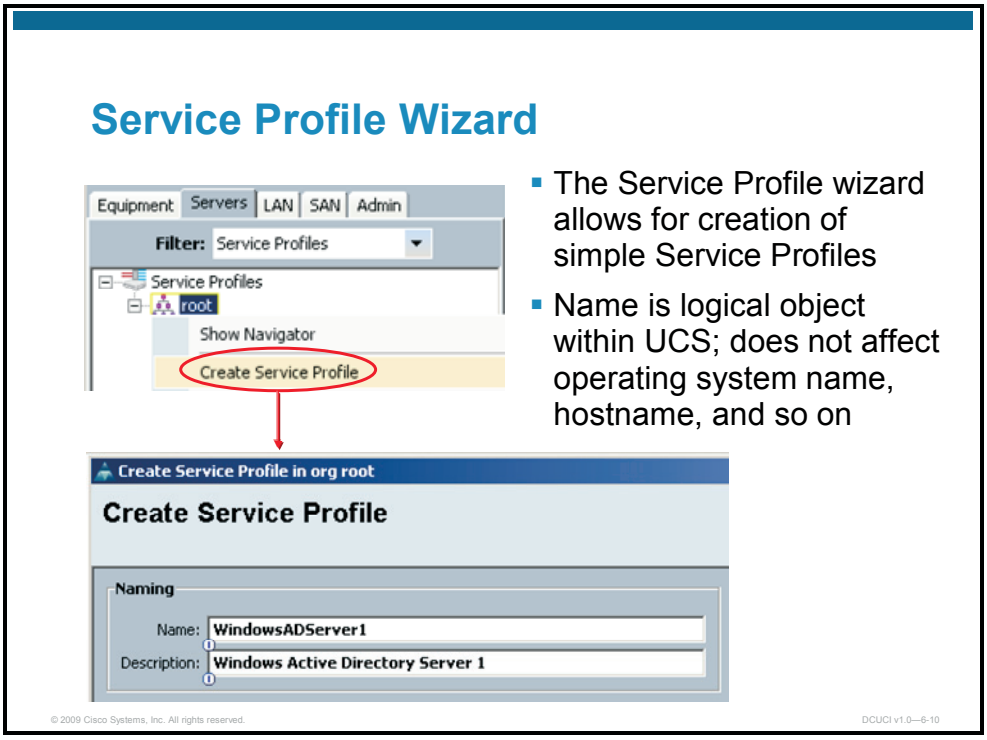

The wizard for creating Service Profiles works well for the basic opt-in model. The wizard will produce a Service Profile with basic characteristics and hardware-derived identifiers. If you want to customize the Service Profile after the wizard is finished (for example, to add more vNICs, or attach vNICs to different or multiple VLANs, and so on) you can still do so.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Service Profile Wizard - Connectivity

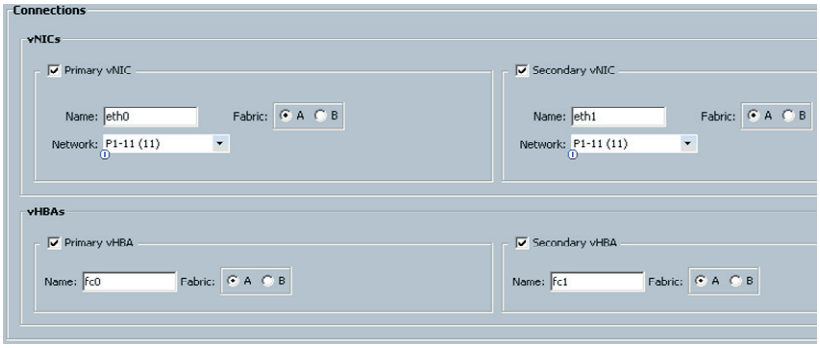

- The simplified wizard allows for the creation of up to two vNICs and up to two vHBAs
- Only one VLAN can be chosen per vNIC
- vHBAs will use default VSAN

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the options for vNIC and vHBA connectivity available in the wizard. Note that only one VLAN can be chosen per vNIC, and no capability for 802.1q VLAN trunking is provided. These configuration options can be changed later once the wizard has finished. Also note that vHBA configuration does not allow the selection of a VSAN. The default VSAN object is automatically chosen.

https://t.me/learningnets

Service Profile Wizard – Boot Order

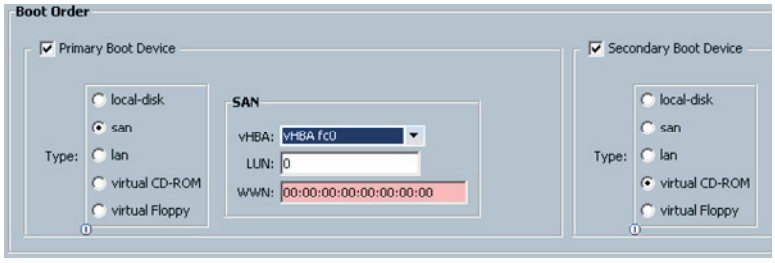

- Only two boot devices may be specified
- If SAN is specified as a boot device, boot target information must be provided

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the configuration options for boot order in the wizard. Up to two boot devices can be specified.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Service Profile Wizard – Server Association

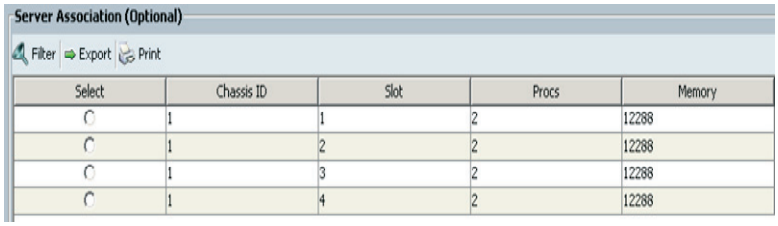

- A physical compute node can be selected at time of profile creation
- If selected, Service Profile will automatically be associated with selected compute node
- Can be associated manually later

© 2009 Cisco Systems, Inc. All rights reserved.

Finally, the wizard allows association with a specific compute node, if desired. Alternately, the profile may be manually associated with a compute node later.

https://t.me/learningnets

Service Profile Expert Wizard

This topic describes how to use the expert wizard to create a Service Profile.

The expert wizard allows a greater degree of customization of the Service Profile at the time of creation. Profiles created using either wizard can be customized after creation in the same manner.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard - UUID

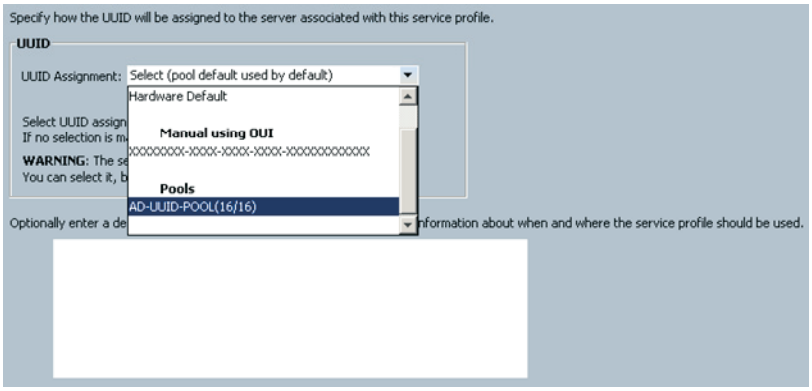

- UUID can be configured manually, chosen from a pool, or use the hardware default (derived) value
- An optional description can be added to the Service Profile

© 2009 Cisco Systems, Inc. All rights reserved.

Once the expert wizard is started, provide a name and optional description for the Service Profile.

https://t.me/learningnets

Expert Wizard – Local Storage

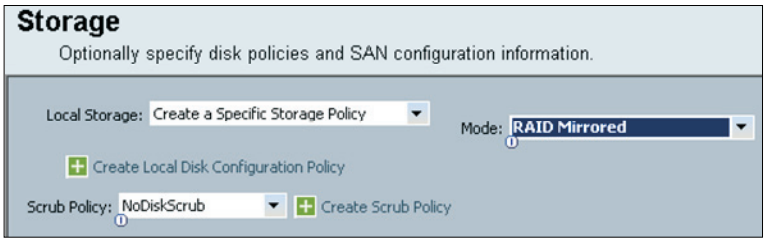

- Local Storage Policies and Scrub Policies can be optionally configured
- Local Storage Policies determine how the local disks on the compute node should be configured
- Scrub Policies determine what actions (if any) should be taken against the local disks upon disassociation
  - BIOS settings
  - Local disk contents

© 2009 Cisco Systems, Inc. All rights reserved.

