DCICT

Introducing Cisco Data Center Technologies

Volume 2

Version 1.0

Student Guide

Text Part Number: 97-3182-01

https://t.me/learningnets



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Table of Contents

Volume 2

Cisco Data Center Storage Networking ***3-1***

Overview 3-1
Module Objectives 3-1

Comparing Storage Connectivity Options in the Data Center **3-3**

Overview 3-3
Objectives 3-3
Comparing Block- and File-Based Network Storage 3-4
NFS, Fibre Channel, iSCSI, and SCSI 3-7
Layered Architecture 3-8
Client/Server Architecture 3-9
Remote Procedure Call 3-10
Portmap 3-10
Automounter 3-11
NFSv2 3-13
NFSv3 3-13
NFSv4 3-13
iSCSI Concepts 3-25
iSCSI Node Names 3-26
Software-Based iSCSI Drivers 3-27
TCP and iSCSI Offload Engines 3-28
iSCSI Drivers and Offload Engines 3-29
SCSI Evolution from DAS to SAN 3-30
Summary 3-34

Describing Fibre Channel Storage Networking **3-35**

Overview 3-35
Objectives 3-35
Fibre Channel SAN Topologies 3-36
Fibre Channel Port Types 3-42
Fibre Channel Addressing 3-45
Switched Fabric Address Space 3-47
Fibre Channel Layered Model 3-48
FCNS and the FLOGI Process 3-49
Fibre Channel Zoning and LUN Masking 3-64
Summary 3-69

Verifying Fibre Channel Communications on Cisco MDS 9000 Series Multilayer Switches **3-71**

Overview 3-71
Objectives 3-71
Configure a Cisco MDS 9000 Series Multilayer Switch from the CLI Setup Script 3-72
Update the Cisco NX-OS on a Cisco MDS 9000 Series Multilayer Switch 3-78
Update Licensed Features on the Cisco MDS 9000 Series Multilayer Switch 3-84
Common Software Across All Platforms 3-85
Manual Installation 3-87
Obtaining the License Key File 3-87
Installing the License Key File 3-88
Verify Initiator and Target Fabric Login 3-90
Verify Fibre Channel Zoning on a Cisco MDS 9000 Series Multilayer Switch 3-100
Uses for Zoning 3-101
Collect Technical Support Data on a Cisco MDS 9000 Series Multilayer Switch 3-103
Summary 3-104
Module Summary 3-105
References 3-106
Module Self-Check 3-107
Answers 3-110

https://t.me/learningnets

Cisco Data Center Unified Fabric 4-1

Overview 4-1
Module Objectives 4-1

Describing DCB 4-3

Overview 4-3
Objectives 4-3
Unified Fabric Benefits 4-4
FCoE Logical End Point 4-8
FCoE Node 4-9
Fibre Channel Forwarder 4-10
Discovery and Virtual Link Establishment 4-13
IEEE Standards That Enable FCoE 4-16
Priority Flow Control 4-18
Link-Level Flow Control 4-18
Enhanced Transmission Selection 4-21
DCB Exchange 4-23
Goals of DCBX 4-24
Summary 4-26

Identifying Connectivity Options for FCoE on the Cisco Nexus 5000 Series Switch 4-27

Overview 4-27
Objectives 4-27
SFP Modules 4-28
1000BASE-T SFP for Copper Networks 4-28
1000BASE-SX SFP for Multimode Fiber 4-28
1000BASE-LX/LH SFP for Multimode and Single-Mode Fiber 4-29
1000BASE-EX SFP for Long-Reach Single-Mode Fiber 4-29
1000BASE-ZX SFP for Long-Reach Single-Mode Fiber 4-29
1000BASE-BX10-D and 1000BASE-BX10-U SFP for Single-Fiber Bidirectional Applications 4-29
Cabling Requirements and Distance Limitations for Common SFP and SFP+ Transceivers 4-33
Connecting the Cisco UCS P81E Virtual Interface Card to Cisco Nexus 5500UP Unified Fabric Switches 4-34
Static Interface and Static vNIC or vHBA 4-36
Dynamic Interface and Dynamic vNIC 4-36
Fixed vEth Interface 4-36
Floating vEth Interface 4-37
Features and Benefits 4-38
Connecting the Cisco Nexus 5500UP Unified Fabric Switch to Northbound LAN and SAN Fabrics 4-43
Unified Fabric with FCoE: I/O Consolidation 4-43
Summary 4-45

Describing Enhanced FCoE Scalability with Cisco Nexus 2232 10GE Fabric Extenders 4-47

Overview 4-47
Objectives 4-47
Scaling the Data Center Virtualized Access Layer with the Cisco Nexus 2232 10GE Fabric Extenders 4-48
Cisco Nexus 2232 10GE Fabric Extender-to-Cisco Nexus 5500 Switch Connectivity 4-50
Adapter FEX on the Cisco Nexus 2232 10GE Fabric Extender 4-55
Verifying Adapter FEX on the Cisco Nexus 2232 10GE Fabric Extender 4-59
Summary 4-62
Module Summary 4-63
References 4-64
Module Self-Check 4-67
Module Self-Check Answer Key 4-69

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS ***5-1***

Overview 5-1
Module Objectives 5-1

Describing the Cisco UCS B-Series Product Family **5-3**

Overview 5-3
Objectives 5-3
Cisco UCS 6100 and 6200 Series Fabric Interconnects 5-4
Cisco UCS 6120XP 20-Port Fabric Interconnect 5-5
Cisco UCS 6140XP 40-Port Fabric Interconnect 5-6
Comparison of Cisco UCS 6120XP and Cisco UCS 6140XP Fabric Interconnects 5-7
Cisco UCS 6100 Series Expansion Modules 5-7
Cisco UCS 6248UP 48-Port Fabric Interconnect 5-8
Cisco UCS 6296UP 96-Port Fabric Interconnect 5-9
Comparison of Cisco UCS 6248UP and Cisco UCS 6296UP Fabric Interconnects 5-10
Cisco UCS 6200 Series Expansion Module 5-11
Cisco UCS 5108 Blade Server Chassis 5-12
Cisco UCS 5108 Blade Server Chassis: Rear 5-13
Cisco UCS I/O Modules 5-13
Power Supply Input Connectivity 5-14
Power Supply Output Connectivity 5-15
Cisco UCS B200 M3 Blade Server 5-16
Cisco UCS B200 M3 Blade Server Motherboard 5-17
Cisco UCS B230 M2 Blade Server 5-18
Cisco UCS B230 M2 Blade Server Motherboard 5-19
Cisco UCS B250 M2 Extended Memory Blade Server 5-20
Cisco UCS B250 M2 Extended Memory Blade Server Motherboard 5-21
Cisco UCS B440 M2 High-Performance Blade Server 5-22
Cisco UCS B440 M2 High-Performance Blade Server Motherboard 5-23
Comparison of Cisco UCS B-Series Blade Servers 5-24
Mezzanine Card Options for Cisco UCS B-Series Blade Servers 5-25
Cisco UCS Virtual Interface Card 1240 5-26
Cisco UCS M81KR Virtual Interface Card 5-27
Cisco UCS Nonvirtualized Adapter Portfolio 5-28
Memory Population Guidelines for Cisco UCS B-Series Blade Servers 5-29
Cisco UCS B200 M3 Blade Server Memory Population 5-30
Cisco UCS B-Series Blade Server Memory Population 5-31
Summary 5-32

Describing the Cisco UCS C-Series Product Family **5-33**

Overview 5-33
Objectives 5-33
Cisco UCS C-Series Product Family 5-34
Cisco UCS C200 M2 High-Density Rack Server 5-35
Cisco UCS C210 M2 General-Purpose Rack Server 5-37
Cisco UCS C220 M3 Rack Server 5-39
Cisco UCS C240 M3 Rack Server 5-41
Cisco UCS C250 M2 Extended-Memory Rack Server 5-43
Cisco UCS C260 M2 Rack Server 5-45
Cisco UCS C460 M2 High-Performance Rack Server 5-47
Cisco USC C-Series Comparison Summary 5-49
PCIe Adapter Options for Cisco UCS C-Series Rack Servers 5-50
RAID Adapter Options for Cisco UCS C-Series Rack Servers 5-53
Data Striping for Increased Bandwidth 5-60
Memory Population Guidelines for Cisco UCS C-Series Rack Servers 5-63
Cisco UCS C220 M3 Rack Server Memory Population 5-64
Summary 5-66

https://t.me/learningnets

Connecting Cisco UCS B-Series Blade Servers 5-67
Overview 5-67
Objectives 5-67
Chassis-to-Fabric Interconnect Physical Connectivity 5-68
I/O Module Architectures 5-76
Cisco Integrated Management Controller Chip on Cisco UCS B-Series Blade Servers 5-77
Three Basic Port Personalities in the Fabric Interconnect 5-78
Discovery Process 5-80
Summary 5-82
Setting up an Initial Cisco UCS B-Series Cluster 5-83
Overview 5-83
Objectives 5-83
Cabling a Cisco UCS Fabric Interconnect Cluster 5-84
Initial Setup Script for the Primary Peer 5-86
Initial Setup Script for the Secondary Peer 5-89
Verifying a Fabric Interconnect Cluster 5-90
Summary 5-93
Describing Cisco UCS Manager Operations 5-95
Overview 5-95
Objectives 5-95
Cisco UCS Manager 5-96
Layout of the Cisco UCS Manager GUI 5-99
Navigation Window Tabs 5-100
Device Discovery in Cisco UCS Manager 5-106
Chassis Discovery 5-106
Discovery of Server in Chassis 5-107
Verifying Device Discovery in Cisco UCS Manager 5-108
Summary 5-111
Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles 5-113
Overview 5-113
Objectives 5-113
Benefits of Stateless Computing 5-114
Using Identity Pools in Service Profiles 5-117
Using Service Profile Templates to Enable Rapid Provisioning and Consistent Application of Policy 5-125
Creation of Policies for Service Profiles and Service Profile Templates 5-130
Chassis and Blade Power Capping 5-132
Summary 5-134
Module Summary 5-135
References 5-136
Module Self-Check 5-139
Module Self-Check Answer Key 5-144

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module 3

Cisco Data Center Storage Networking

Overview

This module describes the function and operation of Cisco data center storage-networking technologies.

Module Objectives

Upon completing this module, you will be able to describe the function and verify correct operation of Cisco data center storage-networking technologies. This ability includes being able to meet these objectives:

- Compare iSCSI, Fibre Channel, and NAS connectivity for remote server storage
- Describe the Fibre Channel topologies, port types, and fabric login
- Configure and verify Fibre Channel communications on Cisco MDS switches

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 1

Comparing Storage Connectivity Options in the Data Center

Overview

This lesson compares Internet Small Computer System Interface (iSCSI), Fibre Channel, and network-attached storage (NAS) connectivity for remote server storage.

Objectives

Upon completing this lesson, you will be able to compare the three storage-connectivity options that are used in the data center. This ability includes being able to meet these objectives:

- Compare block- and file-based network storage
- Describe NFS, Fibre Channel, iSCSI, and SCSI
- Describe the evolution of SCSI from DAS to SAN

https://t.me/learningnets

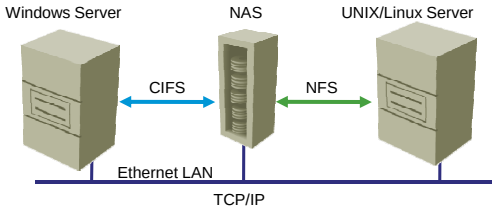
Comparing Block- and File-Based Network Storage

Compare the NAS and SAN approaches to network storage.

File-Based Protocols

CIFS and NFS are file-based protocols used for reading and writing files across a network.

- CIFS is used primarily on Windows servers.
  - Samba service implements CIFS on UNIX systems.
- NFS is used primarily on UNIX and Linux servers.
- Both have high latency over TCP/IP (very chatty protocols).
- Both provide shared access to files.
- Both are suitable for file-based apps:
  - Microsoft Office
  - Microsoft SharePoint
  - File and Print services



© 2012 Cisco and/or its affiliates. All rights reserved.

Common Internet File System (CIFS) and Network File System (NFS) are file-based protocols that are used for reading and writing files across a network. CIFS is found primarily on Microsoft Windows servers (a Samba service implements CIFS on UNIX systems), and NFS is found primarily on UNIX and Linux servers.

Because CIFS and NFS are transported over TCP/IP; latency is high. CIFS and NFS are verbose protocols that send thousands of messages between the servers and NAS devices when reading and writing files.

These file-based protocols are suitable for file-based applications:

- Microsoft Office
- Microsoft SharePoint
- File and Print services

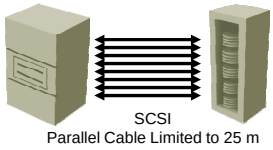
One benefit of NAS storage is that files can be shared easily between users within a workgroup.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Block-Based Protocols

- SCSI is a block-based protocol for writing and reading data blocks to and from a storage device.
- SCSI can be transported in several ways:
  - Parallel SCSI cable
    - Low latency, short distance
    - Half duplex, legacy deployment
  - iSCSI (SCSI over TCP/IP)
    - High latency, low bandwidth, inexpensive
  - Fibre Channel cable (traditional SAN)
    - Low latency, high bandwidth, lossless delivery
  - FCoE
    - Lossless delivery over unified I/O
- Suitable for block-based, structured applications
  - Database applications, high bandwidth, high performance
  - On-line Transaction Processing (OLTP) applications, high IOPS



© 2012 Cisco and/or its affiliates. All rights reserved.

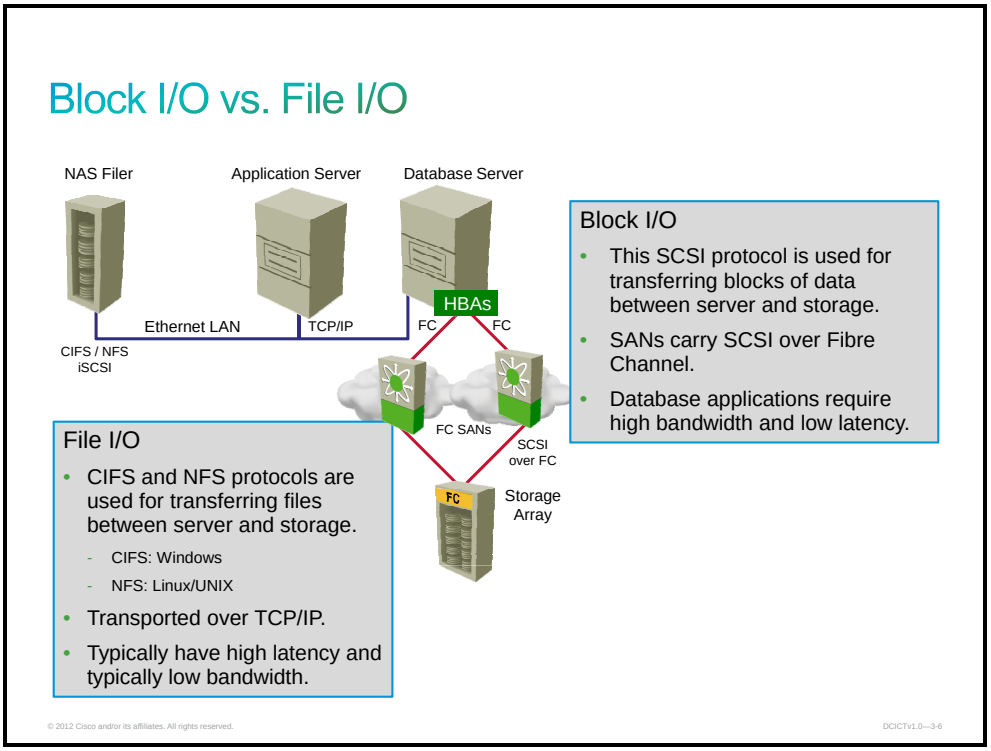
Small Computer Systems Interface (SCSI) is a block-based protocol for writing and reading data blocks to and from a storage device.

SCSI messages and data can be transported in several ways:

- **Parallel SCSI cable:** This transport is mainly used in traditional deployments. Latency is low, but distance is limited to 25 m and is half-duplex, so data can flow in only one direction at a time.
- **iSCSI is SCSI over TCP/IP:** This transport has much higher latency and relatively low bandwidth over Gigabit Ethernet LANs but is inexpensive because the software drivers are free and storage is cheaper.
- **Fibre Channel cable:** This transport is the basis for a traditional SAN deployment. Latency is low and bandwidth is high (8 Gb/s). SCSI is carried in the payload of a Fibre Channel frame between Fibre Channel ports. Fibre Channel has a lossless delivery mechanism by using buffer-to-buffer credits (BB_Credits).
- **Fibre Channel over Ethernet (FCoE):** This transport replaces the Fibre Channel cabling with 10 Gigabit Ethernet cables and provides lossless delivery over unified I/O.

The SCSI protocol is suitable for block-based, structured applications such as database applications that require many I/O operations per second (IOPS) to achieve high performance.

https://t.me/learningnets



Block I/O and file I/O are complementary solutions for reading and writing data to and from storage devices.

Block I/O uses the SCSI protocol to transfer data in 512-byte blocks. SCSI is an efficient protocol and is the accepted standard within the industry.

SCSI may be transported over many physical media. SCSI is mostly used in Fibre Channel SANs to transfer blocks of data between servers and storage arrays in a high-bandwidth, low-latency redundant SAN. Block I/O protocols are suitable for well-structured applications such as databases that transfer small blocks of data when updating fields, records, and tables.

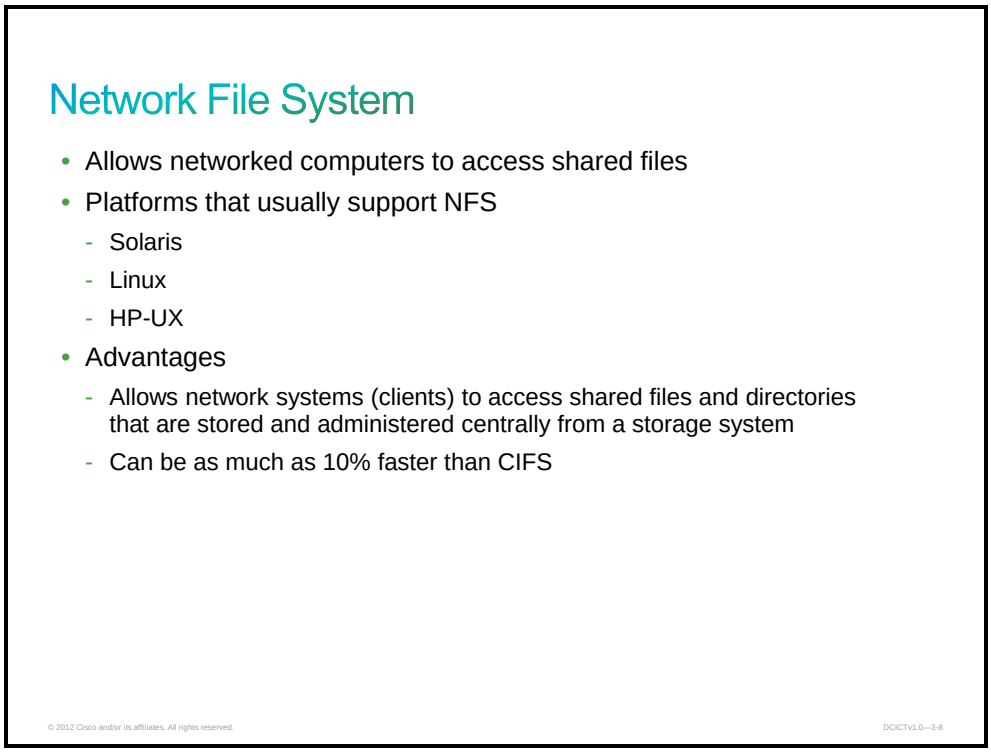
Another block I/O protocol is Serial Advanced Technology Attachment (SATA). This protocol is mostly used in home desktop and laptop computers.

File I/O uses the CIFS protocol primarily for transferring files to and from Windows servers and uses the NFS protocol for transferring files to and from UNIX and Linux servers. Both CIFS and NFS are carried over TCP/IP, usually across the LAN, and so have high latency and relatively low bandwidth. Most applications are file-based and read and write entire files. Examples are the Microsoft Office applications (i.e., Microsoft Word, Excel, and PowerPoint), audio and video applications, and web server applications that transfer HTML files to a browser.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

NFS, Fibre Channel, iSCSI, and SCSI

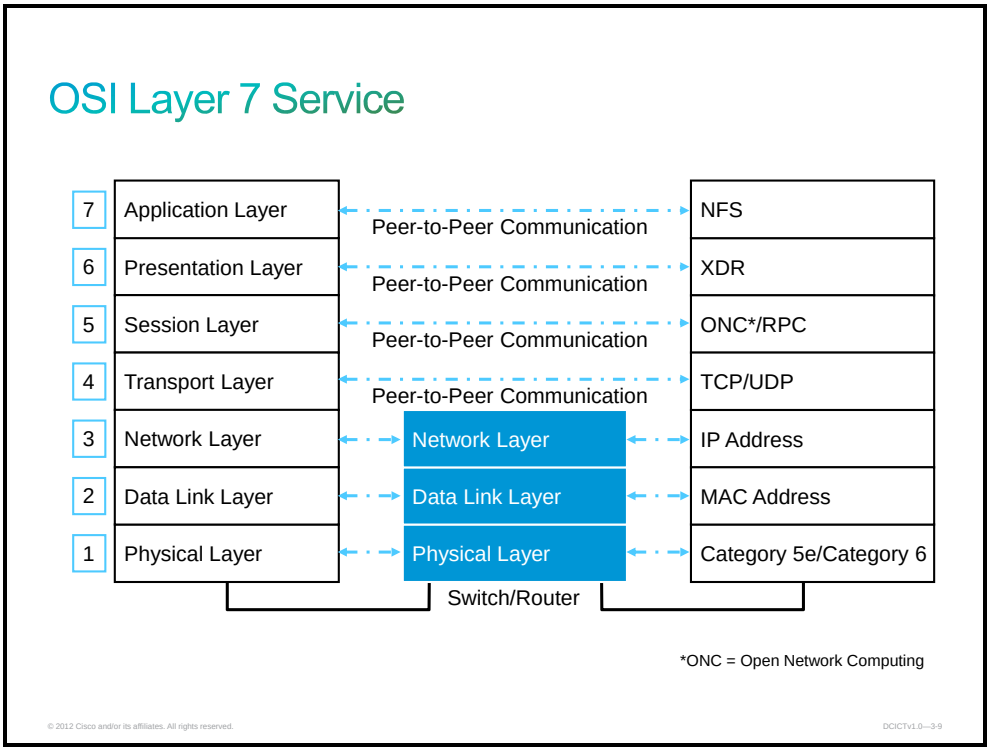
This topic discusses the NFS, Fibre Channel, iSCSI, and SCSI protocols for transporting data between servers and remote storage devices.



NFS is a distributed file system that Sun Microsystems developed in the 1980s to address the need to share resources in a distributed networking environment. Networked computers can share files across networks without being in the same physical location as the server. NFS removed the need to physically move files between computers.

An NFS server has one or more file systems that are mounted by NFS clients. To the NFS clients, the remote disks look like local disks. A NetApp storage system in a NAS implementation can act as the NFS server. There are three versions of NFS: version 2 (NFSv2), version 3 (NFSv3), and version 4 (NFSv4). Different clients running different UNIX or Linux operating systems can exchange files by using NFS.

https://t.me/learningnets

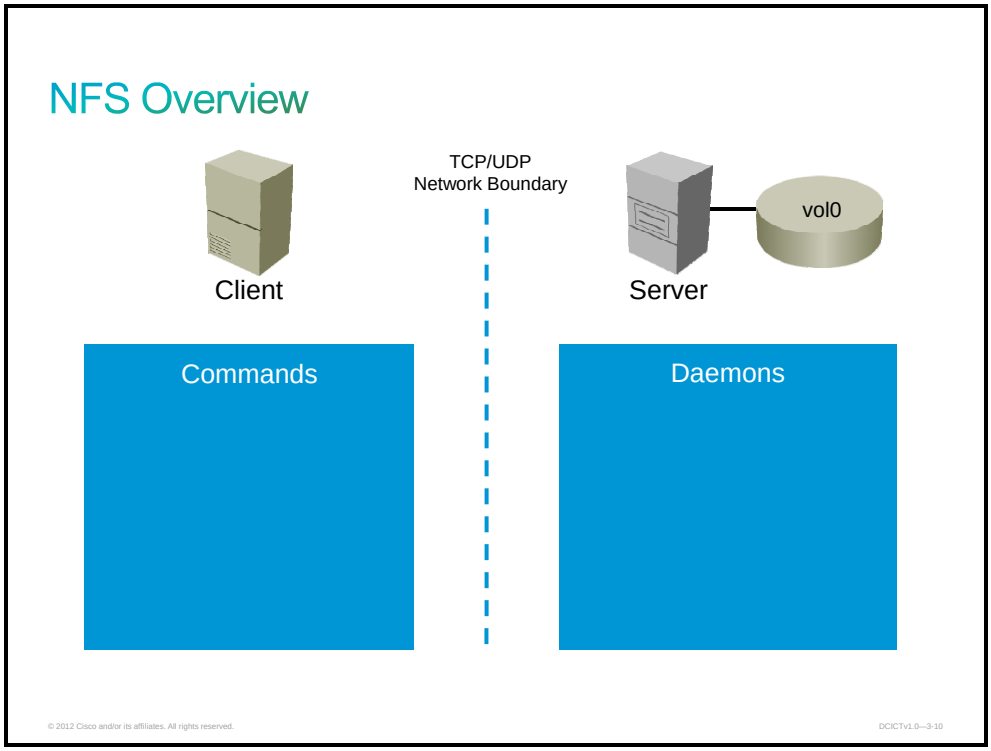


NFS is a service that resides at the application layer (Layer 7) of the Open Systems Interconnection (OSI) protocol stack. Similar to the way that network nodes communicate with one another based on the seven-layer architecture, NFS uses an interprocess communication method that is called a remote procedure call (RPC).

Layered Architecture

eXternal Data Representation (XDR) is a standard for the description and encoding of data. XDR is used to communicate data between different machines types. XDR fits into the presentation layer (Layer 6) of the OSI model. XDR is not a programming language but uses a language to describe data formats.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Client/Server Architecture

The theory of client/server architecture is based on the concept that one computer has the resources that another computer requires. These resources can be made available through NFS. The system with the resources is called the server, and the system that requires the resources is called the client. Examples of resources are email, database, and files. The client and the server communicate with each other through established protocols.

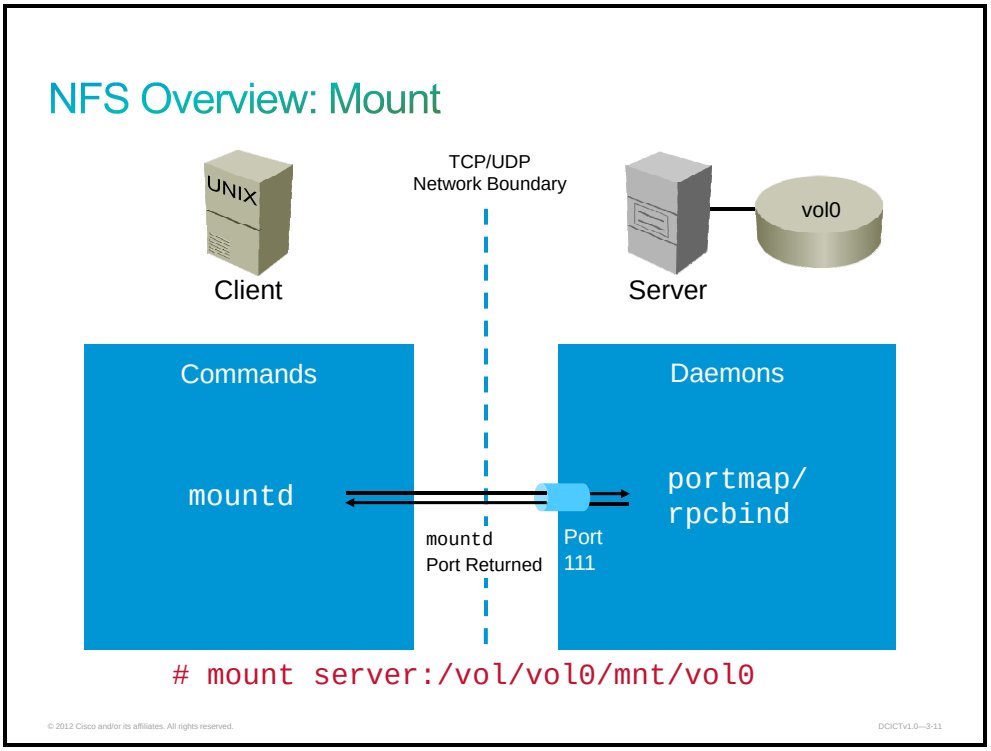
A distributed (client/server) network might contain multiple servers and multiple clients, or multiple clients and one server. The configuration of the network depends on the resource requirement of the environment.

The benefits of client/server architecture include cost reduction because of hardware and space requirements. The local workstations do not need as much disk space because commonly used data can be stored on the server. Other benefits include centralized support (backups and maintenance) that are performed on the server.

NFS is a widely used protocol for sharing files across networks. NFS is stateless, to allow for easy recovery in the event of server failure. In the figure, the server in the network is a network appliance storage system, and the client can be one of many versions of the UNIX or Linux operating system.

The storage system provides services to the client. These services include mount daemon (mountd), Network Lock Manager (nlm_main), NFS daemon (nfsd), Status Monitor (sm_1_main), quota daemon (rquot_1_main), and portmap (also known as rpcbind). Each service is required for successful operation of an NFS process. For example, a client cannot mount a resource if mountd is not running on the server. Similarly, if rpcbind is not running on the server, NFS communication cannot be established between the client and the server.

https://t.me/learningnets



In the figure, a storage system is configured as the NFS server. The NFS client first uses the standard UNIX mount command to mount the required file system.

Next, the client makes an RPC to the portmap or rpcbind daemon that is running on the server.

Remote Procedure Call

RPC is a client/server programming environment that vendors use for developing platform-independent applications. RPC allows applications (programs) to communicate with one another, just like network nodes communicate with one another by using TCP or UDP. RPC provides an environment in which the client process can have the server process execute a procedure call as if the process were a subroutine on the local system.

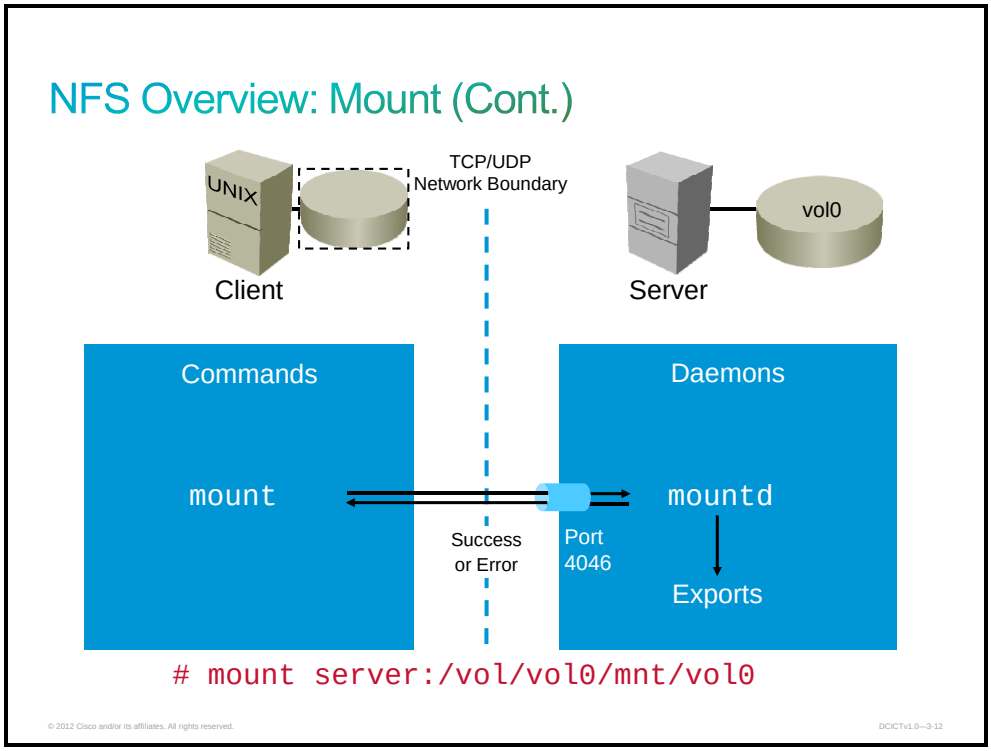
Portmap

A portmap, sometimes known as rpcbind, is an RPC service that allows clients and servers to communicate with one another by using interprocess communication methods. Just like network nodes use IP addresses to communicate with one another, the portmap service allows RPC services (processes or programs) to use assigned port addresses to communicate with one another.

A portmap allows these RPC services to use assigned ports, as long as they are registered with the portmap with program number, version, and transport protocol. The portmap program is usually registered on port 111 of the TCP and UDP transport protocols. Port 111 is also known as a privileged port in UNIX. A privileged port is any port number lower than 1024 and can be used as a source port only by a UNIX superuser. NFS servers usually use port 2049 as the default.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



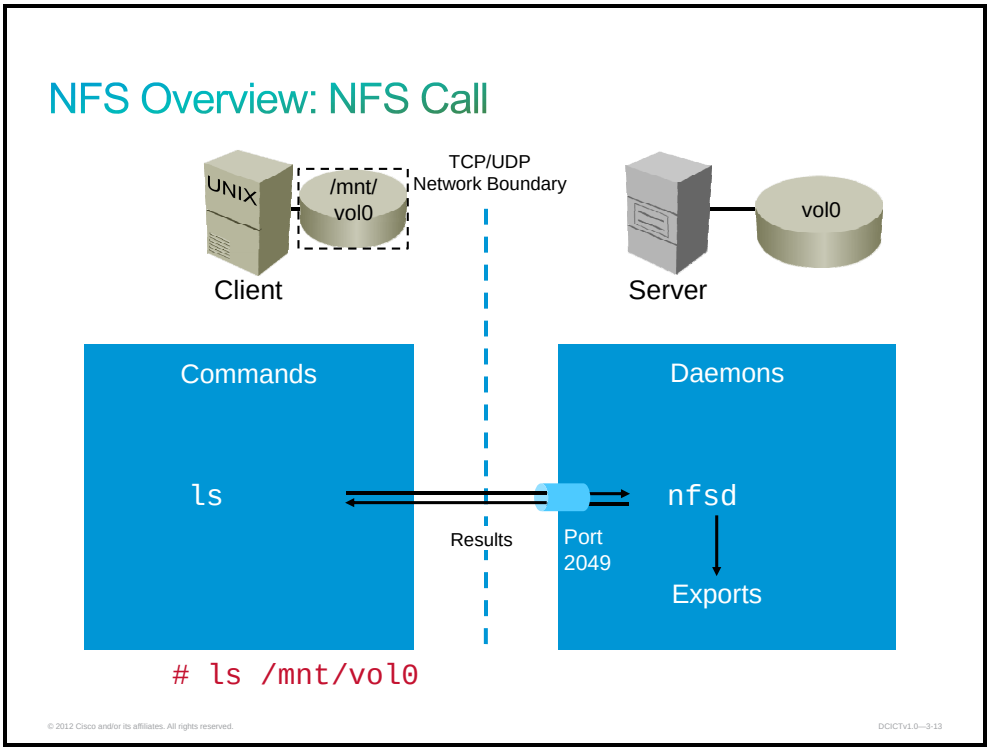
The client then issues a call to the mount daemon (mountd) on the server. Mountd verifies access to the resource and caches the results in the access. Either a successful result or an error is returned. If the mount command was successful, the resource is now accessed at the mountpoint, as shown in the figure.

Automounter

Automounter is an NFS program that mounts the file system on demand and unmounts the files if they are not accessed within a few minutes.

Note	An automounter is a mechanism for mounting resources that are operating system dependent.

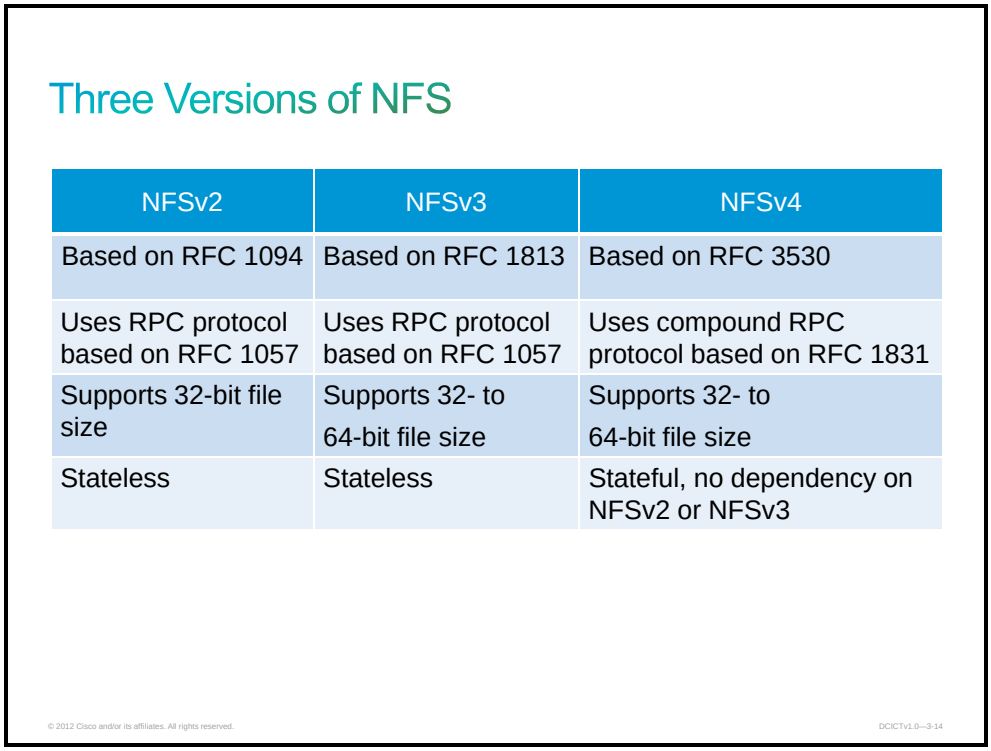
https://t.me/learningnets



The client can now issue file system commands (such as **ls**) within the mountpoint. An RPC call is sent to the NFS daemon (nfsd) on the server, to process the call and return the results.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



NFSv2

The features and functions of NFSv2 are defined in RFC 1094. NFSv2 uses RPC that is defined in RFC 1057 and uses XDR that is defined in RFC 1014. Maximum file size is 4 GB.

NFSv3

NFSv3 was developed to minimize the limitations in NFSv2, especially the file-size limitations (64-bit file size). NFSv3 is defined in RFC 1813 and uses the same RPC as NFSv2.

NFSv4

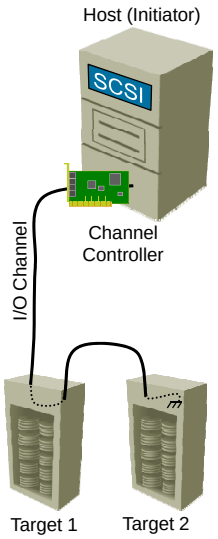
NFSv4 is a distributed file system that is based on RFC 3530. NFSv4 uses RPC that is defined in RFC 183 and supports the same file size as NFSv3. NFSv4 is designed to use the Internet, support traditional file-access methods, and integrate support for file locking and mount protocol. NFSv4 uses XDR that is based on RFC 1832. NFSv4 makes provision for end-to-end security, and Kerberos 5 is one of the supported methods.