Local Storage Policies determine how the local disks on the compute node should be configured. The policy can specify that the local disks not be used, and RAID-0 Striping, or RAID-1 Mirroring.

The Scrub Policy defines what actions should be taken when this Service Profile is disassociated from a compute node, such as erasing the local disks or BIOS settings.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard – SAN Storage

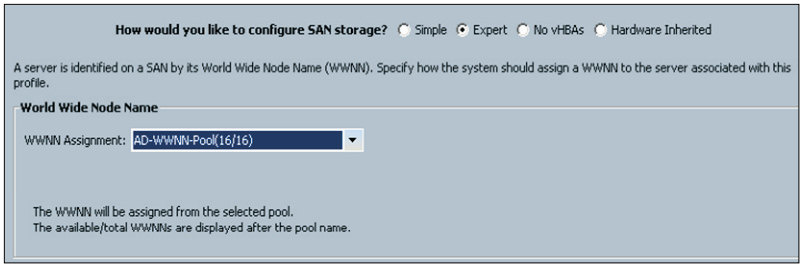

- The WWN can be assigned manually or from a pool
- Simple configuration is the same as the simplified wizard

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the expert options for configuring the WWNN for the Service Profile. The simple options are the same as the simplified wizard shown earlier.

https://t.me/learningnets

Expert Wizard – vHBAs

- As many vHBAs can be created as desired
- Service Profile can only be associated with physical hardware that can support the number of vHBAs:
  - No vHBAs supported on Oplin-based mezzanine cards
  - 2 vHBAs supported per Menlo-based mezzanine card
- Larger numbers of vHBAs require virtual interface card

© 2009 Cisco Systems, Inc. All rights reserved.

In the expert wizard, no vHBAs are created by default. Use the **Add** button to add vHBAs, as necessary.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard – vHBAs (Cont.)

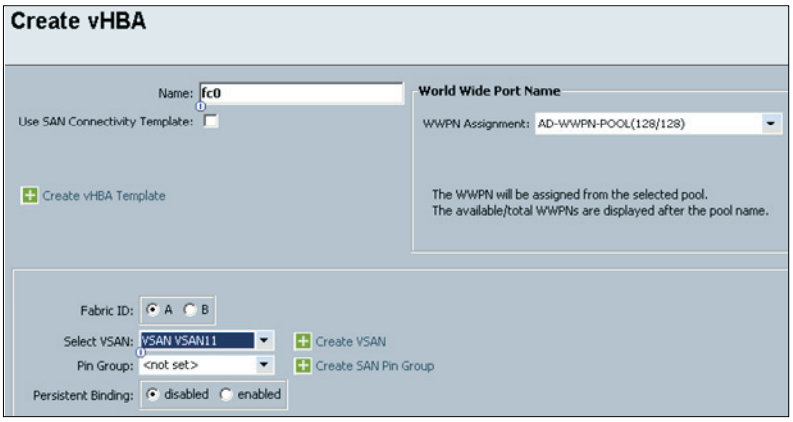

- WWPN can be assigned manually, or from a pool
- Each vHBA is associated with one fabric, and one VSAN
- Pin Group can be chosen, if desired

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the configuration options available when creating a vHBA in the expert wizard. The WWPN can be selected automatically from a pool, or assigned manually. Each vHBA is associated with one VSAN and one Fabric Interconnect.

https://t.me/learningnets

Expert Wizard – vNICs

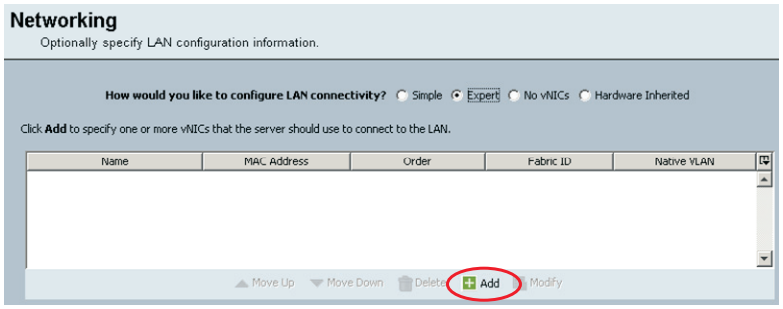

- Expert configuration is similar to storage configuration
- Add vNICs as necessary

© 2009 Cisco Systems, Inc. All rights reserved.

vNICs are added in the same manner as vHBAs. Click the **Add** button to create a vNIC.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard – Single VLAN vNIC

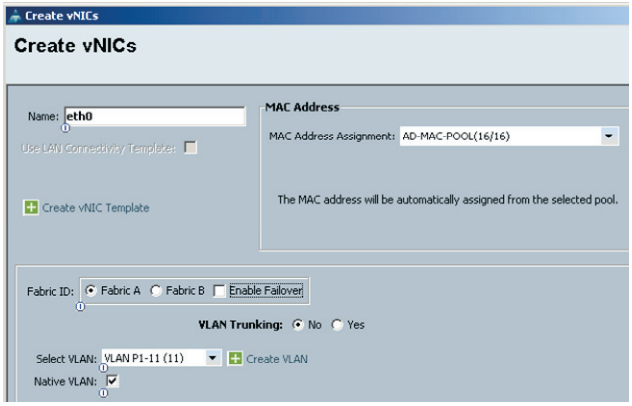

- Still need to specify Native VLAN checkbox if non-tagged Ethernet is desired

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the configuration options when creating a vNIC. Select the MAC address pool to use, or assign one manually. This example shows the creation of a vNIC that will receive only one VLAN as a Native or non-tagged VLAN.

https://t.me/learningnets

Expert Wizard – Trunked VLAN vNIC

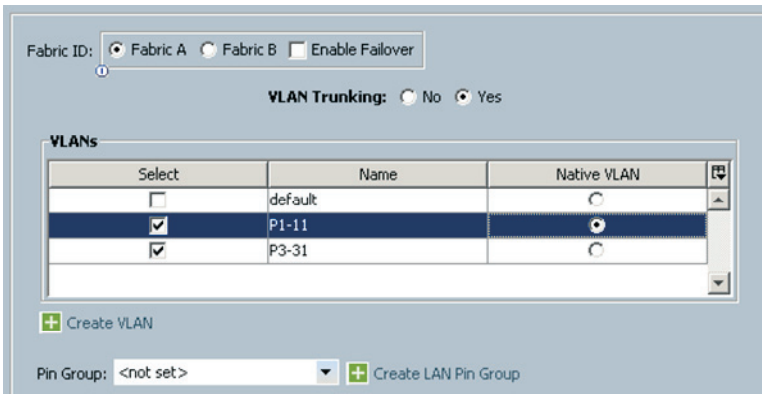

- Multiple VLANs can be sent to vNIC as standard 802.1q VLAN trunk
- One VLAN may optionally be designed as Native and sent untagged

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the configuration options when creating a vNIC. Select the MAC address pool to use, or assign one manually. This example shows the creation of a vNIC that will receive an 802.1q VLAN trunk. In this example, two VLANs will be trunked to this vNIC, with one of them untagged as the Native VLAN.

https://t.me/learningnets

Expert Wizard – Boot Order

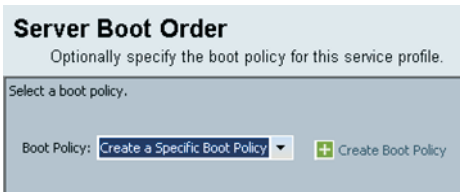

- Boot order can use policy or defined for this specific profile
- Order can include local storage or SAN resources