Data ONTAP 6.4 and later provide complete NFSv4 server and client (Linux, Sun Microsystems Solaris, Hummingbird, and Windows) support.

https://t.me/learningnets

SCSI Overview

- Most storage networks use the SCSI protocol for communication between servers and disk drive devices:
  - The SCSI channel is used to transmit SCSI commands, data, and status.
  - The most common channel is the basic parallel SCSI bus, which can be internal or external to a host.
- Multiple devices are supported on the channel:
  - Enables communication between the SCSI initiator and SCSI target.
- SCSI is a half-duplex parallel standard with some limitations:
  - As many as 16 devices per SCSI bus
  - Bus length up to 82 feet (25 m)
  - Shared channel bandwidth of up to 320 MB/s



© 2012 Cisco and/or its affiliates. All rights reserved.

SCSI technology supports writing blocks of one or more files or databases to an external disk in a SAN. The files are written by using the host system attributes, including permissions, file size, and modification date. The SCSI channel that is used to transmit SCSI commands, data, and status is the foundation for all communications in a SAN (for SCSI devices). Connectivity between a SCSI host and SCSI device can be supported on either an internal or external connection.

SCSI uses a bus technology that supports multiple devices that are attached to a controller and can support daisy-chaining. Through the bus connection, a host can communicate with several SCSI devices by using multiple half-duplex I/O channels. SCSI communication is commonly described as being between an initiator and a target. An initiator is typically the host computer. A target is typically the storage device. In the figure, the initiator is a SCSI adapter, and the target is a Redundant Array of Independent Disks (RAID) array.

The parallel SCSI specification contains these elements:

- Bus length of as much as 82 feet (25 m)
- Shared channel bandwidth of as much as 320 MB/s
- As many as 16 devices per SCSI bus
- Half-duplex operation

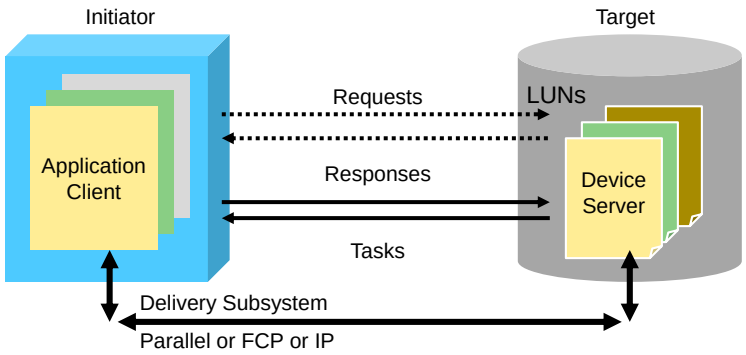
Daisy-chaining slightly increases SCSI scalability. Because daisy-chained storage devices are captive behind a server, this technology still has limitations from a networking perspective.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

SCSI Protocol

- SCSI performs the passing of commands, status, and block data between initiators and targets.
- SCSI is a hierarchy of functions to assemble raw data blocks into application-readable files.



© 2012 Cisco and/or its affiliates. All rights reserved.

SCSI performs the more challenging part of passing commands, status, and block data between platforms and storage devices. One function of operating systems is to hide the complexity of the computing environment from the end user. Management of system resources, including memory, peripheral devices, display, and context-switching between concurrent applications, is generally concealed behind the user interface. The internal operations of the operating system must be robust and must closely monitor changes of state. They must also ensure that transactions are completed within the allowable timeframes and must automatically initiate recovery or retries in the event of incomplete or failed procedures. The SCSI protocol provides these functions for I/O operations for peripheral devices such as disk, tape, optical storage, printers, and scanners. This protocol is typically embedded in a device driver or logic onboard a host adapter.

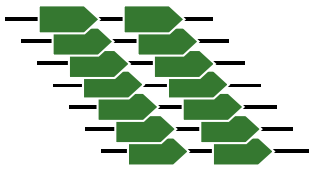
The SCSI protocol layer sits between the operating system and the peripheral resources, so it has different functional components. Applications typically access data as files or records. Although this information might be stored on disk or tape media in the form of data blocks, retrieval of the file requires a hierarchy of functions. These functions assemble raw data blocks into a coherent file that an application can manipulate.

The SCSI architecture defines the relationship between initiators (hosts) and targets (such as disks) as a client/server exchange. The SCSI version 3 (SCSI-3) application client resides in the host and represents the upper layer application, file system, and operating system I/O requests. The SCSI-3 device server sits in the target device, responding to requests.

https://t.me/learningnets

SCSI Protocol (Cont.)

- SCSI uses a parallel architecture in which data is sent simultaneously over multiple wires.
- SCSI is half-duplex—data travels in one direction at a time.
- On a SCSI bus, a device must assume exclusive control over the bus to communicate. (SCSI is sometimes referred to as a simplex channel because only one device can transmit at a time.)



- Half-duplex
- Parallel
- Shared bus

© 2012 Cisco and/or its affiliates. All rights reserved.

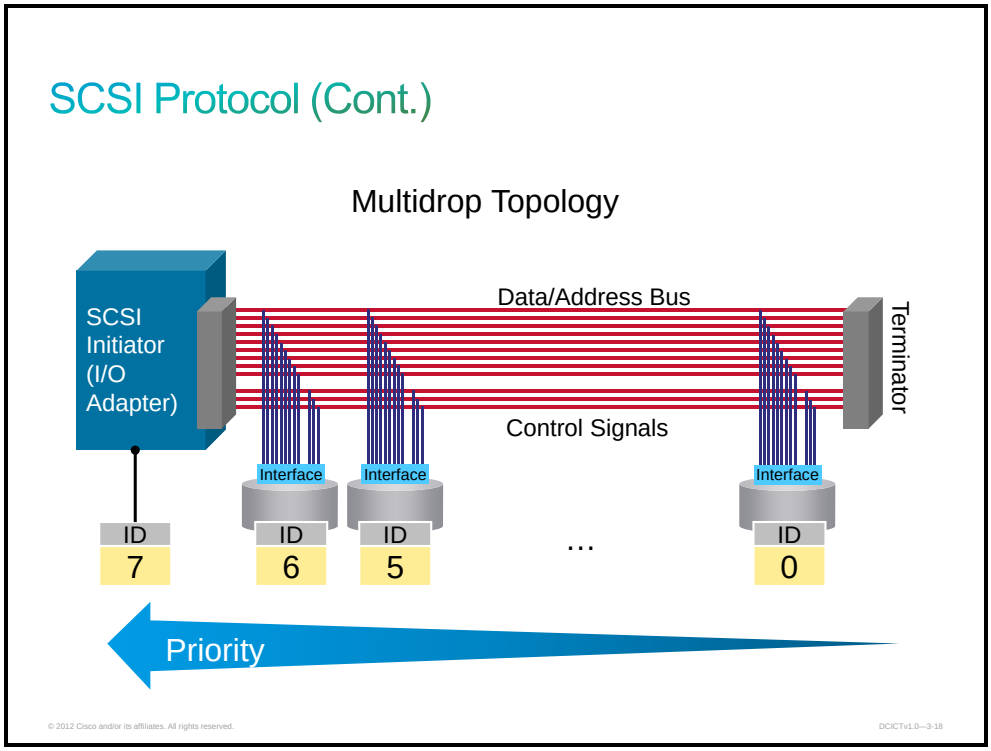
The bus, target, and logical unit number (LUN) triad is defined from parallel SCSI technology. The bus represents one of several potential SCSI interfaces that are installed in the host, each supporting a separate string of disks. The target represents a single disk controller on the string. The LUN designation allows additional disks that are governed by a controller; for example, a RAID device.

These are characteristics of parallel SCSI technology:

- SCSI uses a parallel architecture in which data is sent simultaneously over multiple wires.
- SCSI is half-duplex (data travels in one direction at a time).
- On a SCSI bus, a device must assume exclusive control over the bus to communicate.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

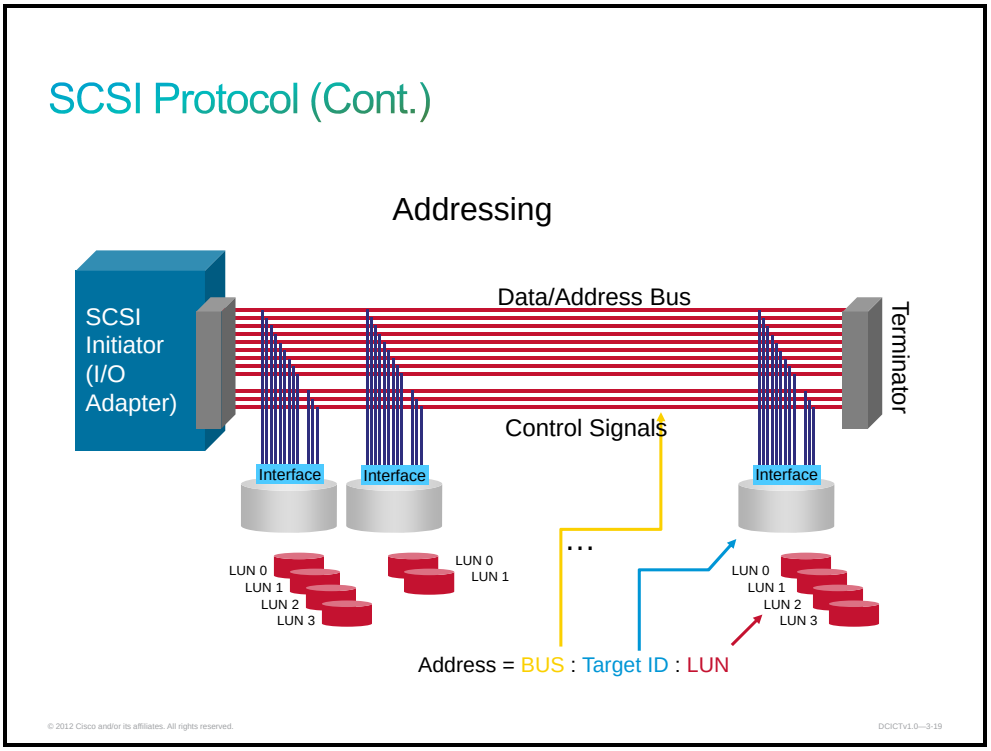


Multidrop Topology and Addressing

All the devices on a SCSI bus connect to one cable. This design is called a multidrop topology and has these characteristics:

- Data bits are sent in parallel on separate wires. Control signals are sent on a separate set of wires.
- Only one device at a time can transmit, and the transmitting device has exclusive use of the bus.
- A special circuit that is called a terminator must be installed at the end of the cable. The cable must be terminated to prevent unwanted electrical effects from corrupting the signal.
- A multidrop topology has inherent limitations:
  - Parallel transmission of data bits allows more data to be sent in a given period but complicates transmitter-receiver synchronization. This complication is caused by control signals, such as clock signals, being sent on a separate set of wires.
  - A multidrop topology is an inefficient way to use the available bandwidth, because only one communication session can exist at any time.
  - Termination circuits are built into most SCSI devices, but the administrator often must set a jumper on the device to enable termination.
  - Incorrect cable termination can cause a severe failure or intermittent, difficult-to-trace errors.

https://t.me/learningnets



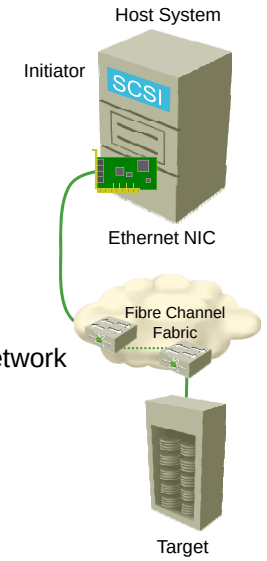
SCSI was designed to support a few devices at most. Therefore, the device addressing scheme is fairly simple and inflexible. SCSI devices use hard addressing, which has these characteristics:

- Each device has a series of jumpers that determine the physical address of the device, or SCSI ID. The ID is software-configurable on some devices.
- Each device must have a unique ID. Before adding a device to the cable, the administrator must know the ID of every other device that connects to the cable. The administrator must select a unique ID for this new device.
- The ID of each device determines its priority on the bus. For example, the SCSI target with ID 7 always has a higher priority than the SCSI initiator with ID 6. Each device must have exclusive use of the bus while transmitting, so ID 6 must wait until ID 7 has finished. Fixed priority makes it more difficult for administrators to control performance and quality of service (QoS).

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Fibre Channel Overview

- Extends and networks SCSI:
  - Provides high-speed transport for the SCSI payload
  - Uses a much more scalable serial standard
- Fibre Channel benefits:
  - Addressing for as many as 16 million devices
  - Loop (shared) and fabric (switched) transport
  - Host speeds of 100–800 MB/s (1–8 Gb/s)
  - Segments of up to 10 km (without extenders)
  - Multiple protocol support
- Combines the best attributes of a channel and a network



© 2012 Cisco and/or its affiliates. All rights reserved.

The serial connectivity of Fibre Channel provides a mechanism for transporting SCSI information across high-speed networks. Fibre Channel provides high-speed transport for the SCSI payload but overcomes the distance and limitations that come with parallel SCSI technology.

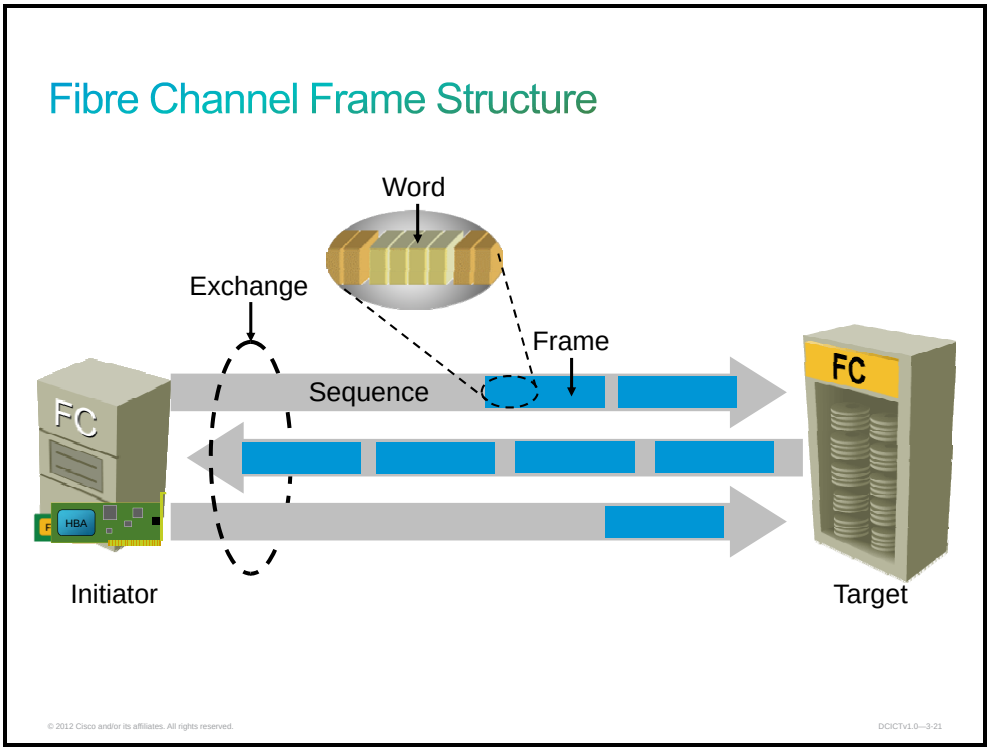
Fibre Channel includes these features:

- Addressing for as many as 16 million nodes
- Loop (shared) and fabric (switched) transport options
- Host speeds of 100 to 800 MB/s, or an effective throughput of 1 to 8 Gb/s on the fabric
- Segment distances of as far as 6.2 miles (10 km)
- Support for multiple protocols

Fibre Channel is the integration of the best attributes of host channel and networking technologies. Fibre Channel implements attributes from channel technology in the mainframe environment, including reliability and scalability. Fibre Channel also implements attributes from networking, including connectionless services, high connectivity, and long distances.

Fibre Channel supports three topologies: point-to-point, arbitrated loop (similar to Token Ring), and switched fabric. Fibre Channel switching is point-to-point oriented. Both the initiator and storage target side are points. Fibre Channel has become the dominant open-systems protocol for connecting servers to storage.

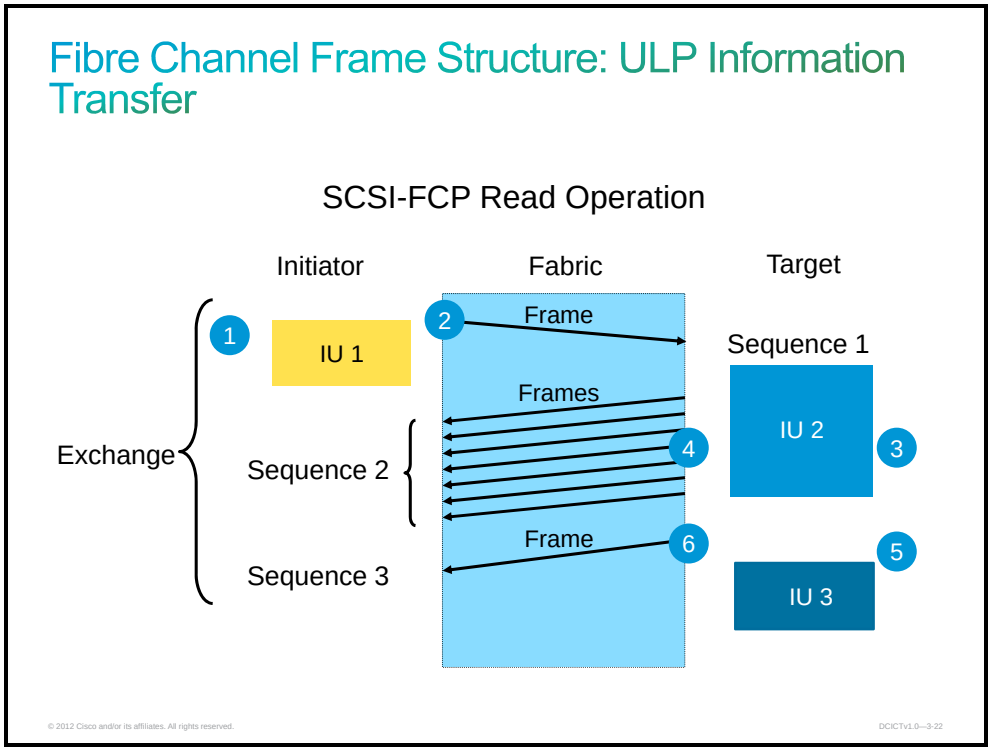
https://t.me/learningnets



The figure shows a transaction between a host (initiator) and a storage device (target). The transaction has these characteristics:

- The smallest unit of data is a word. Words consist of 32 bits (4 bytes) of data that are encoded into a 40-bit form by the 8-bit/10-bit encoding process.
- Words are packaged into frames. A Fibre Channel frame is equivalent to an IP packet.
- A sequence is a series of frames that are sent from one node to another node. Sequences are unidirectional. In other words, a sequence is a set of frames that are issued by one node.
- An exchange is a series of sequences that are sent between two nodes. The exchange is the mechanism that two ports use to identify and manage a discrete transaction. The exchange defines an entire transaction, such as a SCSI read or write request. An exchange is opened whenever a transaction is started between two ports and is closed when the transaction ends. A Fibre Channel exchange is equivalent to a TCP session.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

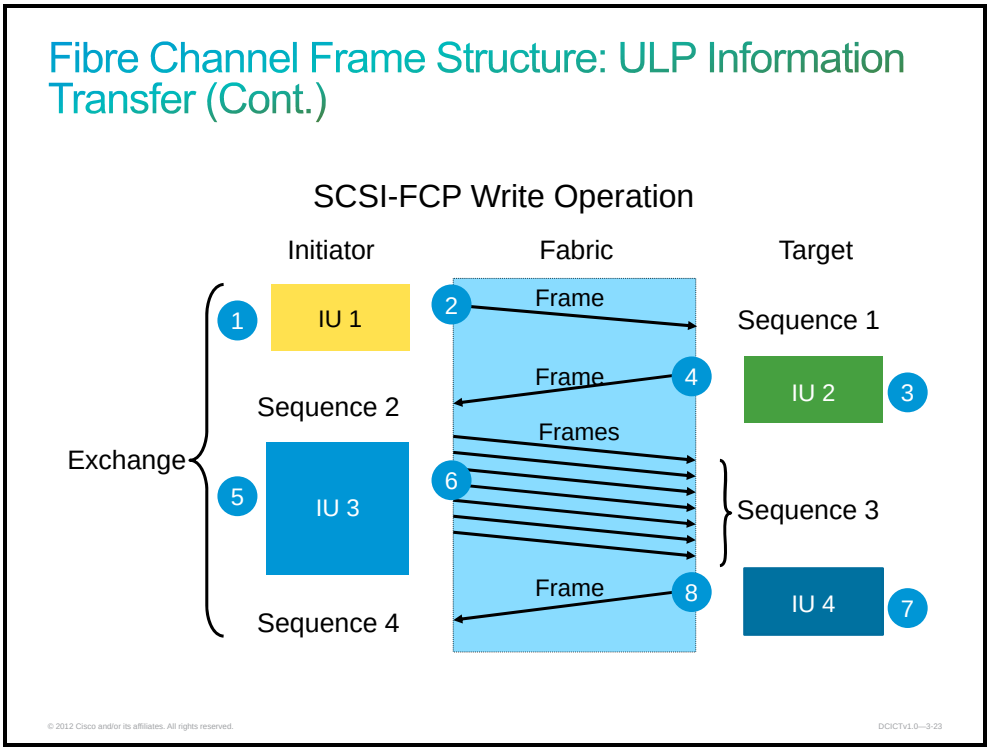


The figure shows a SCSI-Fibre Channel Protocol (FCP) read operation, which consists of these actions:

1. The initiator node generates a SCSI read request (FCP_CMD), which is packaged as information unit (IU) 1.
2. The initiator FC-2 layer converts IU 1 to a single command chunk and sends it across the fabric as a single frame. This process constitutes sequence 1.
3. The target node processes IU 1, retrieves the requested data (FCP_DATA) from storage, and packages the data as IU 2.
4. The target FC-2 layer converts IU 2 to one or more data chunks and sends them across the fabric. This process constitutes sequence 2.
5. The target node generates a status command (FCP_RSP) that informs the initiator that the requested data transmission is complete. The status command is packaged as IU 3.
6. The target FC-2 layer converts IU 3 to a single command chunk and sends it across the fabric. This process constitutes sequence 3.