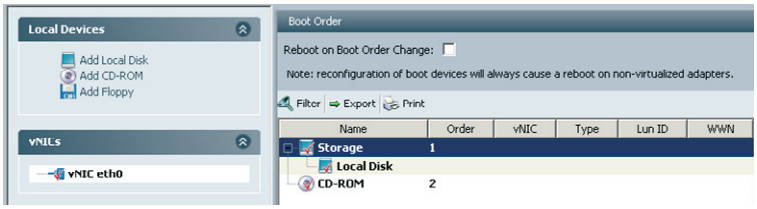

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the options for specifying the boot order for the compute node. Local disk, or virtual CD-ROM and floppy drives (discussed later in this module) are shown here. Booting from SAN is also supported and shown later in this lesson.

https://t.me/learningnets

Expert Wizard – SAN Boot

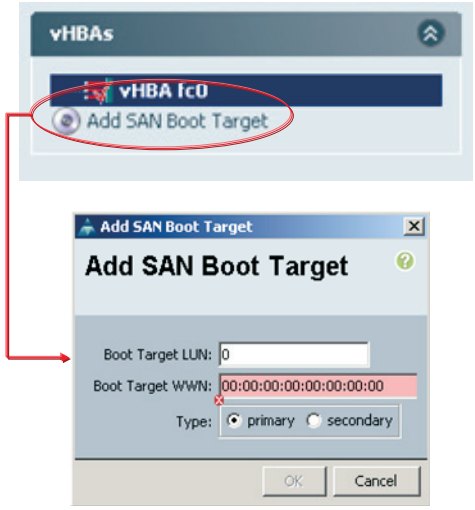

- When SAN boot is selected, you must define the SAN Boot Target
- This WWN is the WWNN or WWPN of the storage device's access port
- The LUN number will depend on the configuration of the storage device

© 2009 Cisco Systems, Inc. All rights reserved.

To configure a compute node to boot from SAN, first add the vHBA to the boot order (as shown previously) and then click the **Add SAN Boot Target** link. In the following dialog box, provide the boot target WWN and LUN number.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard – Server Assignment

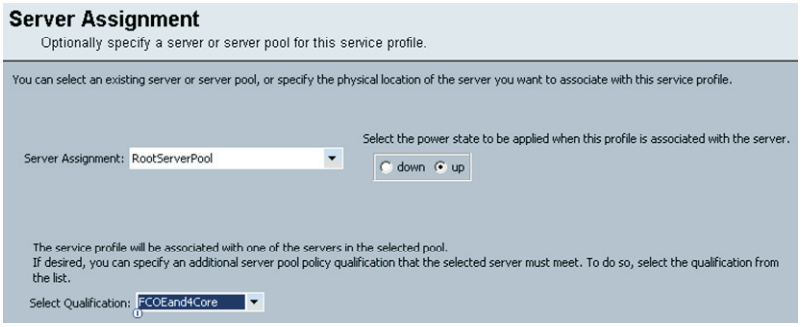

- The profile can be associated with a specific compute node, a server pool, or manually associated later
- Note that an additional qualification policy can be selected, for example when selecting from a manual pool

© 2009 Cisco Systems, Inc. All rights reserved.

The Service Profile can optionally be associated with a compute node or server pool. Additionally, a qualification policy can be selected to further refine the compute node selected.

https://t.me/learningnets

Expert Wizard – Firmware Selection

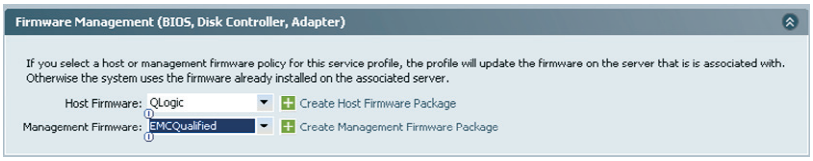

- Optionally, set which firmware versions will be applied to compute node at association
- Useful when deploying certified firmware combinations to meet support requirements

© 2009 Cisco Systems, Inc. All rights reserved.

If desired, the host and management firmware packages (discussed in a later module) can be selected and applied to the target compute node at the time of association.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Expert Wizard – Operational Policies

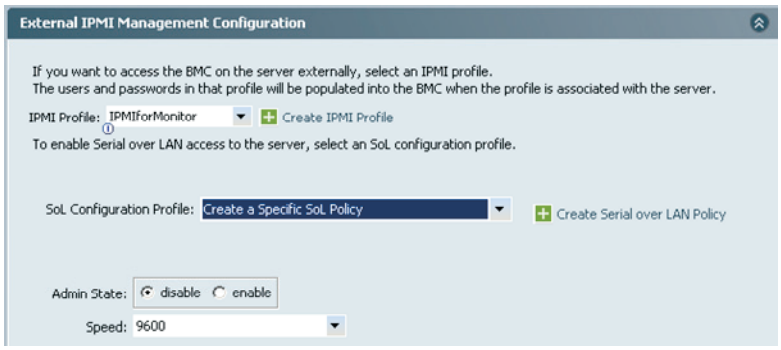

- Optionally, set Intelligent Platform Management Interface (IPMI), or Serial-over-LAN policies
- Both policy types are discussed in a later lesson

© 2009 Cisco Systems, Inc. All rights reserved.

Finally, policies can be selected for IPMI, or Serial-over-LAN, to be applied to the compute node. Both policy types are discussed in a later lesson.

https://t.me/learningnets

Considerations for SAN Boot

- Storage arrays may have differing requirements for SAN boot
- LUNs may be Active/Active or Active/Passive on redundant controllers
- May fail to boot on Active/Passive (or “Trespass”) type arrays
- Use Primary and Secondary boot targets for Active/Passive arrays
- Consult with your storage manufacturer for more details about your particular array
- Even arrays that claim to be Active/Active may require secondary target configuration

© 2009 Cisco Systems, Inc. All rights reserved.

SAN boot configurations vary, depending on the target storage device. It is important to consult with the administrators of the storage environment to determine which targets and settings are appropriate for each compute node.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Profile Templates and Clones

This topic describes Service Profile templates and clones.

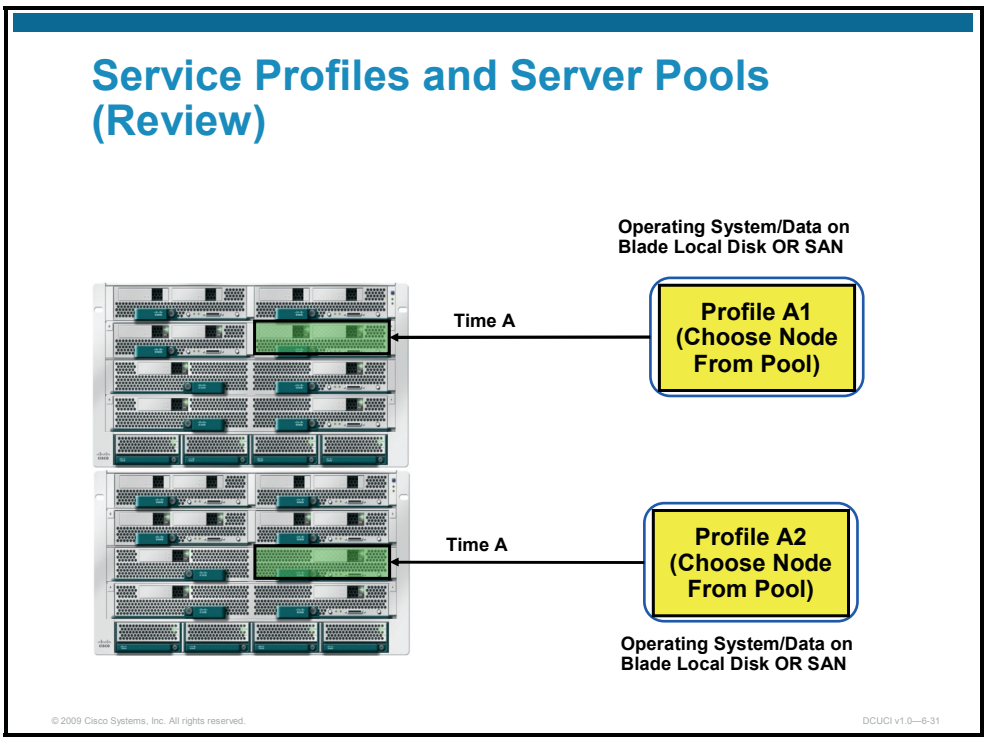

Every compute node that is booted concurrently must have its own Service Profile. Templates allow for the creation of multiple similar Service Profiles.

https://t.me/learningnets

Service Profile Cloning

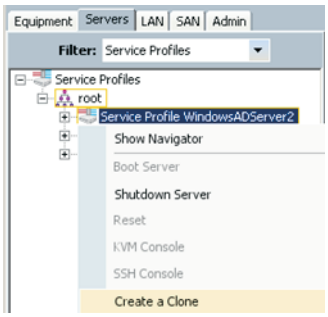

- All pool associations are replicated to a cloned template
- Specific new values for MAC, WWN will be immediately assigned to the profile from the appropriate pool

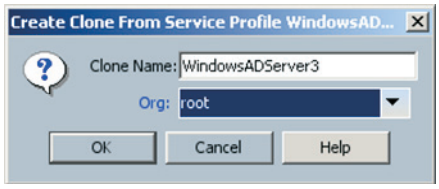

© 2009 Cisco Systems, Inc. All rights reserved.