At this point, the I/O operation is complete. The collection of three sequences constitutes a single exchange.

https://t.me/learningnets

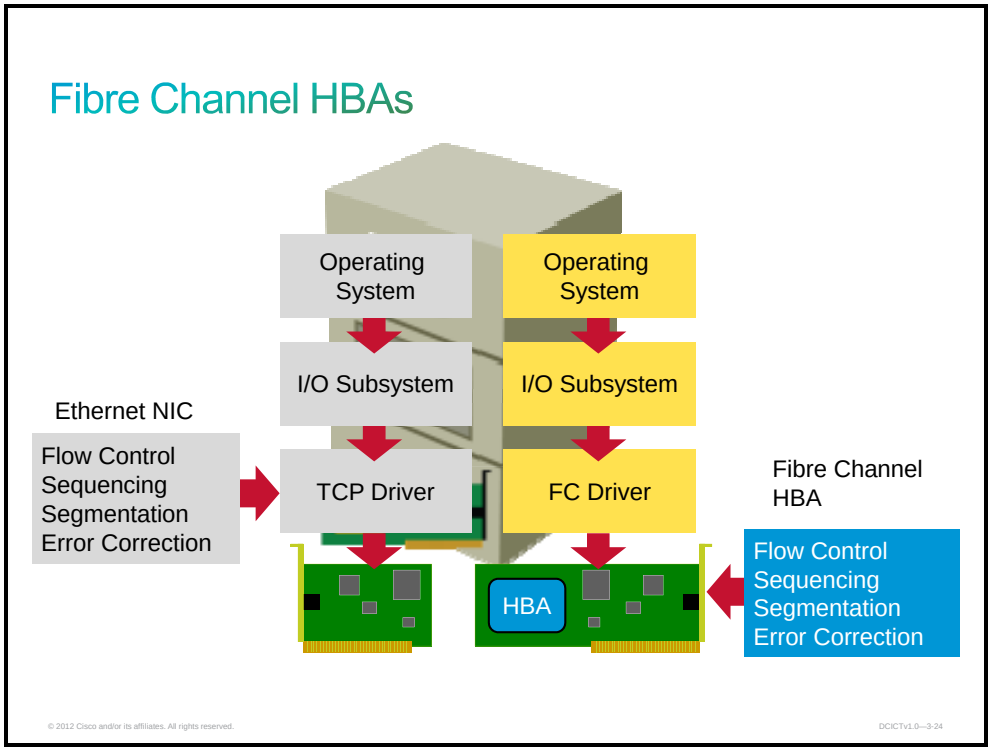


The figure shows a SCSI-FCP write operation, which consists of these actions:

1. The initiator node generates a SCSI write request (FCP_CMD), which is packaged by the FC-4 layer as IU 1.
2. The initiator FC-2 layer converts IU 1 to a single command chunk and sends it across the fabric as a single frame. This process constitutes sequence 1.
3. The target node responds with a SCSI write request response (FCP_XFR_RDY), which is packaged as IU 2. The write request response is required for synchronization between the initiator and target.
4. The target FC-2 layer converts IU 2 to a single command chunk and sends it across the fabric. This process constitutes sequence 2.
5. The initiator node retrieves the data (FCP_DATA) from its upper-layer protocol (ULP) buffers and packages it as IU 3.
6. The initiator FC-2 layer converts IU 3 to one or more data chunks and sends them across the fabric. This process constitutes sequence 3.
7. The target generates a status command (FCP_RSP) to confirm the end of the exchange. The command is packaged as IU 4.
8. The target FC-2 layer converts IU 4 to a single command chunk and sends it across the fabric. This process constitutes sequence 4.

The collection of four sequences constitutes a single exchange.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Host bus adapters (HBAs) are I/O adapters that are designed to maximize performance by performing protocol-processing functions in silicon. HBAs are roughly analogous to network interface cards (NICs), but HBAs are optimized for SANs and provide features that are specific to storage.

The figure contrasts HBAs with NICs, illustrating that HBAs offload protocol-processing functions into silicon.

With NICs, software drivers perform protocol-processing functions such as flow control, sequencing, segmentation and reassembly, and error correction. The HBA offloads these protocol-processing functions onto the HBA hardware—usually some combination of an ASIC and firmware. Offloading these functions is necessary to provide the performance that storage networks require.

NICs can utilize more than 80 percent of a server CPU capacity (measured with a 1-GHz Intel Pentium CPU) to deliver 50 to 80 MB/s on a Gigabit Ethernet link. I/O processing adds considerable real cost to what might appear to be an inexpensive NIC.

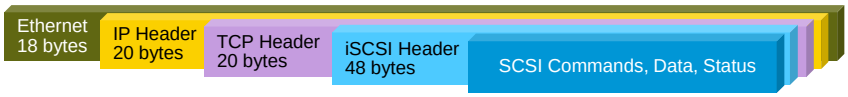
HBAs manage I/O transactions with little or no involvement of the server CPU. Fibre Channel HBAs can provide throughput at nearly 95 percent of the link speed with less than 10-percent server CPU utilization.

https://t.me/learningnets

iSCSI Protocol Overview

iSCSI encapsulates SCSI commands, data, and status over an IP transport for linking storage devices with servers:

- Allows IP hosts to gain access to Fibre Channel-based storage targets
- Leverages existing IP-based infrastructures and management facilities
- TCP (port 3260) provides congestion control and in-order delivery of error-free data
- Addresses distance limitations
- Cisco MDS 9222i or MDS 9000 18/4-Port multiservice modules provide transparent SCSI routing



© 2012 Cisco and/or its affiliates. All rights reserved.

iSCSI is a transport protocol that operates in addition to TCP and encapsulates SCI-level commands and data into IP for a TCP/IP byte stream. iSCSI is a means of transporting SCI packets over TCP/IP, providing for an interoperable solution that can take advantage of existing IP-based infrastructures and management facilities. The SCSI protocol has been mapped over various transports, including Parallel SCSI, Intelligent Peripheral Interface (IPI), IEEE 1394, FireWire, and Fibre Channel. These transports are I/O-specific and have limited distance capabilities.

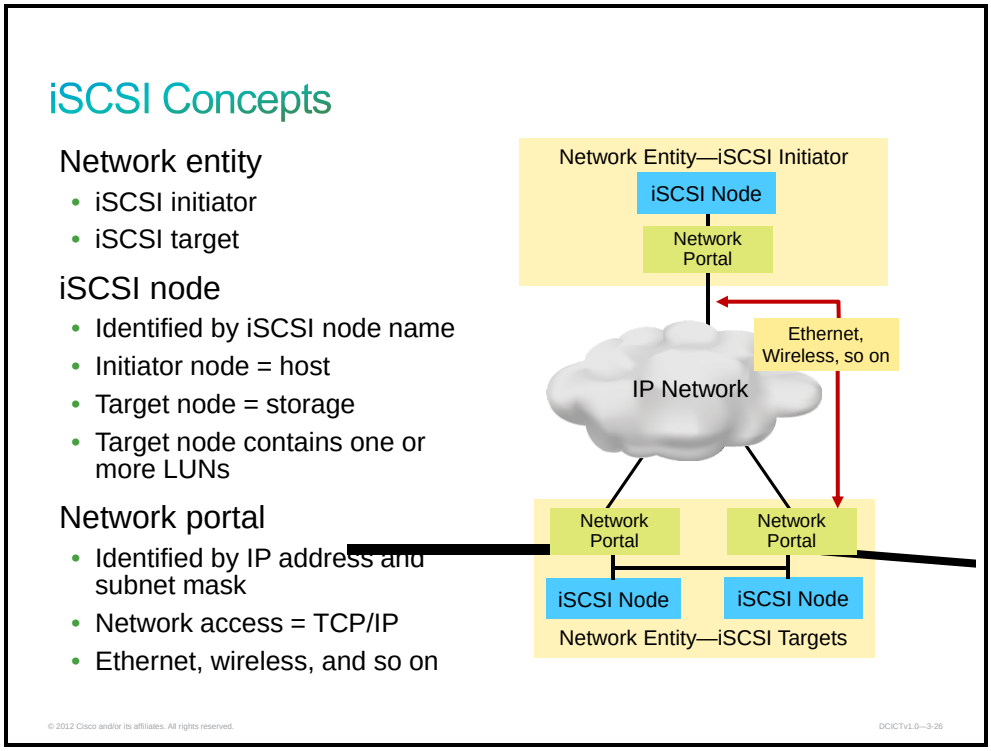
Mapping SCSI (I/O) over TCP ensures that high-volume storage transfers have in-order delivery and error-free data with congestion control. This process allows IP hosts to gain access to previously isolated Fibre Channel-based storage targets and overcomes distance limitations.

iSCSI is an end-to-end protocol with human-readable SCSI device, node, and naming. It includes these base components:

- IP Security (IPsec) connectivity security
- Authentication for access configuration
- Discovery of iSCSI nodes
- Remote boot process
- iSCSI MIB standards

The iSCSI protocol was defined by the Internet Engineering Task Force (IETF) IP Storage Working Group.

https://t.me/learningnets



iSCSI Concepts

SCSI standards define a client/server relationship between the SCSI initiator and the SCSI target. iSCSI standards define this concept as the network entity.

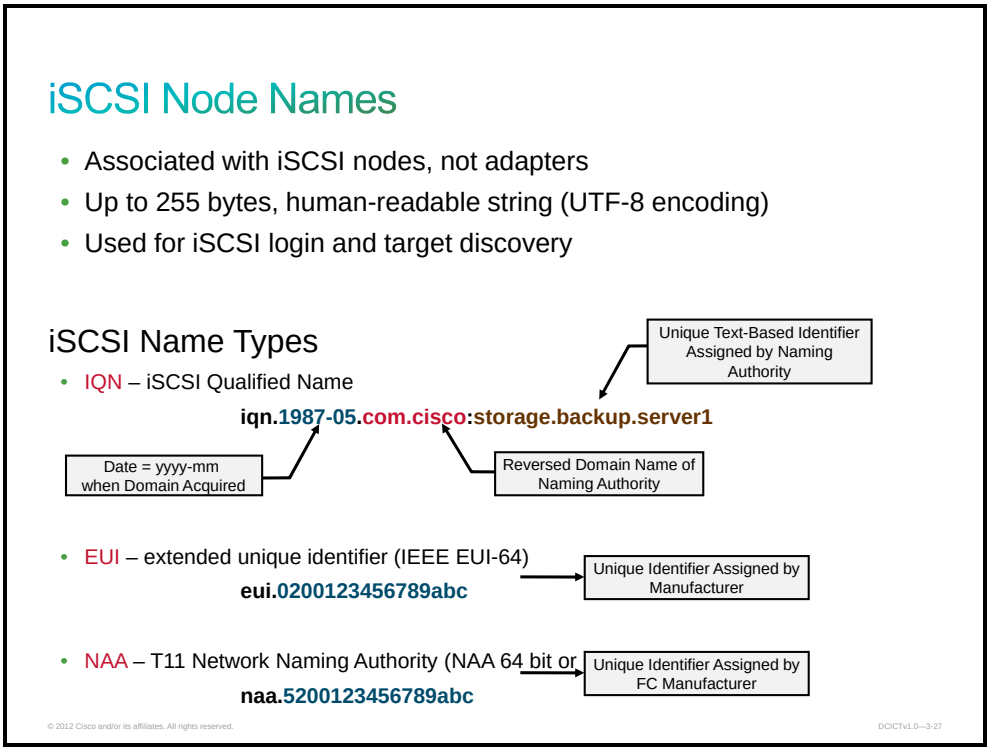
The iSCSI network entity contains an iSCSI node, which is either the initiator or target. iSCSI nodes are identified by an iSCSI node name.

If the target node is a storage array, it may contain one or more SCSI LUNs.

iSCSI initiator nodes communicate with iSCSI target nodes through network portals. Network portals connect to the IP network and are identified by an IP address.

Note that network portals can also be wireless Ethernet ports.

https://t.me/learningnets



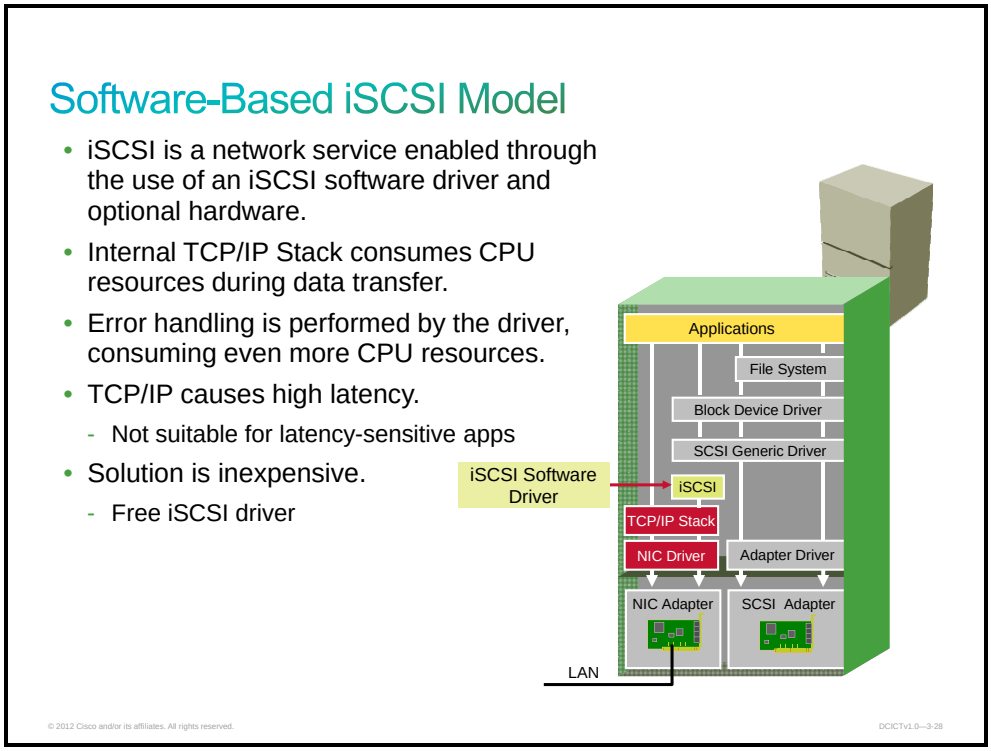
iSCSI Node Names

Every iSCSI node is identified by an iSCSI node name in one of three formats:

- **IQN:** An iSCSI Qualified Name (IQN) can be as much as 255 bytes and is a human-readable UTF-8 encoded string.
- **EUI:** An extended universal identifier (EUI) is an 8-byte hexadecimal number that is defined and allocated by IEEE.
- **NAA:** A Network Address Authority is an 8- or 16-byte hexadecimal number that is defined and allocated by the Fibre Channel T11 committee

The iSCSI driver typically uses the IQN format. Manufacturers of native iSCSI devices use the EUI format.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



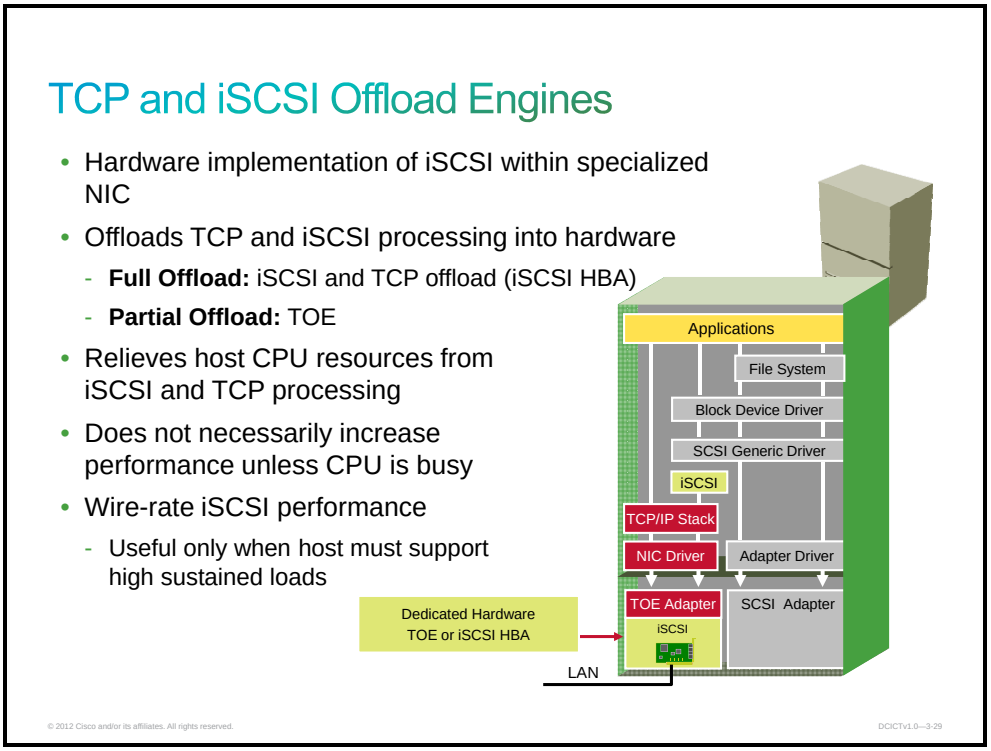
Software-Based iSCSI Drivers

iSCSI drivers are typically free and provide a low-cost solution to customers that do not require high performance or low latency.

The iSCSI driver performs all SCSI processing, TCP/IP processing, and error recovery.

An iSCSI driver running on a 1-GHz CPU spends nearly 95 percent of its CPU cycles when moving data through a NIC at 1 Gb/s. Current host CPUs are much faster and thus use far fewer of their CPU cycles moving data at 1 Gb/s, which makes this low-cost solution practical in some environments.

https://t.me/learningnets



TCP and iSCSI Offload Engines

Some NICs have a TCP Offload Engine (TOE) to offload TCP/IP processing from the host CPU and reduce the CPU load.

Some cards have partial offload and some have full offload.

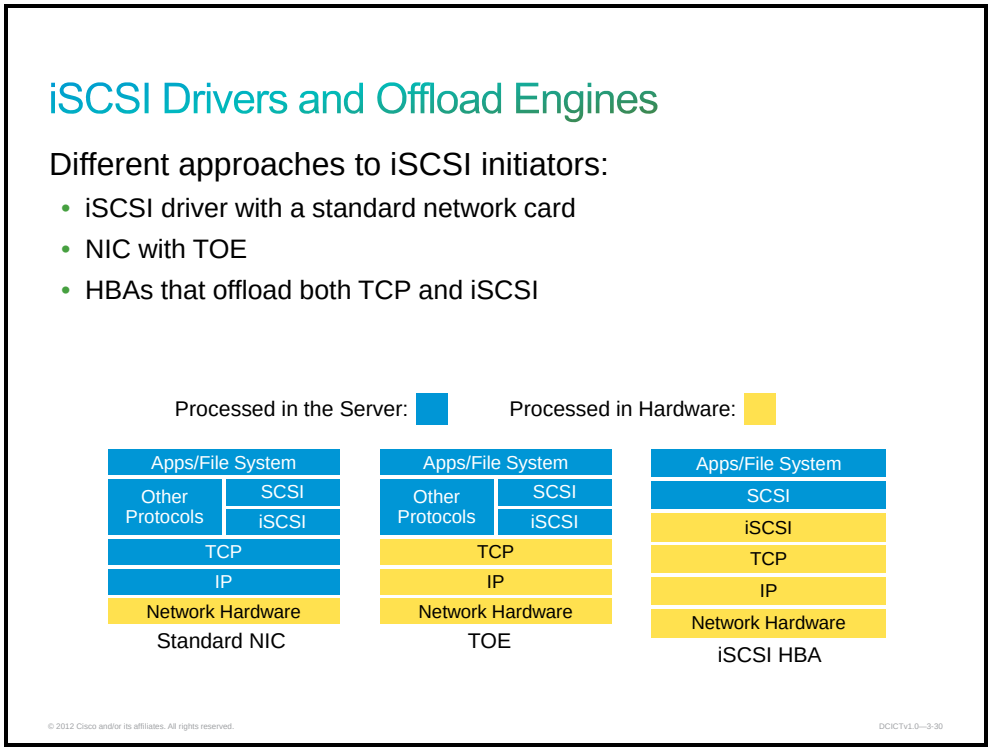
Partial-offload TOE cards offload TCP/IP processing to the TOE but pass all errors (packet loss) to the driver running on the host CPU. In a lossy network, partial-offload TOE cards may perform worse than before.

Full-offload TOE cards offload both TCP/IP processing and error recovery to the TOE card. The host CPU is still responsible for SCSI and iSCSI processing.

The iSCSI HBA offloads TCP/IP processing and iSCSI processing to coprocessors and custom ASICs on the iSCSI HBA. Although relatively expensive, iSCSI HBAs provide lower latency and higher throughput than iSCSI software drivers or TOE cards.

Current host CPU processors have faster performance. Using NICs with TOE is rarely are cost effective as using software iSCSI drivers.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



iSCSI Drivers and Offload Engines

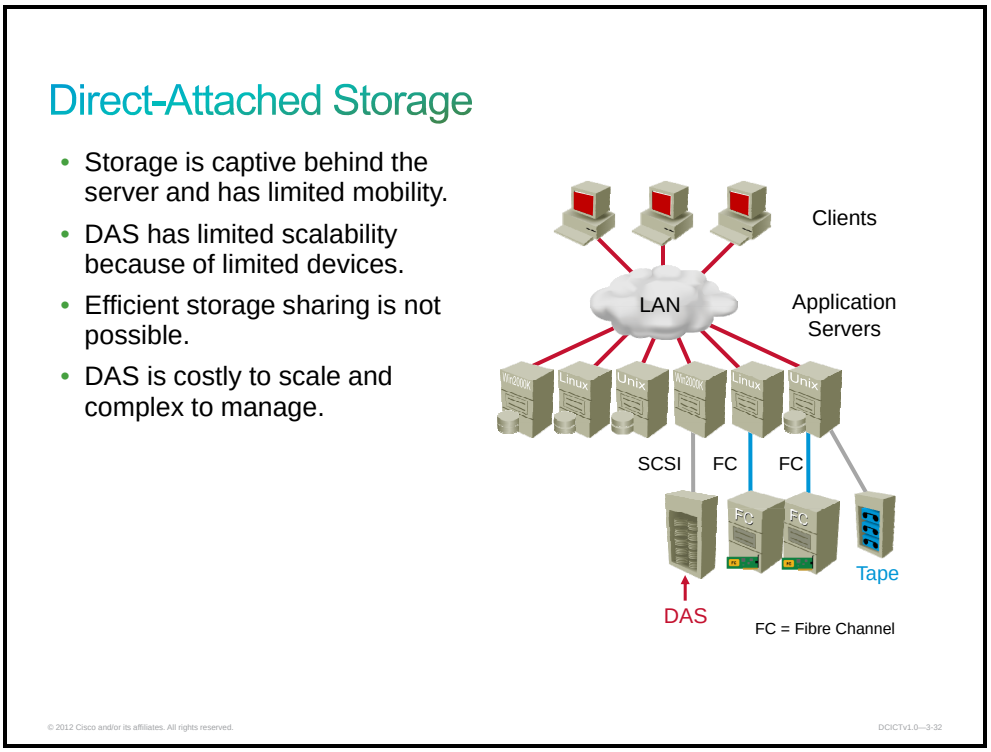
iSCSI drivers running on the host perform all SCSI and iSCSI processing, using the host CPU through the standard NIC. As I/O loads increase, the host consumes more CPU cycles and struggles to deliver throughput. In a congested IP network in which packets are frequently discarded, the TCP stack running on the host CPU must also recover lost packets.

To achieve maximum performance, it is necessary to offload TCP/IP processing, iSCSI processing, and error recovery from the host CPU onto the iSCSI HBA. The host is still responsible for SCSI processing.

https://t.me/learningnets

SCSI Evolution from DAS to SAN

Remote storage has been through several different topologies, from Direct Attached Storage (DAS) to SAN. SAN comprises one or more Fibre Channel switches, attached hosts, and storage devices. This topic describes how SCSI can be more scalable when using SAN or NAS.



DAS is commonly described as captive storage. DAS is often implemented within a parallel SCSI implementation. Devices in a captive storage topology do not have direct access to the storage network and do not support efficient sharing of storage. For example, the only way to access the DAS devices in the figure is by using the host CPU.

To access data with DAS, a user must go through some sort of front-end network. DAS devices provide little or no mobility to other servers and little scalability. DAS devices limit file sharing and can be complex to implement and manage.

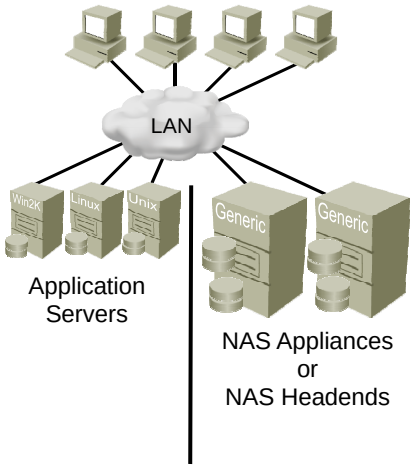
For example, to support data backups, DAS devices require resources on the host and spare disk systems that cannot be used on other systems.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Network-Attached Storage

- Storage is accessed over an IP network.
  - Storage devices can be shared between servers.
  - Files can be shared between users.
- Storage is accessed at a file level via NFS or CIFS.
  - TCP can be tuned to optimize file transport.
- NAS supports slower file retrieval than DAS.



© 2012 Cisco and/or its affiliates. All rights reserved.

An NAS device is dedicated to file sharing. NAS devices have direct IP capabilities that allow access at a file level, using a protocol such as NFS or CIFS across an IP network. NAS devices provide data storage and access as well as the management of these functions. NAS devices support file-serving requests so that resources can be freed up on other servers. NAS servers commonly employ SCSI and RAID technologies internally. An NAS device is directly connected to a front-end network. Storage devices can be shared between servers and between users.

NAS devices respond to requests by providing portions of the file system. To retrieve a file, an NAS device must open a directory, read it, locate the file, check permissions, and then transfer the file. If the file is several directories deep, the NAS device needs to perform repeated operations to retrieve one file. This process can require considerable network traffic, although TCP can be tuned to help optimize the file transport.

Note	The block-level access operations on DAS devices are much faster than operations that require a search of the file or directory system on a volume.

https://t.me/learningnets

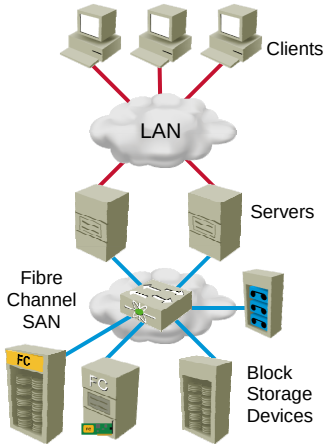
Storage Area Network

Benefits:

- Separates storage from the server
- High-performance interconnection can provide high I/O throughput
  - Storage accessed at a block level via SCSI
- Storage can be shared among servers
- Has lower TCO than DAS
- SAN fabric is interconnection hardware

Disadvantage:

- Might have limited vendor interoperability



© 2012 Cisco and/or its affiliates. All rights reserved.

A SAN is a high-speed subnetwork of storage devices. SAN technology is designed to remotely attach storage devices such as disk drives to host computers. SAN technology allows high throughput and therefore high performance. For example, an advantage of a SAN is that data access is provided at a block level via the SCSI protocol.

In a SAN, the clients request files from the servers, which have knowledge of the file system structure. The servers request specific blocks of data from a storage element in the back-end storage network.

SAN technology supports shared storage that is separate from the server or servers that use it. Storage resources are easily moved or allocated with no interruption of service on the servers. Managing a SAN tends to be more complex than managing a DAS. However, SANs provide a lower total cost of ownership (TCO) because the storage is not captive to one server, so there is a more efficient utilization of storage resources. The consolidation of multiple lower-end storage systems into centralized, higher-end systems allows reduced administrative overhead. Backup is easier with SAN technology, because it does not tie up host capacity or require a dedicated network or bandwidth on the main network. SAN technology also supports the fast implementation of high availability.

The hardware that connects servers to storage devices in a SAN is referred to as a fabric. The SAN fabric enables connectivity of any server to any storage device by using Fibre Channel switching technology. Fibre Channel is a standard from the Technical Committee T11 of the International Committee for Information Technology Standards (INCITS), an ANSI-accredited standards committee. The SAN fabric provides a high-speed dedicated network, with high-availability features, high throughput, and extremely low latency. When storage resources are consolidated into the same fabric, organizations can more easily make use of many features, including data replication, data sharing, and centralized backups.

https://t.me/learningnets

These are some business motivations for SAN deployments:

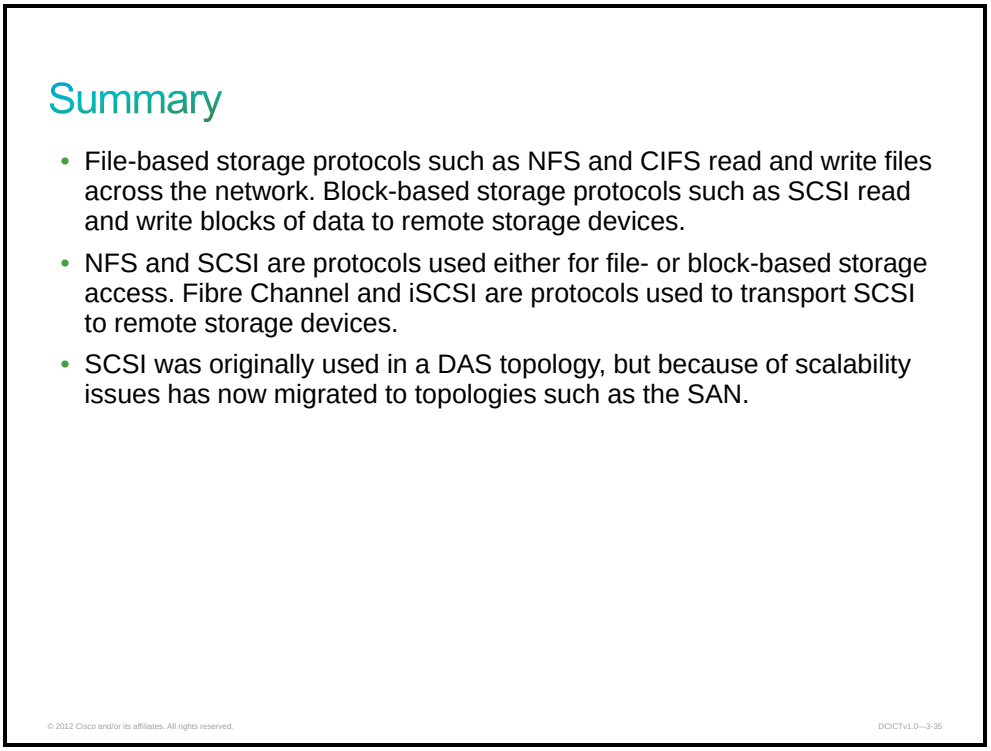
- Escalating storage requirements
- Rising storage management costs
- Business continuity
- Requirement to share information
- Increasing levels of underutilized disk-storage resources, resulting in increased backup storage and backup handling costs

One disadvantage to SANs is that there is limited vendor interoperability.

https://t.me/learningnets

Summary

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



© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Describing Fibre Channel Storage Networking

Overview

This lesson explains the three Fibre Channel topologies that can be found in storage networking. It also explains the use of the different port types and how a device performs a fabric login and port login to establish initiator-to-target connection and communication.

Objectives

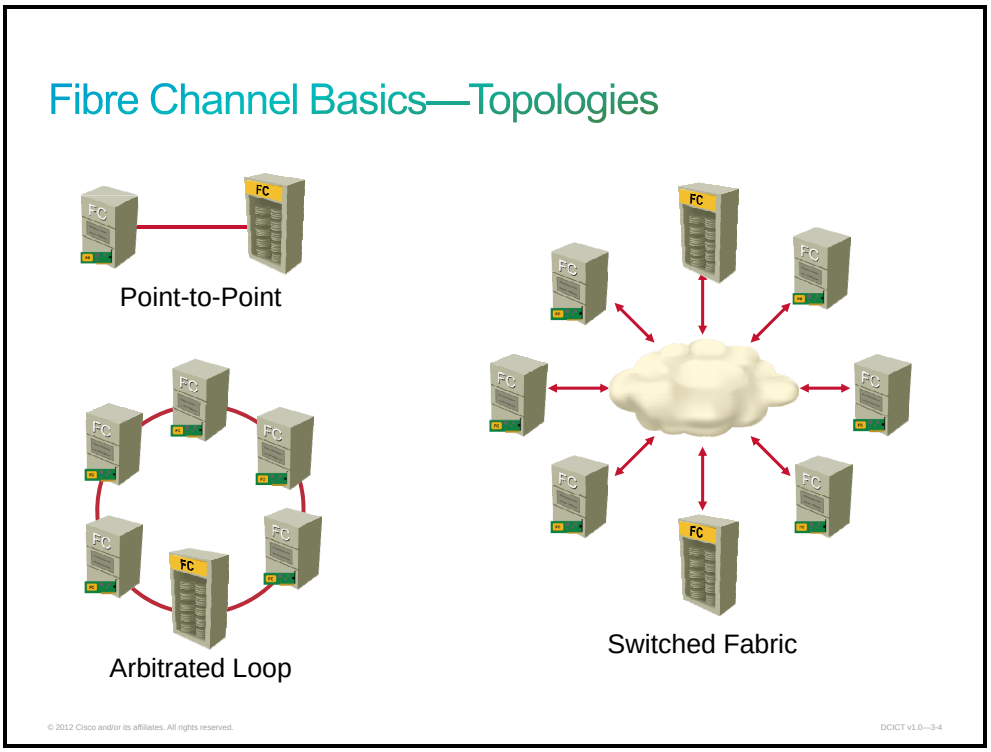
Upon completing this lesson, you will be able to explain the Fibre Channel topologies and how communication is established between an initiator and target device on the SAN. This ability includes being able to meet these objectives:

- Describe Fibre Channel SAN topologies
- Describe Fibre Channel port types
- Describe Fibre Channel addressing
- Describe the Fibre Channel layered model
- Describe the FCNS and FLOGI process
- Describe the use of Fibre Channel zoning and LUN masking

https://t.me/learningnets

Fibre Channel SAN Topologies

This topic explains the difference between the three Fibre Channel SAN topologies that may be used for initiator-to-target connectivity.



In a storage environment, communications takes place between an initiator and a target device. Three main topologies can be used to connect these two devices:

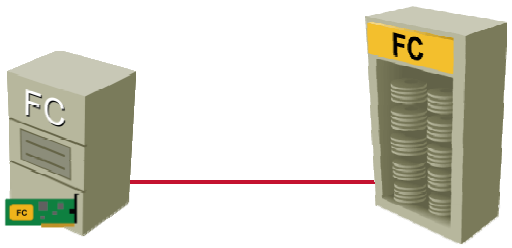
- Point-to-point
- Fibre Channel Arbitrated Loop (FC-AL)
- Switched fabric network, usually referred to as the SAN

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

What Is the Point-to-Point Topology?

- Dedicated one-to-one connection between two nodes
- DAS architecture but with better performance and flexibility than SCSI



© 2012 Cisco and/or its affiliates. All rights reserved.

The point-to-point topology is the simplest Fibre Channel storage configuration. This topology is often referred to as direct-attached storage (DAS) architecture. Originally, this architecture was achieved by using a SCSI bus, but SCSI has been superseded by Fibre Channel cabling.

As the name suggests, a point-to-point configuration is a one-to-one connection between a host and a storage device. In this architecture, the storage is dedicated to the server. This design does not provide a very scalable solution, but does provide very good security because no other device has access to the same storage device.

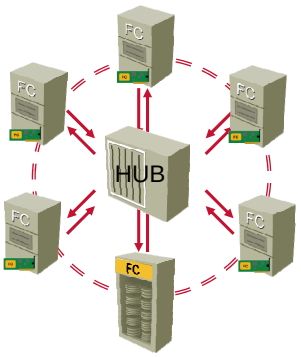
The figure illustrates a server that is connected to a storage device through a Fibre Channel link.

https://t.me/learningnets

What Is the Arbitrated Loop Topology?

Limitations of arbitrated loops:

- Performance:
  - One data path means only one pair of devices can communicate at a time (shared bandwidth)
  - Higher latency than fabrics
- Scalability:
  - 127 addressable ports:
    - 126 available for nodes
    - 1 reserved for fabric-attach port
  - About a dozen nodes in practice
- Reliability:
  - If one device fails, the entire loop can fail



© 2012 Cisco and/or its affiliates. All rights reserved.

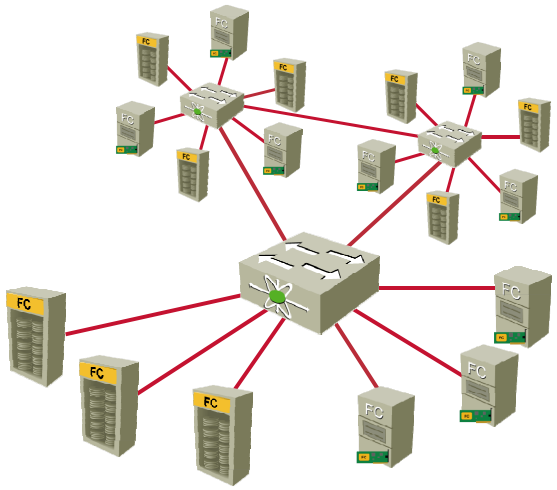
FC-AL is more scalable than the point-to-point topology. In the FC-AL topology, all servers have access to all storage devices, allowing storage resources to be better utilized. The I/O speed between the initiator and target depends on their locations in the loop. Although FC-AL is scalable to 127 devices, typically 12 or fewer devices are attached.