If a profile is associated with pools (associated with a server pool, or its identity is associated with MAC, WWN, and UUID suffix pools), the resulting clone will have the same pool associations as the original profile:

- If the original is associated with the server pool, the clone will also be associated with the server pool, and a blade server will be selected for association immediately on creation of the cloned profile.
- If the original is associated with (MAC, WWN) pools, specific values will immediately be assigned out of the pool as the clone is created.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating a Template

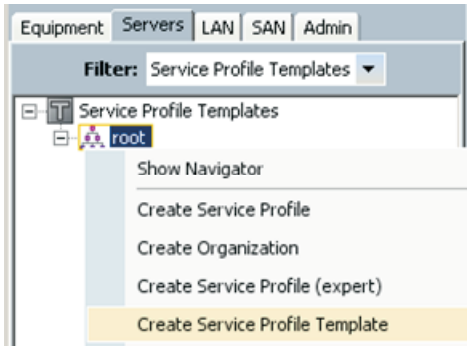

- Process is similar to Service Profile creation
- Can choose server pool (but not individual blade)
- Can associate virtual adapters (vNIC, vHBA) with MAC and WWN pools

© 2009 Cisco Systems, Inc. All rights reserved.

The process of creating a template is nearly identical to that of creating a blade server profile. You can associate the template to a server pool—in this case, a profile that is created as an instance of the template will have a blade server selected from the pool immediately. Once again, because a profile is created as an instance of the template, actual values will immediately be selected from the pools.

https://t.me/learningnets

Service Profile Templates

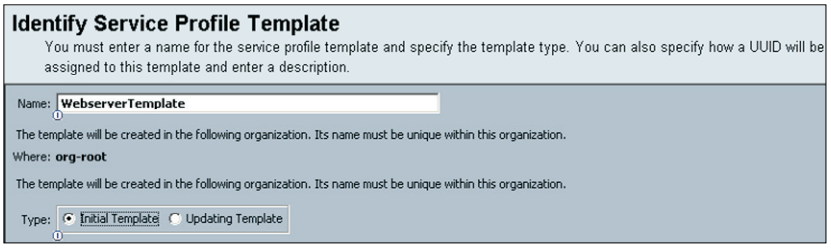

- Initial templates:
  - Updates to template do not affect child profiles
- Updating templates:
  - Updates to template are propagated to child profiles
- Remaining steps use Service Profile (expert) wizard

© 2009 Cisco Systems, Inc. All rights reserved.

Templates are an alternate method for creating large numbers of similar profiles.

Templates are available in two varieties:

- **Initial** templates are used to set the initial properties of an instance of the template, but modification of the template does not modify existing instances. A Service Profile created from an Initial template is not affected if the parent template is later changed.
- An **Updating** template is a slight variation. Modification of the template does modify existing child instances of the template. A Service Profile created from an Updating template will receive any changes made to the parent template.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Service Profiles From Template

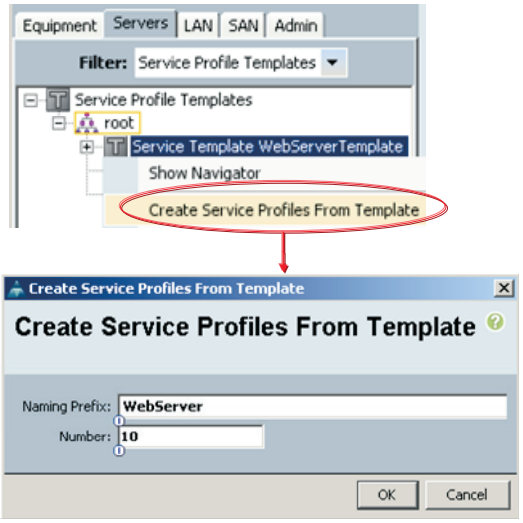

- Creating from template allows creation of multiple profiles
- Provide naming prefix and number of desired children profiles
- Example would create WebServer1, Webserver2, and so on

© 2009 Cisco Systems, Inc. All rights reserved.

After a template is created, you can create as many instances of the template as desired. Each instance will be a separate logical Service Profile.

A template is named using the specified prefix and a numerical suffix.

https://t.me/learningnets

Scrub Policies

This topic describes Scrub Policies and how they are used in Cisco UCS.

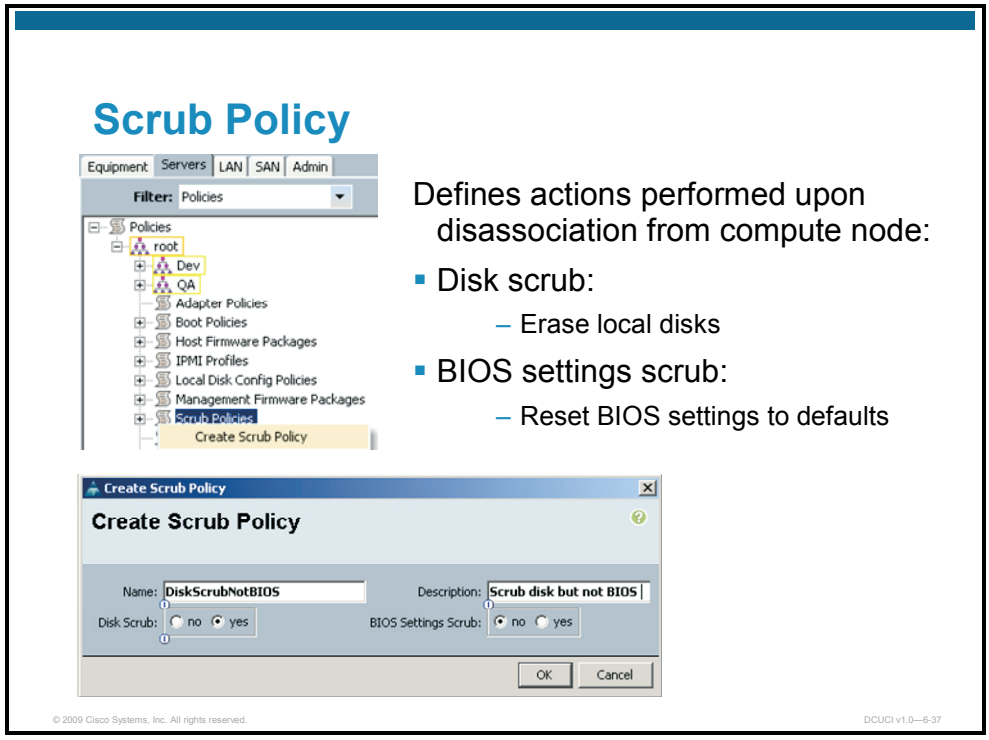

Scrub Policies can be defined to erase local disks or BIOS settings when disassociating a Service Profile from a compute node.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