These are the key limitations of the FC-AL topology:

- Loops suffer from poor performance. Because there is only one data path, only one pair of devices can communicate at any given time. Therefore, all the devices on the loop share the available bandwidth.
- Loops have a higher latency than other topologies. Devices must negotiate for control of the loop before they can send any traffic to their target devices.
- Loops are not very scalable. The FC-AL protocol provides for only 127 unique addresses, one of which is reserved for attaching the loop to a Fibre Channel switched fabric. The other 126 addresses are usable by nodes (servers and storage devices). In practice, only about a dozen devices can be connected to the loop before performance drops below acceptable levels.
- Loop configurations are susceptible to device failures. If one device fails on the loop, many devices are affected.

https://t.me/learningnets

What Is the Switched Fabric Topology?



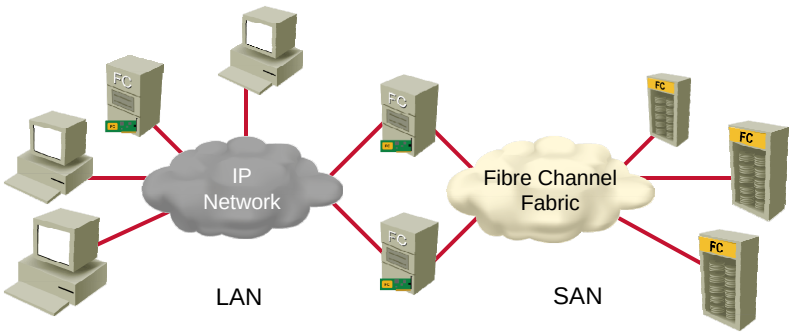
© 2012 Cisco and/or its affiliates. All rights reserved.

The most scalable topology for initiator-to-target connectivity is the switched fabric topology, often referred to as the SAN or fabric. The switched fabric topology incorporates one or more high-bandwidth Fibre Channel switches to manage data traffic between multiple host and storage devices.

https://t.me/learningnets

Fibre Channel Overview

- Fibre Channel is a technology for transporting data between devices.
- Fibre Channel is the transport technology most commonly used for SANs.



© 2012 Cisco and/or its affiliates. All rights reserved.

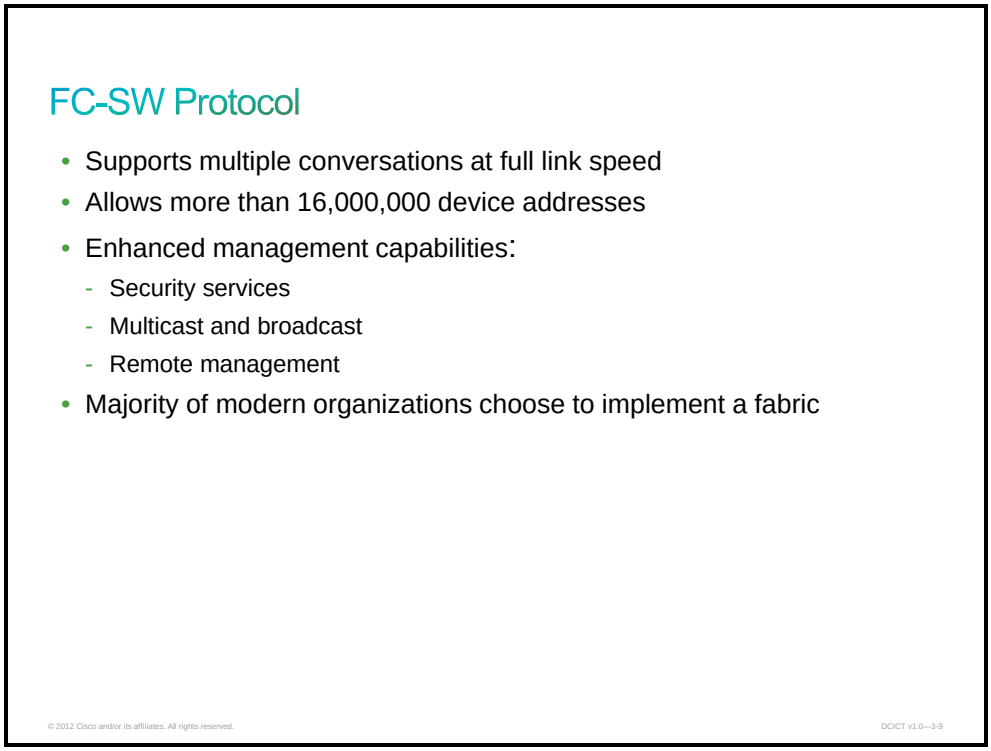
Fibre Channel is a technology for transporting data between devices. Fibre Channel is the network interconnect technology that is commonly used for SANs. Traditional storage technologies, such as SCSI, are designed for controlled, local environments. They support few devices and only short distances, but they deliver data quickly and reliably. Traditional data network technologies, such as Ethernet, were originally designed for distributed environments. Ethernet supports many devices and long distances, but delivery of data can be delayed and often relies on upper-layer protocols (ULPs) to provide reliability and redelivery of traffic.

Fibre Channel combines the best of the SCSI and Ethernet worlds. Fibre Channel is designed to support many devices and longer distances, plus it provides the reliable data delivery that the SCSI protocol requires between an initiator and target.

Servers often connect to both the LAN and the SAN, as depicted in the figure.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



The Fibre Channel Switched Fabric (FC-SW) protocol differs from the FC-AL topology in several important ways:

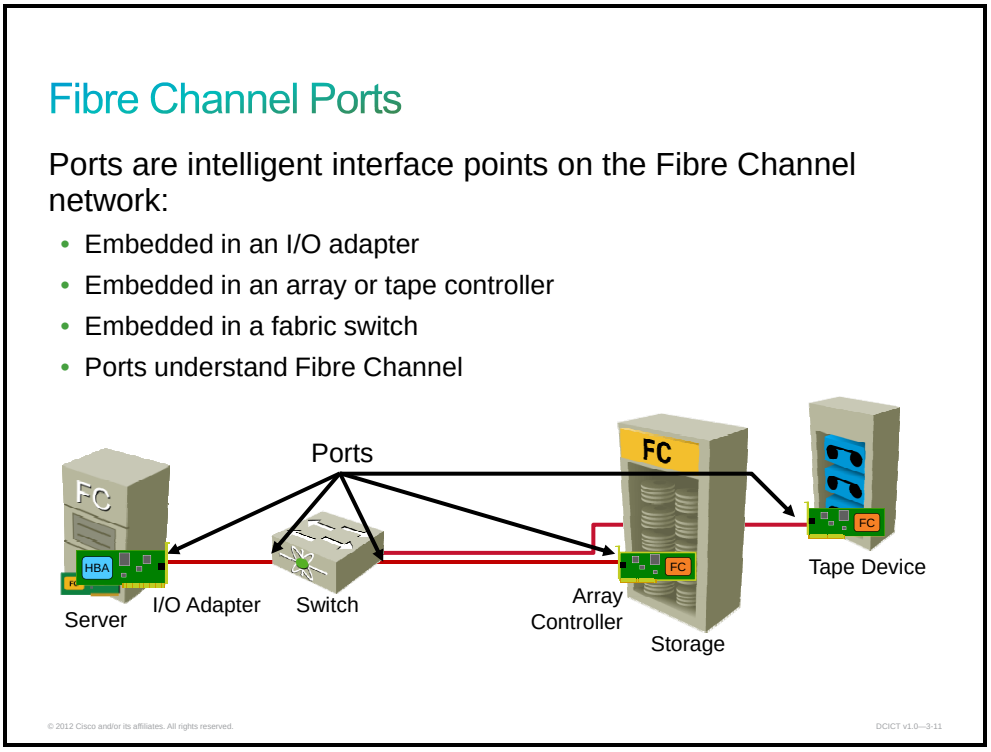
- Switches can support multiple simultaneous conversations. Each conversation between two devices can use the full link bandwidth.
- The FC-SW protocol device-addressing scheme allows more than 16 million ports to connect to the SAN. Existing implementations can support hundreds or thousands of nodes by using large director-class switches.
- The FC-SW protocol defines several management services that increase the scalability, manageability, and security of the SAN.

Because of the limitations of the FC-AL topology, most modern organizations choose to implement a switched fabric network topology, which offers greater scalability, performance, reliability, and manageability.

https://t.me/learningnets

Fibre Channel Port Types

This topic looks at the different Fibre Channel ports that can be found in the SAN.



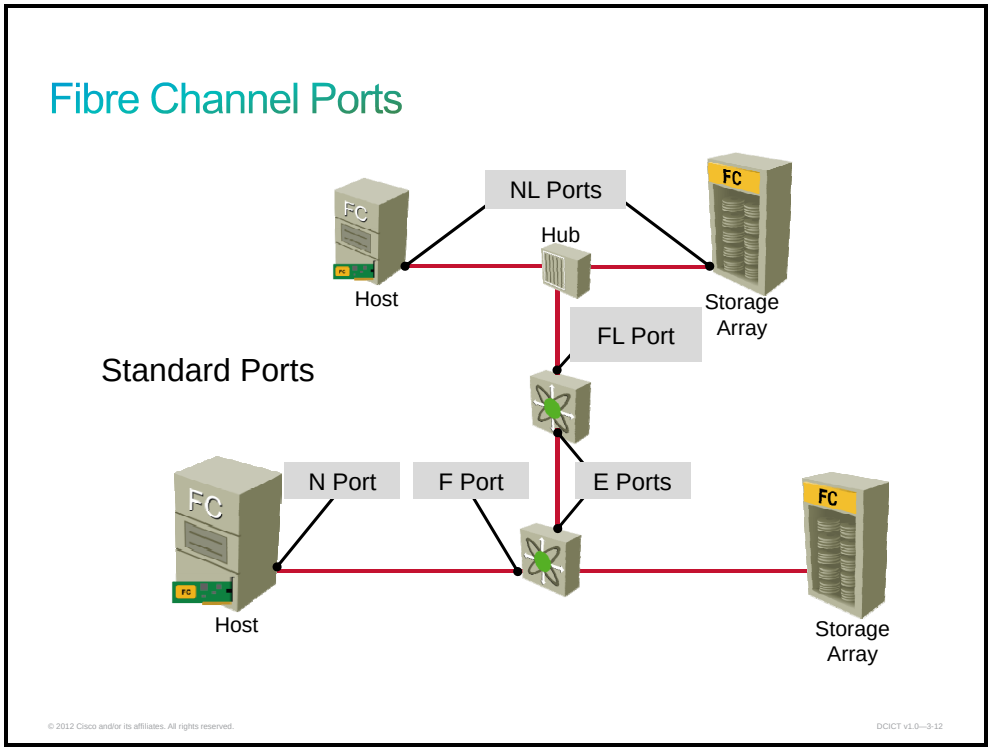
Fibre Channel ports are intelligent interface points on the Fibre Channel SAN. They are found embedded in various devices:

- I/O adapter
- Array or tape controller
- Switched fabric

When a server or storage device communicates, the interface point that acts as the initiator or target for the connection. The server or storage device issues SCSI commands, which the interface point formats to be sent to the target device.

These ports understand Fibre Channel. The Fibre Channel switch recognizes which type of device is attaching to the SAN and configures the ports accordingly.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



As depicted in the figure, there are various Fibre Channel port types:

- **Expansion port (E Port):** In E Port mode, an interface functions as a fabric expansion port. This port connects to another E Port to create an interswitch link (ISL) between two switches. E Ports carry frames between switches for configuration and fabric management. They also serve as a conduit between switches for frames that are destined to remote node ports (N Ports) and node loop ports (NL Ports). E Ports support class 2, class 3, and class F service.
- **Fabric port (F Port):** In F Port mode, an interface functions as a fabric port. This port connects to a peripheral device (such as a host or disk) that operates as an N Port. An F Port can be attached to only 1 N Port. F Ports support class 2 and class 3 service.
- **Fabric loop port (FL Port):** In FL Port mode, an interface functions as a fabric loop port. This port connects to one or more NL Ports (including FL Ports in other switches) to form a public FC-AL. If more than one FL Port is detected on the FC-AL during initialization, then only one FL Port becomes operational; the other FL Ports enter nonparticipating mode. FL Ports support class 2 and class 3 service.
- **Trunking expansion port (TE Port):** In TE Port mode, an interface functions as a trunking expansion port. This port connects to another TE Port to create an extended ISL (EISL) between two switches. TE Ports are specific to Cisco MDS 9000 Series Switches and expand the functionality of E Ports to support virtual SAN (VSAN) trunking, transport quality of service (QoS) parameters, and the Fibre Channel Traceroute (fctrace) feature. When an interface is in TE Port mode, all frames that are transmitted are in the EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link.
- **Node-proxy port (NP Port):** An NP Port is a port on a device that is in N-Port Virtualization (NPV) mode and connects to the core switch via an F Port. NP Ports function like node ports (N Ports) but in addition to providing N Port operations, they also function as proxies for multiple physical N Ports.

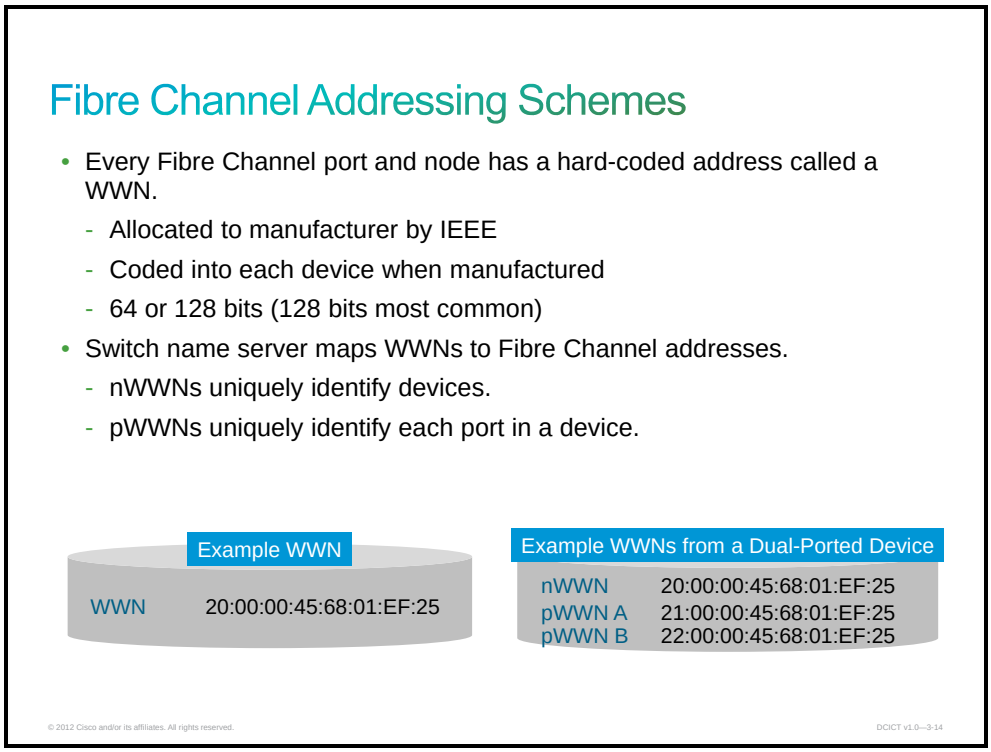
https://t.me/learningnets

- **Trunking fabric port (TF Port):** In TF Port mode, an interface functions as a trunking expansion port. This interface connects to another trunking node port (TN Port) or trunking node-proxy port (TNP Port) to create a link between a core switch and an NPV switch or a host bus adapter (HBA) to carry tagged frames. TF Ports are specific to Cisco MDS 9000 Series Switches and expand the functionality of F Ports to support VSAN trunking. In TF Port mode, all frames are transmitted in the EISL frame format, which contains VSAN information.
- **TNP Port:** In TNP Port mode, an interface functions as a trunking expansion port. This interface connects to a TF Port to create a link to a core N Port ID Virtualization (NPIV) switch from an NPV switch to carry tagged frames.
- **Switched Port Analyzer (SPAN) destination port (SD Port):** In SD Port mode, where an interface functions as a SPAN. The SPAN feature is specific to Cisco MDS 9000 Series Switches. An SD Port monitors network traffic that passes through a Fibre Channel interface. Monitoring is performed using a standard Fibre Channel analyzer (or a similar SwitchProbe) that is attached to the SD Port. SD Ports cannot receive frames and only transmit a copy of the source traffic. This feature is nonintrusive and does not affect switching of network traffic for any SPAN source port.
- **SPAN tunnel port (ST Port):** In ST Port mode, an interface functions as an entry-point port in the source switch for the Remote SPAN (RSPAN) Fibre Channel tunnel. ST Port mode and the RSPAN feature are specific to Cisco MDS 9000 Series Switches. When a port is configured as an ST Port, it cannot be attached to any device and therefore cannot be used for normal Fibre Channel traffic.
- **Fx Port:** An interface that is configured as an Fx Port can operate in either F or FL Port mode. Fx Port mode is determined during interface initialization, depending on the attached N or NL Port.
- **Bridge port (B Port):** Whereas E Ports typically interconnect Fibre Channel switches, some SAN extender devices implement a B Port model to connect geographically dispersed fabrics. This model uses B Ports that are as described in the T11 Standard Fibre Channel Backbone 2 (FC-BB-2).
- **Auto mode:** An interface that is configured in auto mode can operate in one of the following modes: F Port, FL Port, E Port, TE Port, or TF Port, with the port mode being determined during interface initialization.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Fibre Channel Addressing

Fibre Channel uses World Wide Names (WWNs) and Fibre Channel IDs (FCIDs). WWNs are unique identifiers that are hardcoded into Fibre Channel devices. FCIDs are dynamically acquired addresses that are routable in a switch fabric.



Each Fibre Channel port has at least one WWN. Vendors buy blocks of WWNs from the IEEE and allocate them to devices in the factory.

WWNs are important for enabling Cisco Fabric Services because they have these characteristics:

- Guaranteed to be globally unique
- Permanently associated with devices

These characteristics ensure that the fabric can reliably identify and locate devices. This ability is an important consideration for Cisco Fabric Services. When a management service or application needs to quickly locate a specific device, the following occurs:

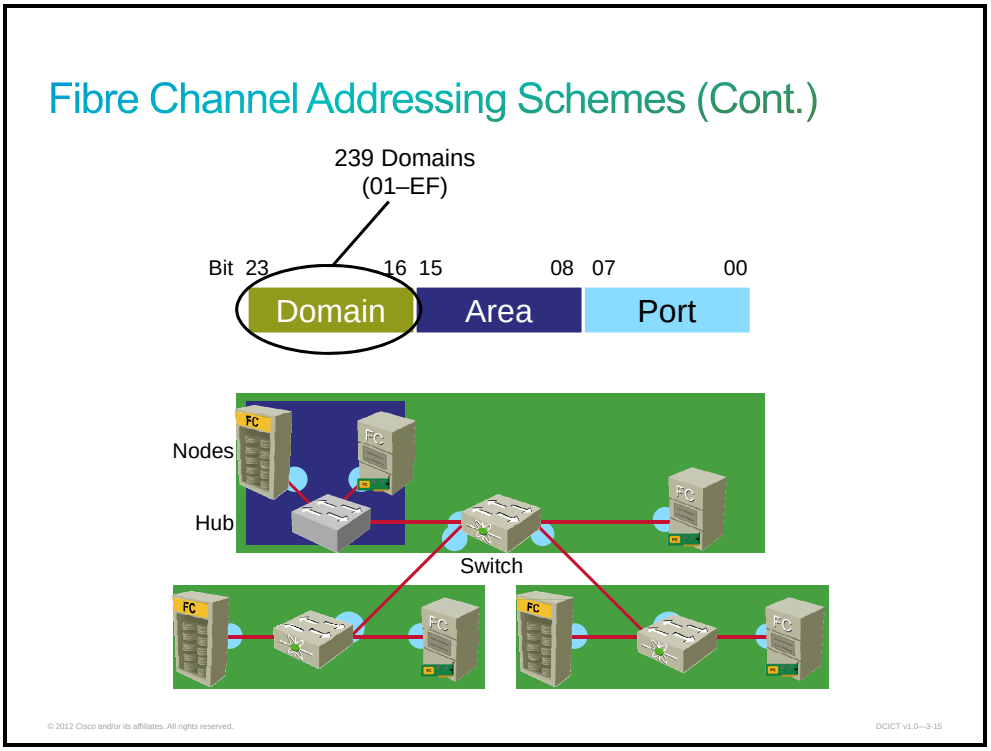
1. The service or application queries the switch name server service with the WWN of the target device.
2. The name server looks up and returns the current port address that is associated with the target WWN.
3. The service or application communicates with the target device, using the port address.

The two types of WWNs are node WWNs (nWWNs) and port WWNs (pWWNs):

- nWWNs uniquely identify devices. Every HBA, array controller, switch, gateway, and Fibre Channel disk drive has a single unique nWWN.
- pWWNs uniquely identify each port in a device. A dual-ported HBA has three WWNs: one nWWN, and one pWWN for each port.

https://t.me/learningnets

nWWNs and pWWNs are both needed because devices can have multiple ports. On a single-port device, the nWWN and pWWN may be the same. On multiport devices, the pWWN is used to uniquely identify each port. Ports must be uniquely identifiable because each port participates in a unique data path. nWWNs are required because the node itself must sometimes be uniquely identified. For example, path failover and multiplexing software can detect redundant paths to a device by observing that the same nWWN is associated with multiple pWWNs.



The Fibre Channel point-to-point topology uses a 1-bit addressing scheme. One port assigns itself an address of 0x000000 and then assigns the other port an address of 0x000001.

The FC-AL topology uses an 8-bit addressing scheme:

- The arbitrated loop physical address (AL-PA) is an 8-bit address, which provides 256 potential addresses. However, only a subset of 127 addresses is available because of the 8B/10B encoding requirements.
- One address is reserved for an FL Port, so 126 addresses are available for nodes.
- Addresses are cooperatively chosen during loop initialization.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Switched Fabric Address Space

The 24-bit Fibre Channel address consists of three 8-bit elements:

- The domain ID is used to define a switch. Each switch receives a unique domain ID.
- The area ID is used to identify groups of ports within a domain. Areas can be used to group ports within a switch. Areas are also used to uniquely identify fabric-attached arbitrated loops. Each fabric-attached loop receives a unique area ID.

Although the domain ID is an 8-bit field, only 239 domains are available to the fabric:

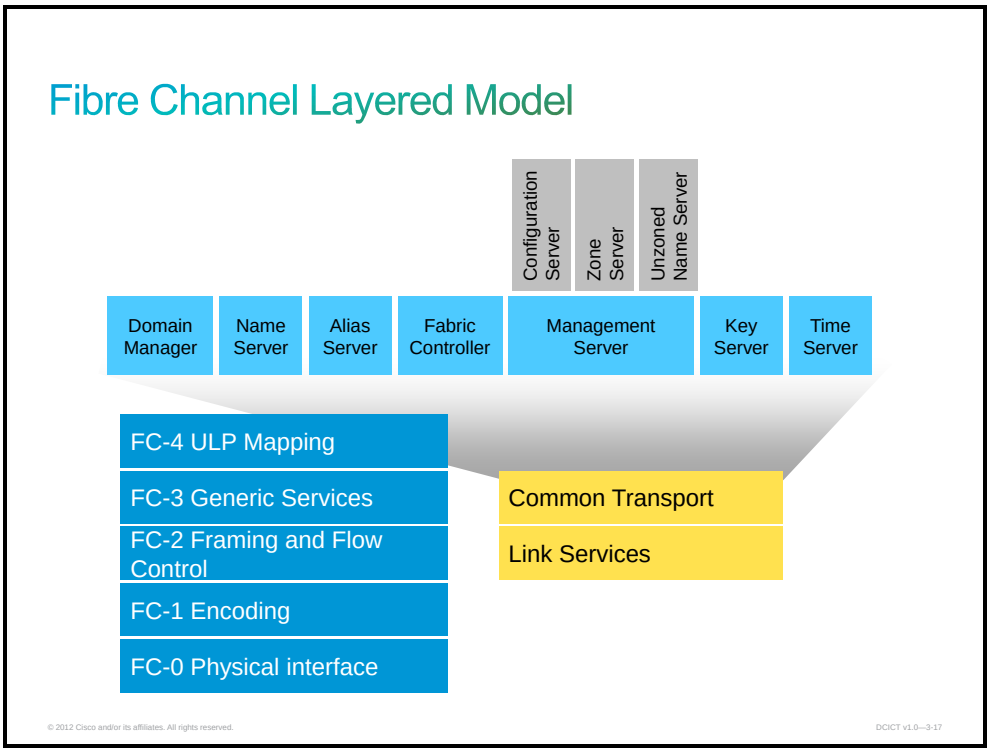
- Domains 01 through EF are available.
- Domains 00 and F0 through FF are reserved for use by switch services.

Each switch must have a unique domain ID, so there can be no more than 239 switches in a fabric. In practice, the number of switches in a fabric is far smaller because of storage-vendor qualifications.

https://t.me/learningnets

Fibre Channel Layered Model

The Fibre Channel Protocol (FCP) comprises of five layers, FC-0 through FC-4. This topic examines these layers and some of the Fibre Channel Generic Services (FC-GS) that are provided to the SAN infrastructure.



These are the main functions of each Fibre Channel layer:

- **FC-4 upper-layer protocol (ULP) mapping:** Provides protocol mapping to identify the ULP that is encapsulated into a protocol data unit (PDU) for delivery to the FC-2 layer
- **FC-3 generic services:** Provides the Fibre Channel Generic Services (FC-GS) that are required for fabric management
- **FC-2 framing and flow control:** Provides the framing and flow control that is required to transport the ULP over Fibre Channel
- **FC-1 encoding:** Provides line coding of signals
- **FC-0 physical interface:** Provides physical connectivity, including cabling, connectors, and so on

The Fibre Channel Switch Fabric 2 (FC-SW-2) specification defines several services that are required for fabric management:

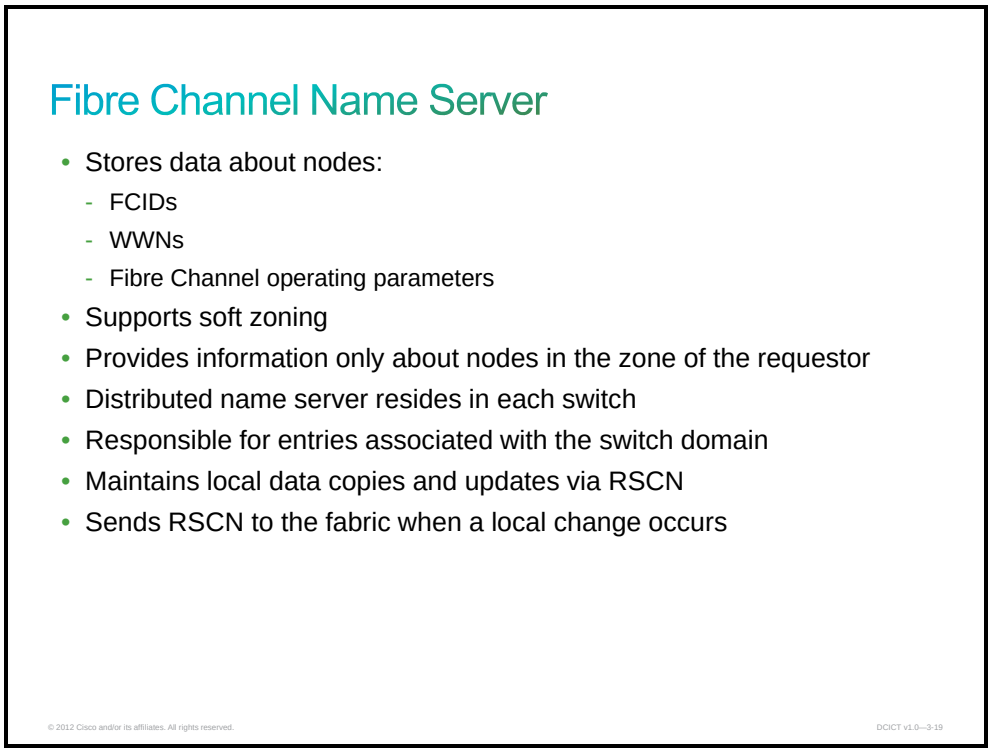
- Name Service
- Login Service
- Address Manager
- Alias Server
- Fabric Controller
- Management Server
- Key Distribution Server
- Time Server

The FC-SW-2 specification does not require switches to implement all these services. Some services can be implemented as an external server function.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

FCNS and the FLOGI Process

The Fibre Channel Name Server (FCNS) is a distributed database that the switch implements. The fabric login (FLOGI) process is the process by which a device logs in to the fabric. This topic examines these services.



The FCNS is a distributed service that implements a database in the fabric for storing information about each node:

- FCIDs
- pWWNs and nWWNs
- Fibre Channel operating parameters, such as ULPs and classes of service

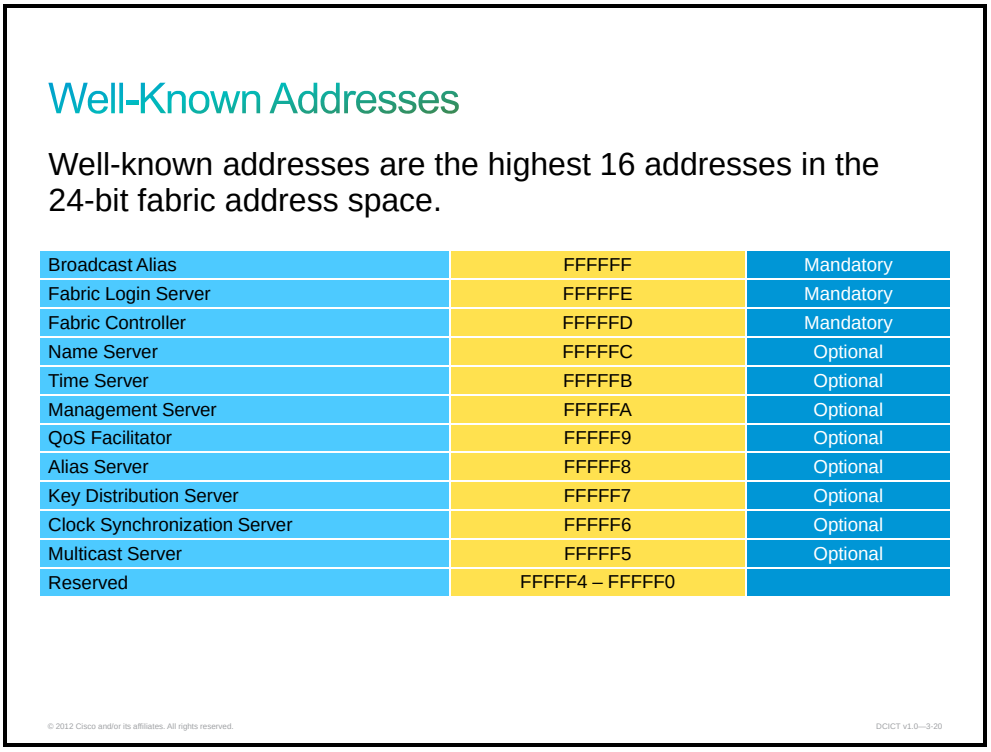
The FCNS supports soft zoning by performing WWN lookups to verify zone membership. Zoning is enforced by providing information only about nodes that are in the zone of the requestor. Management applications that need to obtain information about the fabric use the FCNS.

Each switch in a fabric contains its own resident name server, which is called a distributed name server. Each distributed name server within a switch is responsible for the name entries that are associated with the domain that is assigned to the switch. The distributed name server instances synchronize their databases by using the Registered State Change Notification (RSCN) process.

When a client Nx port wants to query the name server, it submits a request to its local name server via the well-known address for the name server. If the required information is not available locally, then the distributed name server within the local switch responds to the request by making the necessary inquiries of other distributed name server instances in the other switches. The communication between switches that is performed to acquire the requested information is transparent to the original requesting client.

https://t.me/learningnets

Partial responses to distributed name server queries are allowed. If an entry switch sends a partial response back to an Nx port, it must set the partial response bit in the Fibre Channel header.

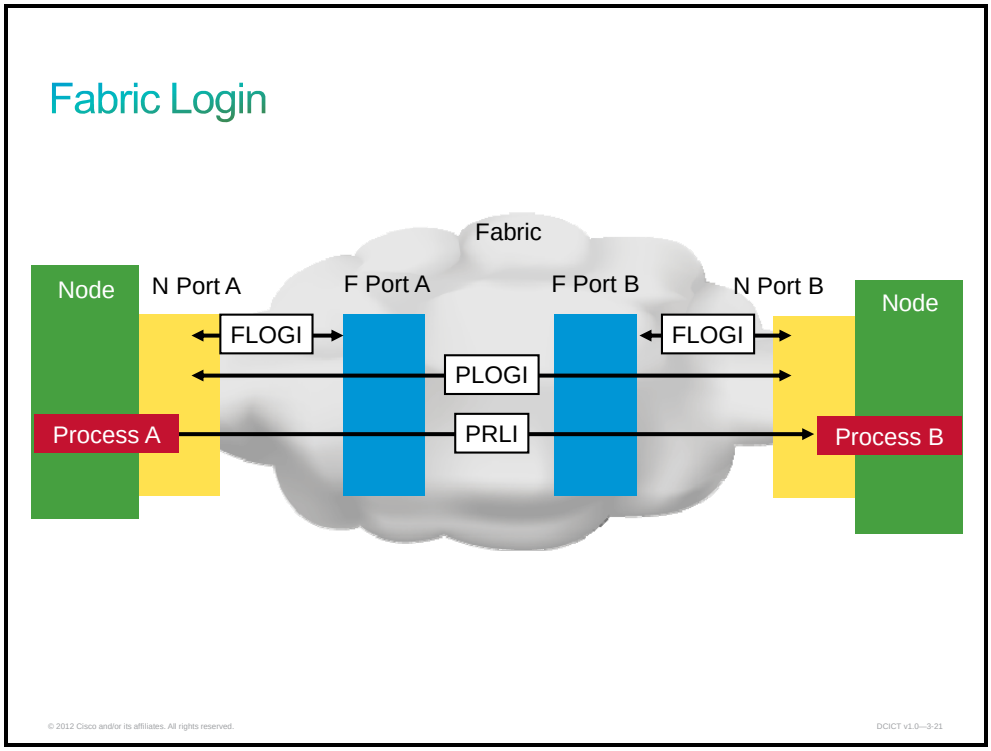


Well-known addresses allow devices to reliably access switch services. All services are addressed in the same way that an N Port is addressed. Nodes communicate with services by sending and receiving Extended Link Services (ELS) commands (frames) to and from well-known addresses.

Well-known addresses are the highest 16 addresses in the 24-bit fabric address space.

© 2012 Cisco Systems, Inc.

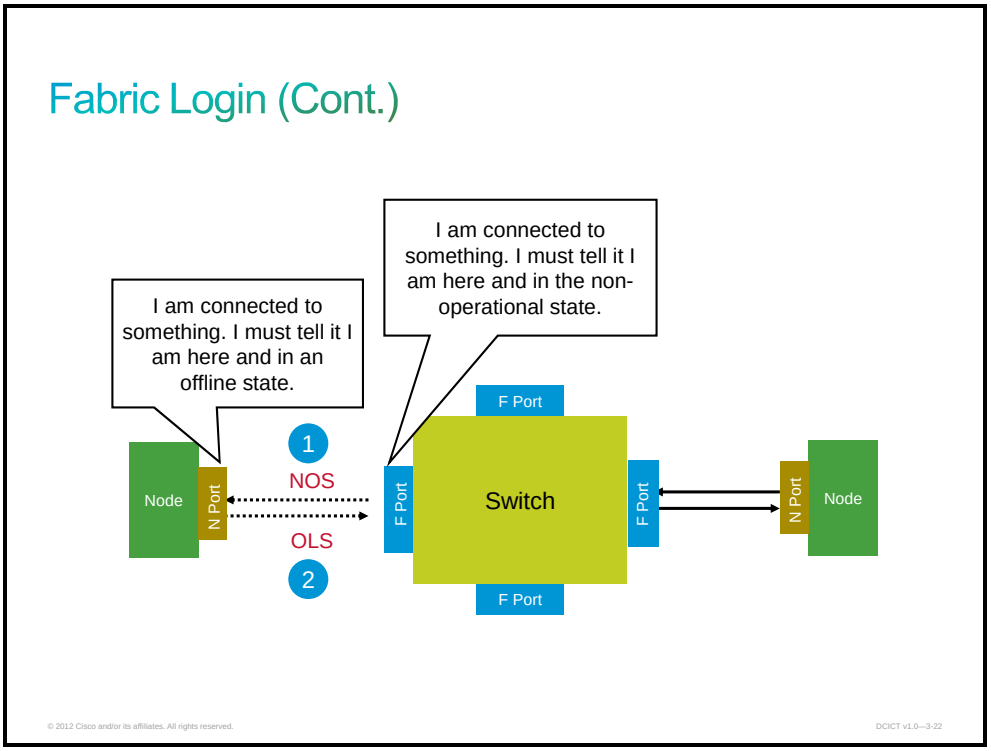
https://t.me/learningnets



Before an N Port can begin exchanging data with other N Ports, three processes must occur:

4. The N Port must log in to its attached F Port. This process is known as the FLOGI.
5. The N Port must log in to its target N Port. This process is known as the port login (PLOGI).
6. The N Port must exchange information about ULP support with its target N Port, to ensure that the initiator and target process can communicate. This process is known as the process login (PRLI).

https://t.me/learningnets

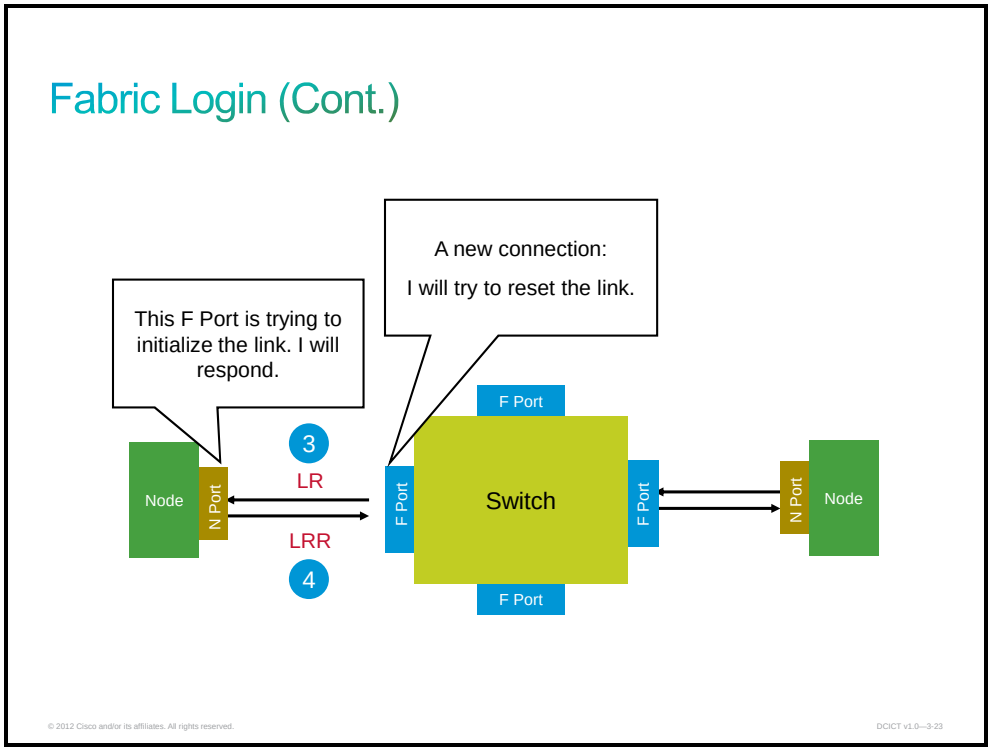


FLOGI is the initial bootstrap process that occurs when an N Port is connected to an F Port. FLOGI is mandatory for N Ports and optional for NL Ports. The N Port uses FLOGI to discover whether a fabric is present. Communication with other N Ports may not be attempted until the FLOGI process is complete.