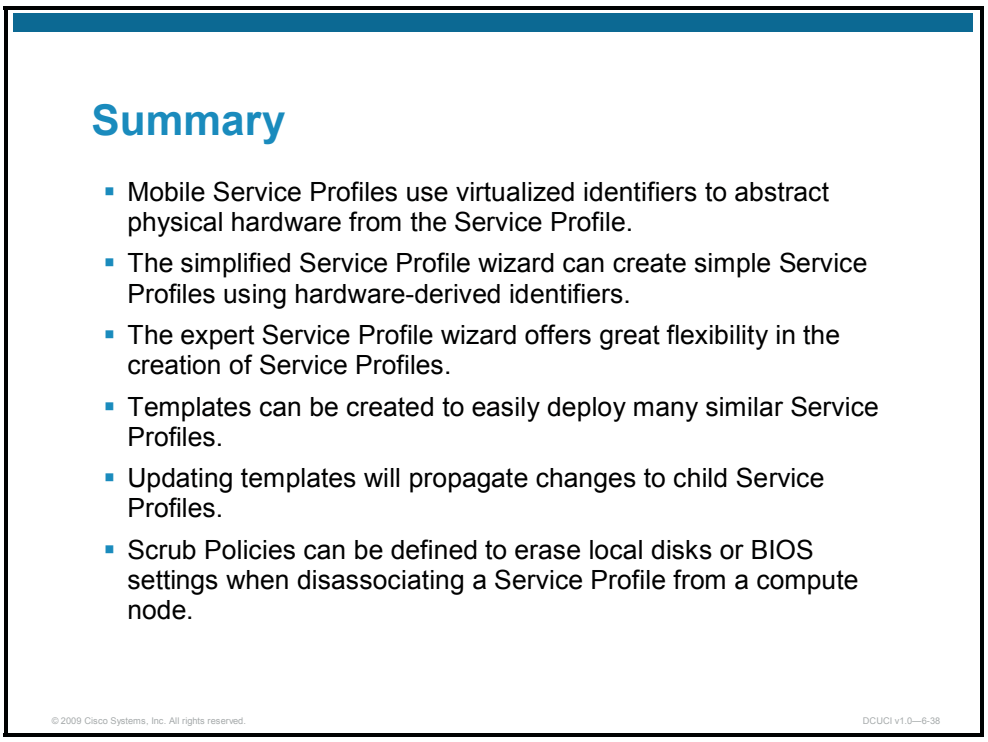

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Managing Service Profiles

Overview

This lesson discusses the full implementation of the stateless computing model using Service Profiles.

The discussion is divided into two themes:

- Mobile (re-locatable) Service Profiles; the same logical server can be booted on different blade servers at different times. When a blade server is associated with a Service Profile, it inherits all of its identity and boot information from the profile. Booting off a SAN logical unit number (LUN) may be used as an easy way to achieve statelessness of the operating system.
- Server farms using server pools; multiple, very similar logical Service Profiles can be used to create servers in a farm. The actual blade server hardware is chosen from pools.

Objectives

Upon completing this lesson, you will be able to describe how to manage Service Profiles. This includes being able to meet these objectives:

- Describe the process of associating Service Profiles with compute nodes
- Describe the process for moving a logical server
- Use KVM and virtual media utilities to access a compute node

https://t.me/learningnets

Associating Service Profiles

This topic describes the process of associating Service Profiles with compute nodes.

Profile: Created But Not Associated

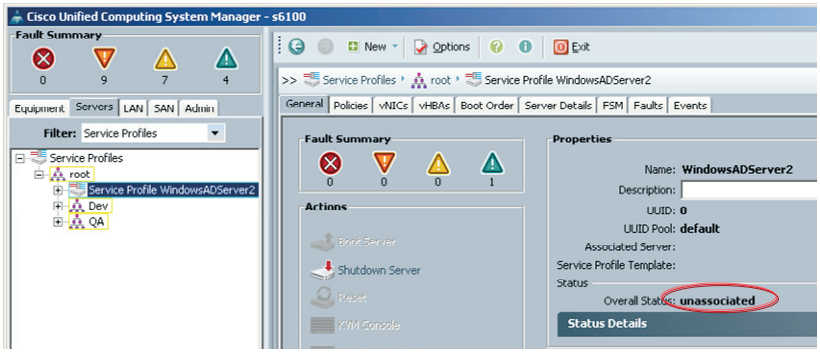

- This profile is not associated with any physical compute node

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows a Service Profile that is not associated with a physical compute node.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Associating Blade Server

- Two step process performed by **associate** command:
  - Assigning compute node to Service Profile
  - Compute node configuration through Cisco UCS Utility Operating System (UUOS)
- Fails if profile is incompatible with compute node chosen

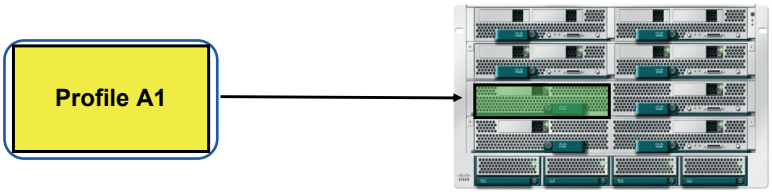

© 2009 Cisco Systems, Inc. All rights reserved.

During association, Cisco UCS Manager tries to assign the compute node to the Service Profile. This does not modify the compute node, but rather verifies that the compute node is compatible with the profile. If it is not compatible, the association action will fail.

Examples of incompatible blades:

- Too few physical network adapters for number of virtual network interface cards (vNICs)
- vNIC specified for failover, but compute node mezzanine adapter is Oplin-based
- vNIC specified for failover, but single Fabric Interconnect

Once the compute node is successfully assigned, the actual association configuration process begins. This involves the Cisco UCS Manager booting the Cisco UUOS to configure the compute node.

https://t.me/learningnets

Configuring Compute Node Through UUOS

- Process of applying profile to compute node:
  - Powers on compute node, if necessary
  - Boots or reboots compute node
  - Provides network boot of UUOS to compute node
  - Cisco UCS Manager communicates over network to UUOS
  - UUOS writes identity and boot order configuration
- Association process also configures the Cisco UCS fabric for proper VLAN and VSAN traffic to reach adapter on compute node

© 2009 Cisco Systems, Inc. All rights reserved.

UUOS serves as a pre-operating system configuration agent for the compute node. Cisco UCS Manager automatically powers on the compute node and sets it up for a PXE (network) boot. Cisco UCS Manager provides the DHCP response and UUOS download on a special dedicated management VLAN (4047) to the compute node. Once UUOS is booted, the Cisco UCS Manager communicates with it to configure the compute node according to the specifications of the profile.

When UUOS has finished configuring, the compute node boots using its BIOS boot order, which may or may not have been provided by the profile.

https://t.me/learningnets

Associating Compute Node with Profile

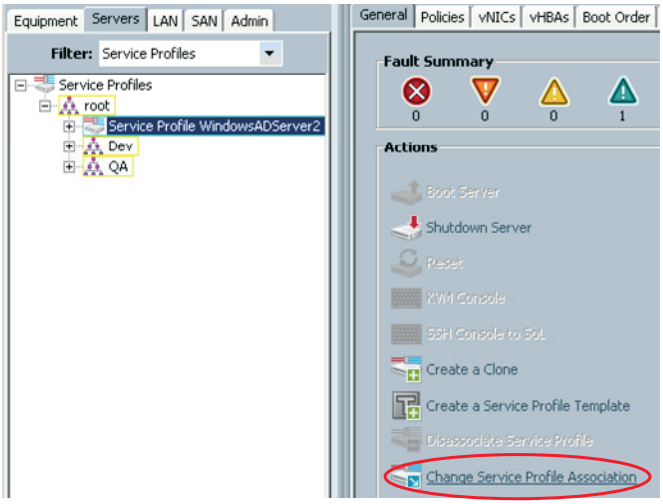

© 2009 Cisco Systems, Inc. All rights reserved.

As depicted in this example, click the **Change Service Profile Association** link to access a form to select the compute node. When committed, this will initiate the assignment and association of the profile.

https://t.me/learningnets

Associating the Service Profile

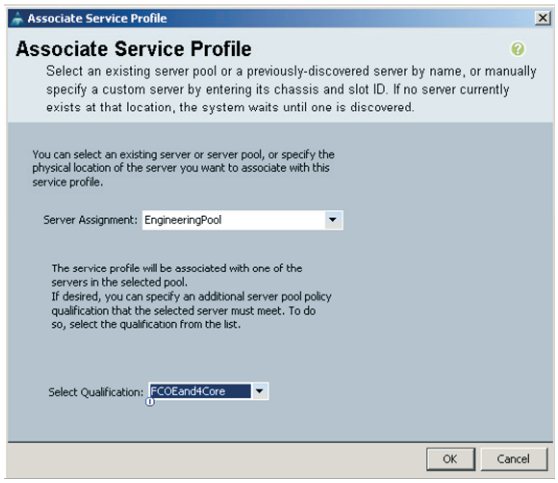

- Association can be to specific compute node or to pool
- Additional qualification can be defined to refine node selection

© 2009 Cisco Systems, Inc. All rights reserved.

In this example, a server pool is selected from which to select a compute node. Note that additional qualification can also be selected to further refine the compute node selection.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Assigned and Associating States of Profile

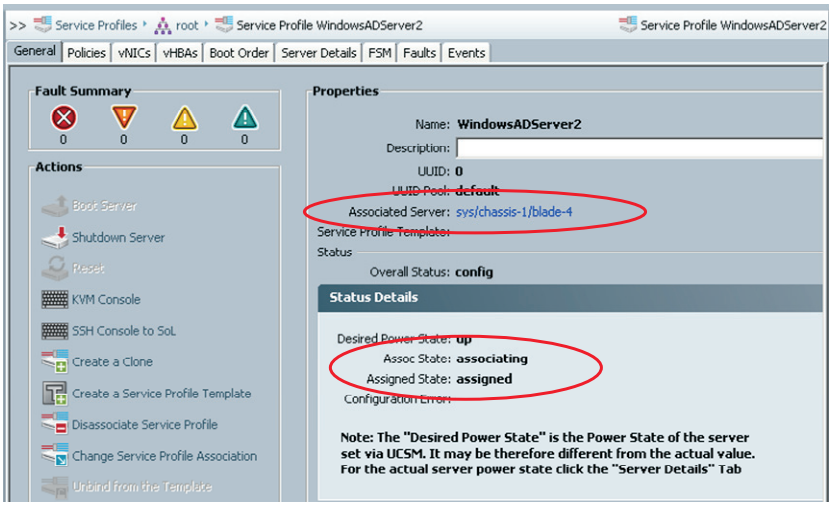

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS Manager GUI indicates that the selected compute node has been successfully assigned. There were no incompatibilities between the compute node and the profile. It is in the process of “associating”. Once the association process is complete, the ‘Assoc State” will change to “associated”.

https://t.me/learningnets

Moving a Logical Server

This topic describes the process for moving a logical server.

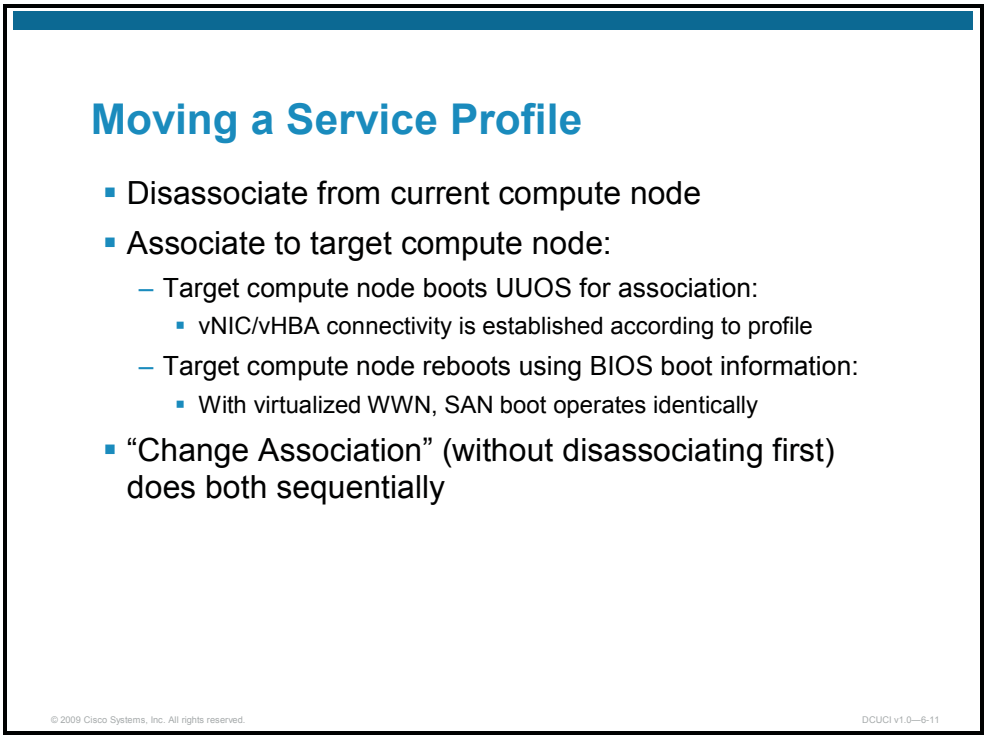

Note that true logical server mobility is likely to work flawlessly only when the source and destination compute nodes have the same type of mezzanine card (for example, both Cisco UCS CNA M71KR, or both Cisco UCS VIC M81KR).

While the same profile can easily associate with blade servers that have different types of adapters, the operating system for the profile will have storage and network adapters that are already associated with particular drivers. Unfortunately, the drivers are different for Cisco UCS CNA M71KR and Cisco UCS VIC M81KR.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Disassociating a Service Profile

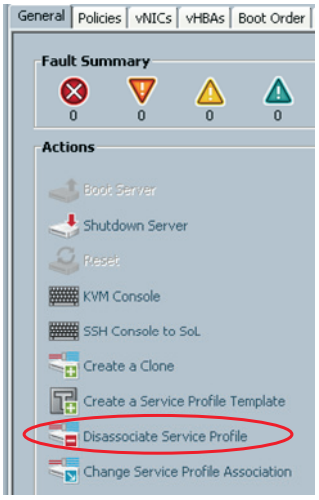

- Disassociating a Service Profile will:
  - Reboot the compute node
  - Boot UUOS to clear identity information
  - Execute Scrub Policy, if any
  - Power down the compute node
- Compute node is now available for assignment to another Service Profile

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the Disassociate Service Profile link.

https://t.me/learningnets

KVM and Virtual Media

The topic describes the use of the Keyboard/Video/Mouse (KVM) facility of Cisco UCS, and the use of virtual media.

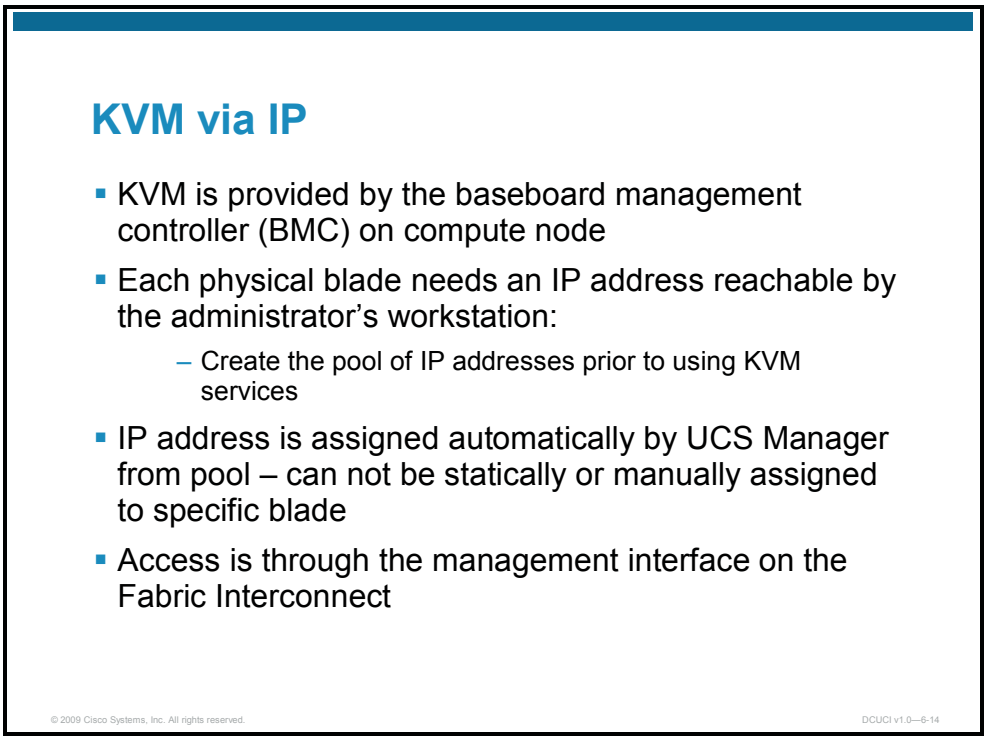

KVM functionality is provided by the BMC on the compute node. To use this functionality, the BMC needs an IP address that is routable on the management interface of the Fabric Interconnect. The IP address is assigned automatically to the BMC from a pool (described later in this lesson).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Server External Management Addresses

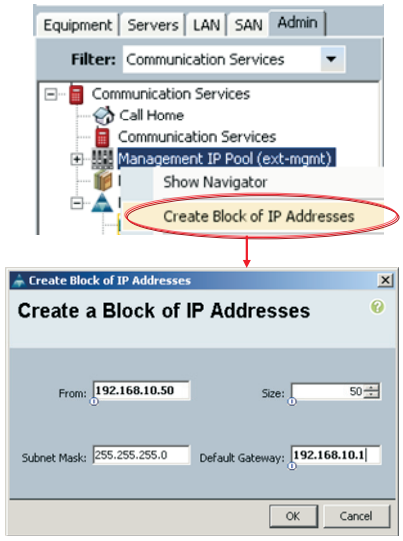

- IP address range must be routable on Fabric Interconnect management interface network
- Same address is used for optional Intelligent Platform Management Interface (IPMI) or Serial-over-LAN connectivity to compute node

© 2009 Cisco Systems, Inc. All rights reserved.

This series of slides shows the procedures for creating an external management IP pool for use with the KVM access method.

Step 1 From the navigation pane, click the **Pool ext-mgmt** object. Then, in the content pane, click the **Create block of IP Addresses** link. The "Create a Block of IP Addresses" panel appears.

Step 2 Enter the required values and click **OK**.

https://t.me/learningnets

Viewing KVM From Compute Node

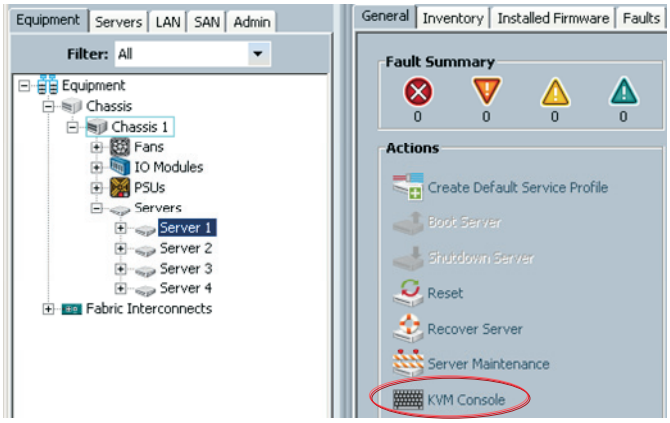

- Click the **Equipment** tab to access KVM for a specific physical compute node

© 2009 Cisco Systems, Inc. All rights reserved.

Access the KVM Console from the compute node view by clicking that link from the General tab of the content pane. After profile association is complete, you will also be able to access that same console from the Service Profile view.

The KVM Console is a Video-over-IP representation of the video output on the blade. The service is provided by the blade server's BMC via the external IP address only. You need not specify that IP address to bring up the KVM Console, but it must be up and reachable from the system where you are running the Cisco UCS Manager GUI.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Viewing KVM From Service Profile

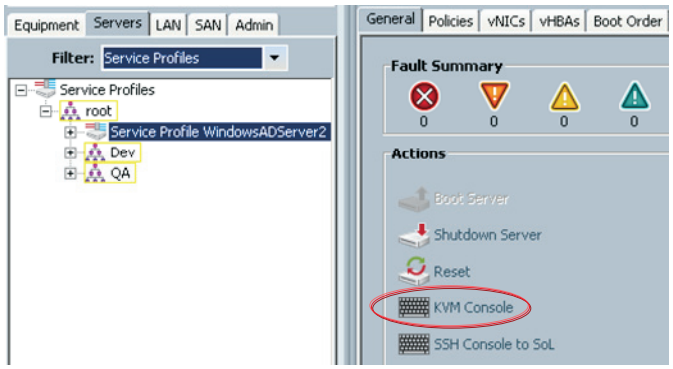

- Click the **Servers** tab to access KVM for the compute node to which this profile is associated

© 2009 Cisco Systems, Inc. All rights reserved.

You can also access the KVM Console from the Service Profile view, by clicking that link from the General tab of the content pane. Because the IP address is associated with the physical compute node and not the Service Profile, this method is preferred because it will always connect with the proper compute node, regardless of association.

https://t.me/learningnets

KVM Console

- KVM operates exactly as a physical keyboard and monitor would

© 2009 Cisco Systems, Inc. All rights reserved.

This slide shows the KVM Console. You will interact with it exactly as if it were the real server console.

This Cisco splash screen is generated by the compute node BIOS during the boot process.

https://t.me/learningnets

KVM: Accessing BIOS

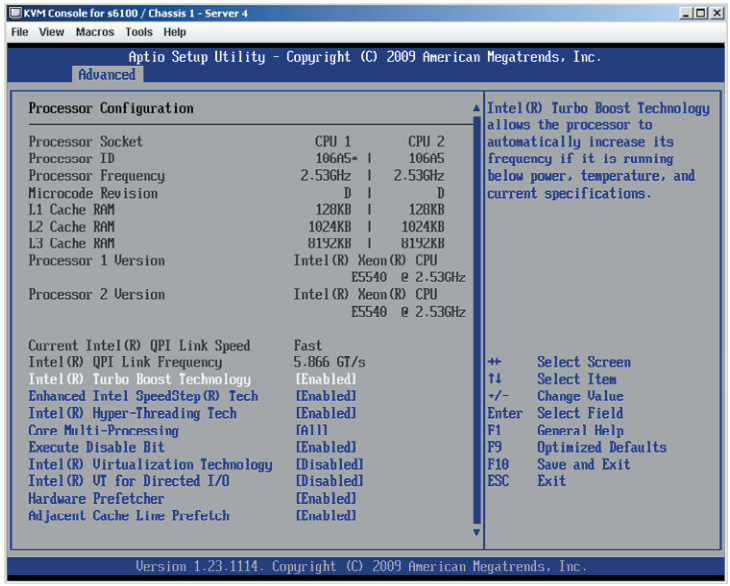

© 2009 Cisco Systems, Inc. All rights reserved.

You can access the BIOS at the beginning of the reboot process just as you do for any x86 server, by pressing F2 when prompted at the beginning of the boot process.

The figure shows the advanced processor configuration options in the BIOS.

https://t.me/learningnets

KVM Console Virtual Media

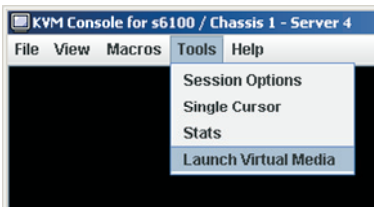

- Single virtual CD/DVD on blade:
  - Mapped to physical CD/DVD media on client
  - Mapped to .iso file on client
- Single virtual floppy on blade:
  - Mapped to physical floppy on client
  - Mapped to floppy image on client
  - Mapped to USB flash medium on client

© 2009 Cisco Systems, Inc. All rights reserved.

The KVM Console provides a sub-utility that allows you to use a physical CD/DVD, or an .iso file on the client (the machine where you are running the GUI and KVM Console) that presents itself to the blade as a virtual CD medium.

The compute node can boot from this medium for installations, or for later emergency repairs.

Note that if you are using VMware on the compute node, the Cisco UCS KVM Console virtual media is used by the compute node itself (for example, for installing VMware), not for the guest VM machines. The guest machines have their own Virtual Console, or virtual media methods, just as they do in any VMware installation.

https://t.me/learningnets

Virtual Media Manager

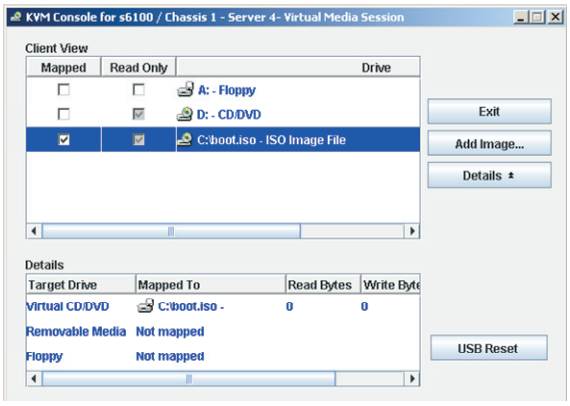

- Virtual Media Manager allows mapping of physical or logical devices to compute node
- Once presented to the compute node, access is exactly like a physical device

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows how to invoke the Virtual Media Manager from the KVM Console. This example uses an .iso file on the client as virtual media for the compute node.

https://t.me/learningnets

Booting a Blade from Virtual Media

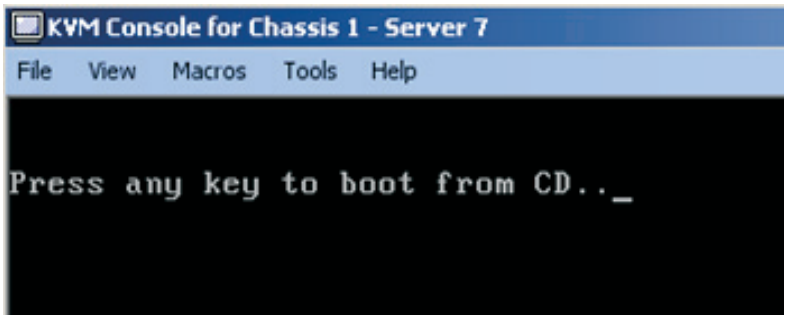

© 2009 Cisco Systems, Inc. All rights reserved.

Note that after we configure networking in the operating system, we are less likely to need the KVM Console, because we can simply access the operating system over the network (with Remote Desktop for Windows, for example).

https://t.me/learningnets

Summary

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

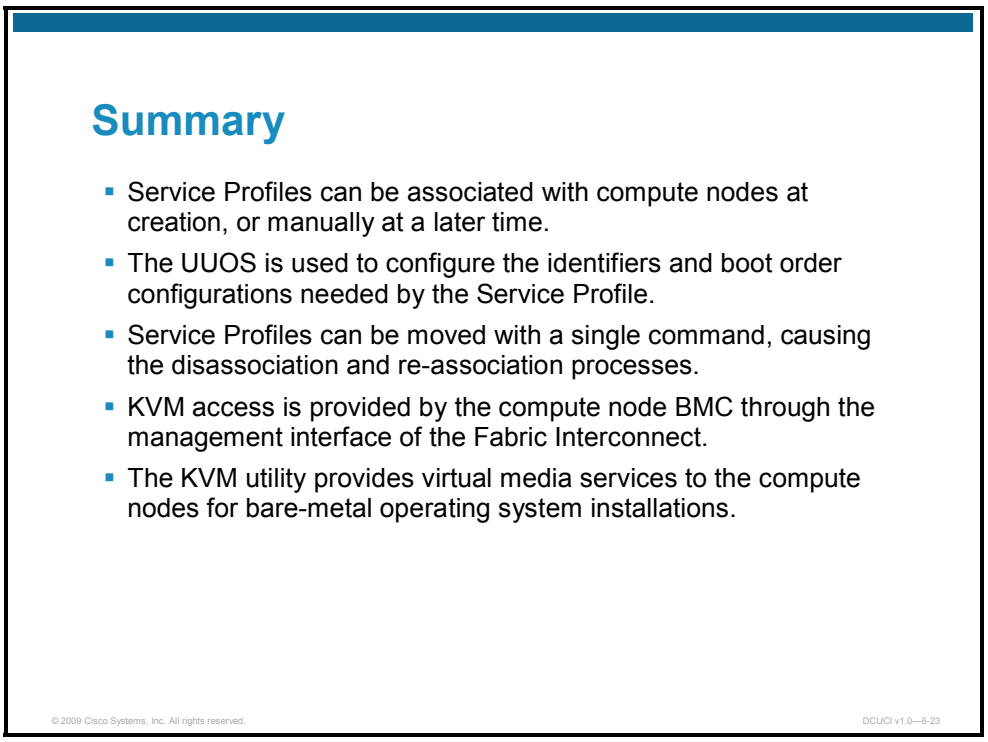

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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

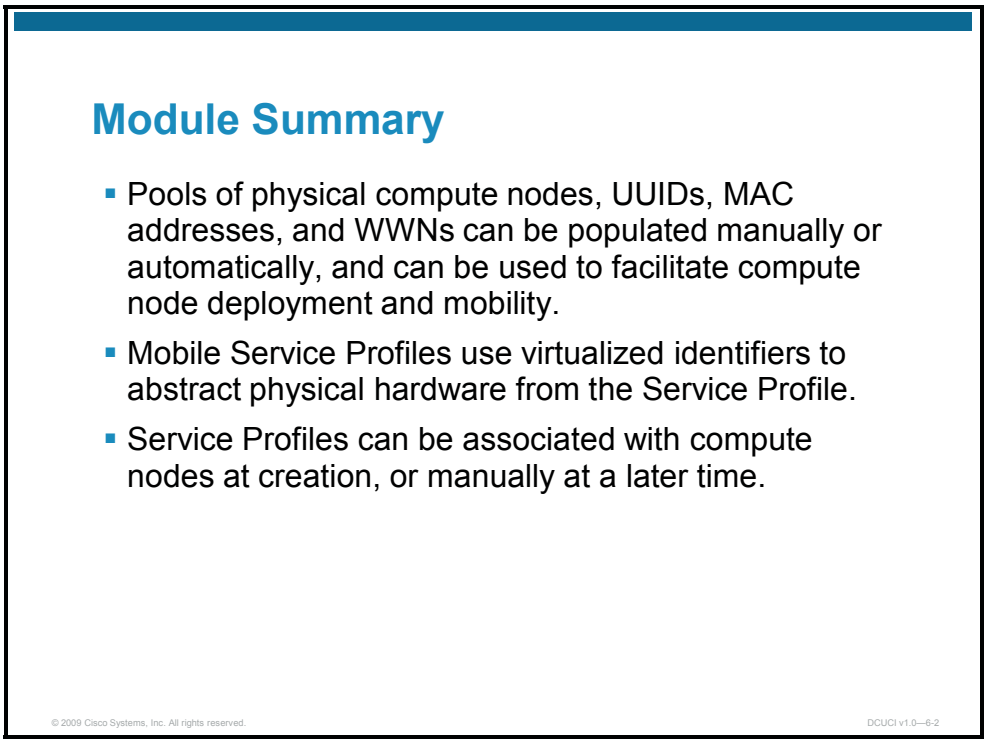

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

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 of the following is true about Server Pools? (Source: Identity and Resource Pools)

A) Server Pools belong to a locale
B) Server Pools can be automatically populated
C) Server Pools cannot contain compute nodes from multiple chassis
D) Server Pools are not inherited by child organizations

Q2) Which of the following statements about Service Profiles is correct? (Source: Creating Service Profiles)

A) A given Service Profile can be simultaneously booted on multiple compute nodes.
B) Service Profiles define the port security policies applied to the host's LAN and SAN ports.
C) Service Profiles are required to provide LAN and SAN connectivity to a compute node.
D) Service Profiles are required only if you are booting from a SAN.

Q3) Keyboard/Video/Mouse (KVM) access to a blade server is provided by which component? (Source: Managing Service Profiles)

A) Fabric Interconnect
B) Chassis Management Controller
C) IO Module
D) Baseboard Management Controller

https://t.me/learningnets

Module Self-Check Answer Key

Q1) B

Q2) C

Q3) D

© 2009 Firefly Communications, LLC

https://t.me/learningnets