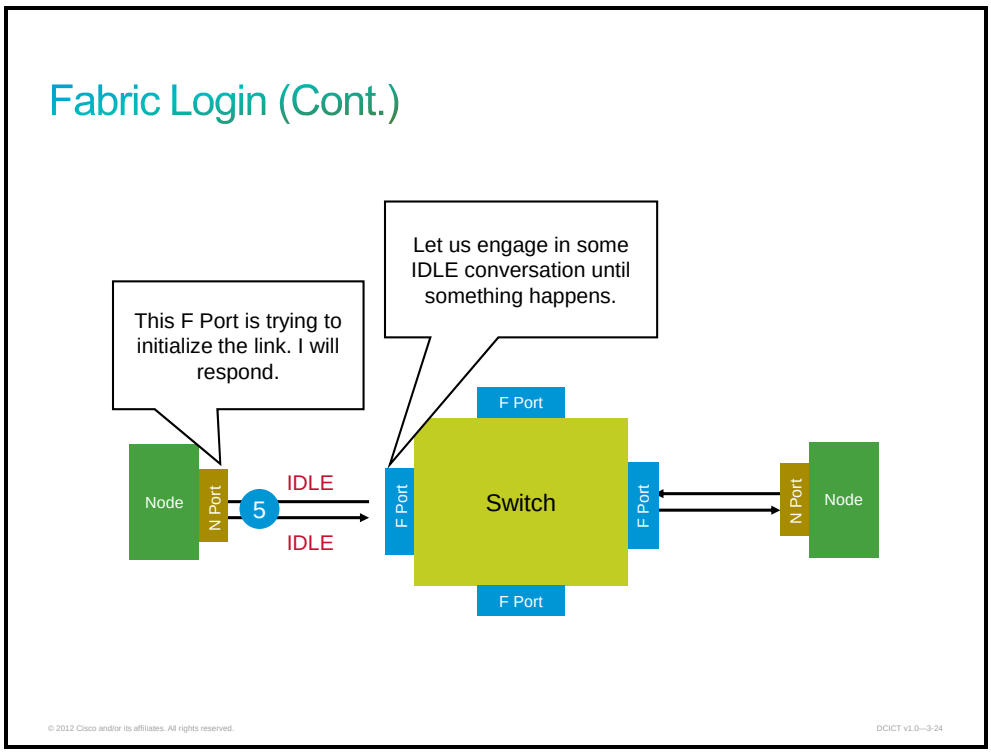
The FLOGI protocol follows this process:

1. The F Port sends a primitive not-operational sequence to the N Port.
2. When the N Port receives the not-operational sequence, it responds with a primitive offline state sequence to begin link initialization.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



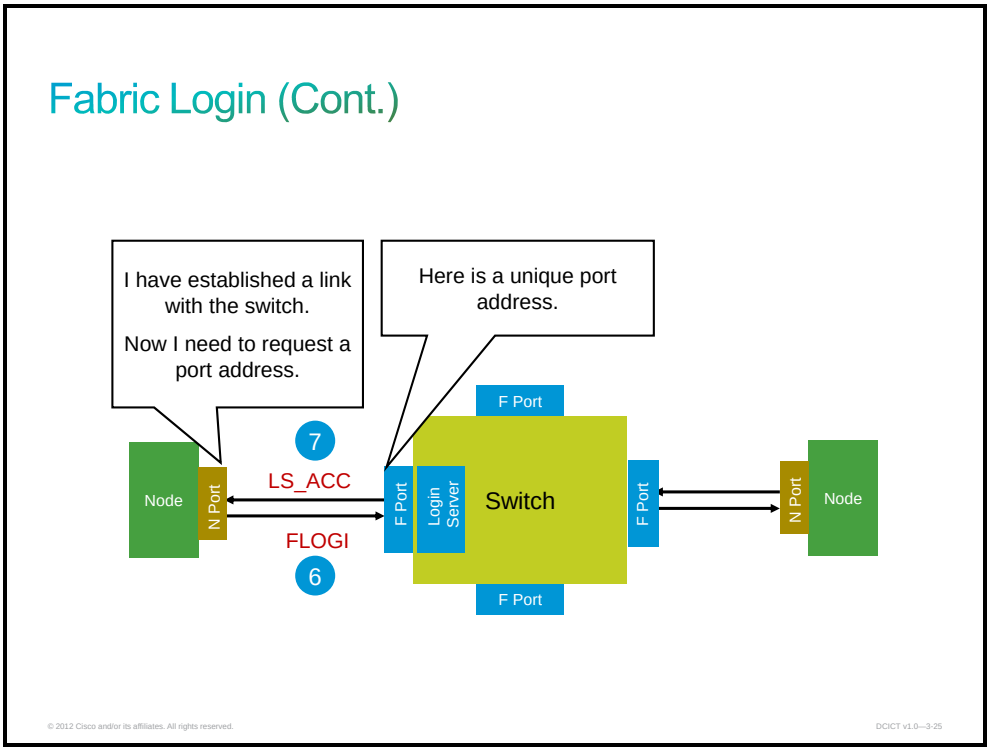
3. After the N Port begins the initialization process by sending an offline state sequence, the F Port tries to reset the port by sending a link reset (LR) command.
4. The N Port responds with a link reset response (LRR) command.



5. From this point on, the link is active and IDLE fill words flow in both directions on the link.

Following link initialization, a new N Port uses a Source ID (SID) of 0x000000 or AL-PA of 0x0000 to indicate that the port is unidentified during the FLOGI. An existing N Port uses its existing port address as its SID.

https://t.me/learningnets

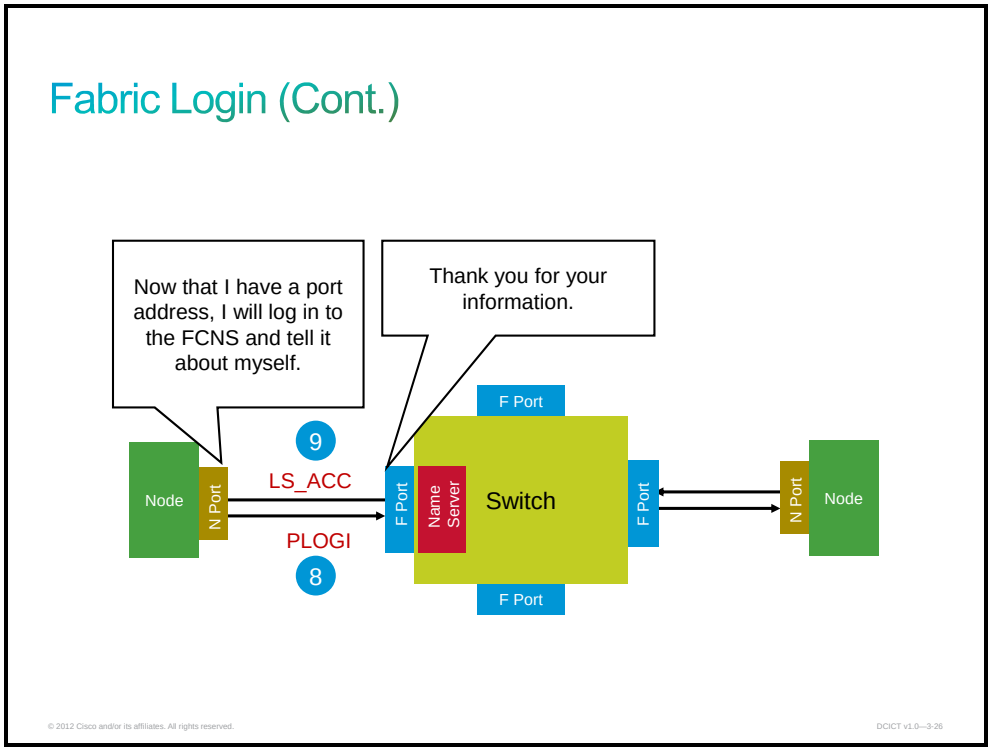


6. After the N Port has established a link to its F Port, the N Port obtains a port address. The N Port performs this process by sending a FLOGI link services command to the switch login server (at well-known address 0xFFFFFE).

7. The login server sends an accept (ACC) reply that contains the N Port address in the Destination ID field.

When an N Port performs a FLOGI and receives an ACC frame, that action indicates that the ACC came from another N Port. Then, the N Port that is logging in assumes that it is in a point-to-point configuration. In this case, the N Port immediately initiates PLOGI with the other N Port after completing the FLOGI.

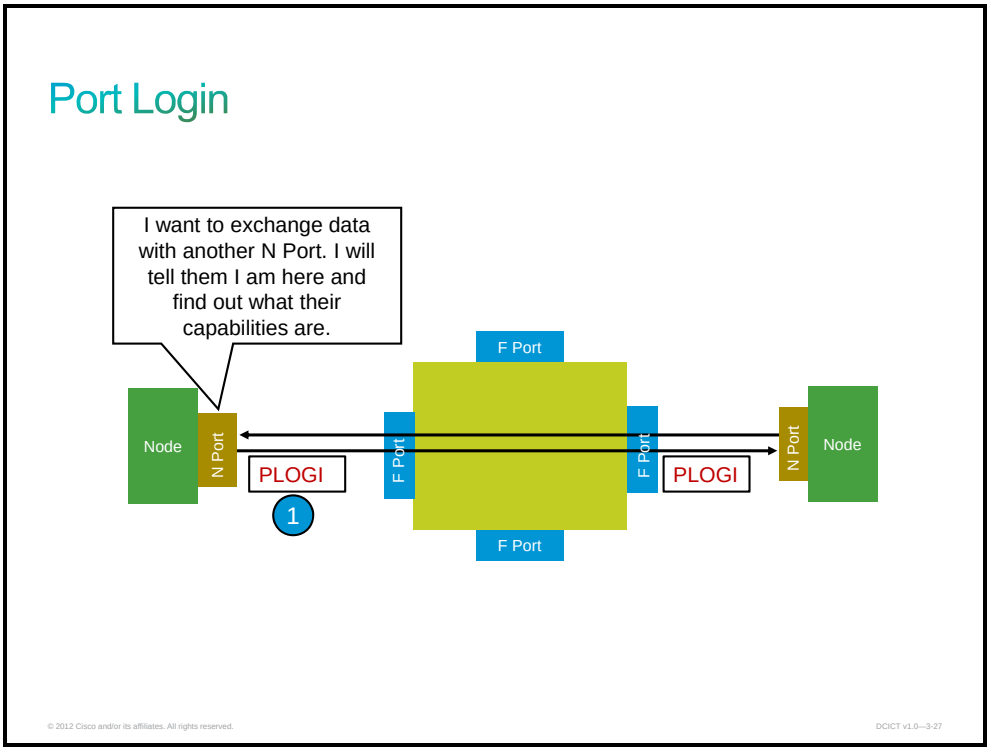
© 2012 Cisco Systems, Inc.
https://t.me/learningnets



8. After receiving a port address, the N Port logs in to the fabric name server at address 0xFFFFFC. The N Port transmits its service parameters, such as the number of buffer credits it supports, its maximum payload size, and supported CoSs.

9. The name server responds with a link service accept (LS_ACC) frame.

https://t.me/learningnets



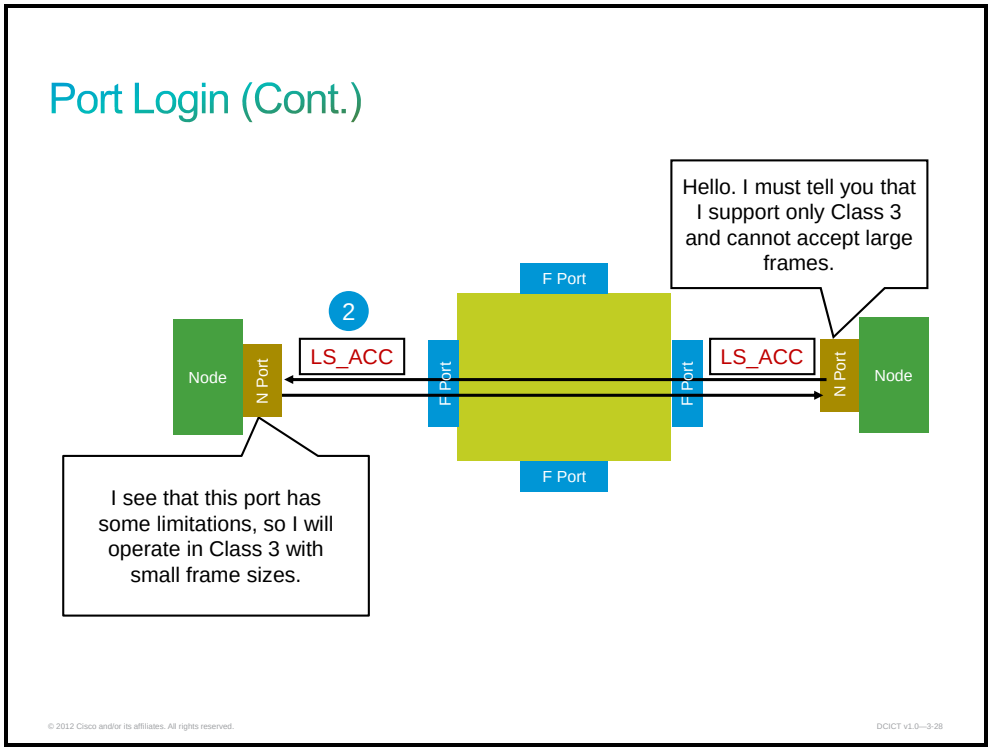
PLOGI enables the requested N Port to create a communication channel with the remote N Port with which it wants to communicate, by setting and exchanging operational parameters.

After completing the FLOGI process, the N Port can log in to another N Port by using the PLOGI protocol. PLOGI must be completed before the nodes can perform a ULP operation.

The PLOGI protocol follows this process:

1. The initiator N Port sends a PLOGI frame that encapsulates the N Port operating parameters in the payload.

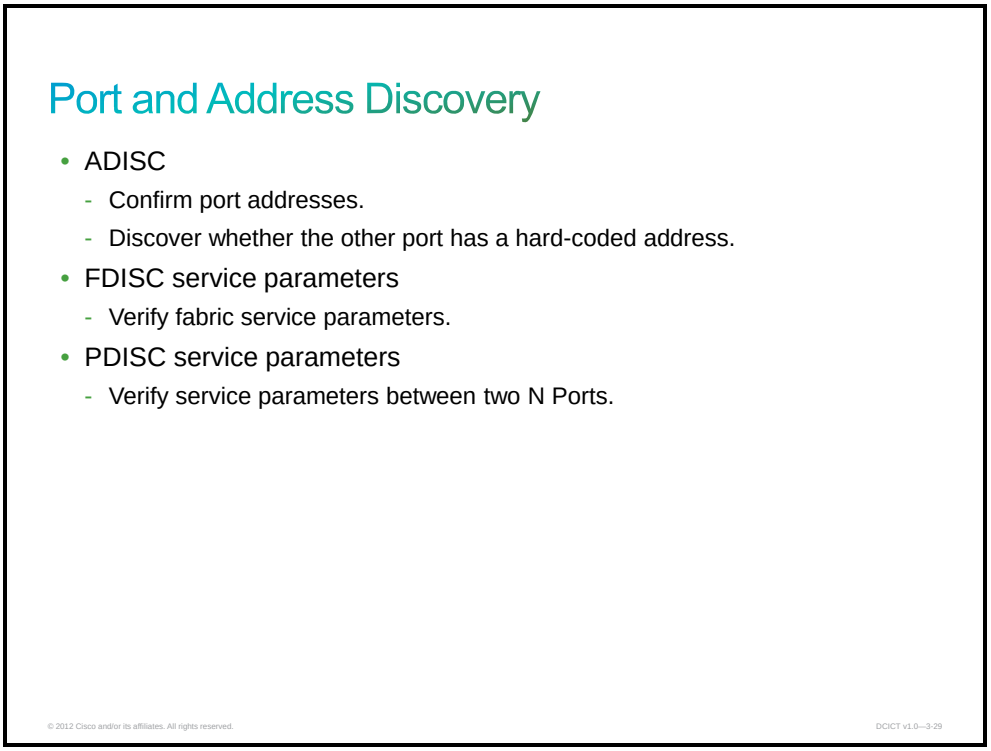
© 2012 Cisco Systems, Inc.
https://t.me/learningnets



2. The target N Port responds to the initiator N Port by sending an ACC frame that specifies the target N Port operating parameters. The operating system driver that manages the initiator N Port stores this information in a parameter block.

An N Port can be logged in to multiple N Ports simultaneously. N Ports typically perform a process logout (PRLO) only when one of the nodes goes offline.

https://t.me/learningnets

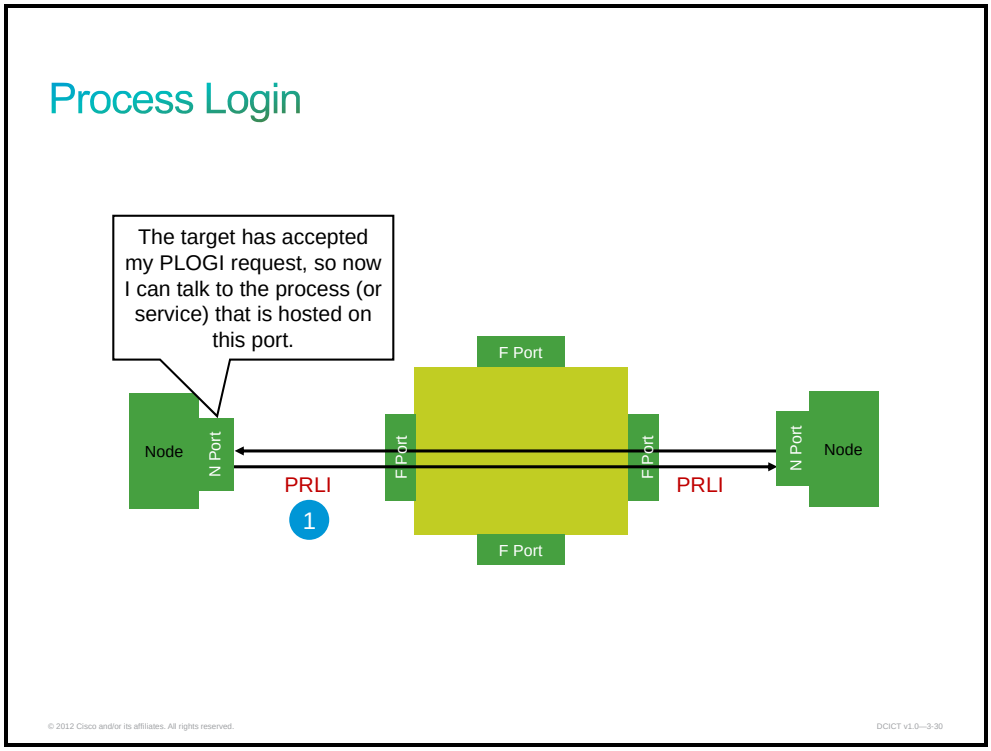


After FLOGI and PLOGI are complete, the N Port can use the following ELS commands. These commands retrieve updated information or verify information about port addresses and service parameters:

- Discover Address (ADISC) can be used to confirm the address of another port or to discover whether the other port has a hardcoded address.
- Discover Fabric (FDISC) service parameters can be used to verify Cisco Fabric Service parameters.
- Discover Port (PDISC) service parameters can be used to verify the service parameters of another N Port.

These commands allow ports to query and verify fabric and port parameters without performing PLOGI and thus forcing a logout of the current session.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



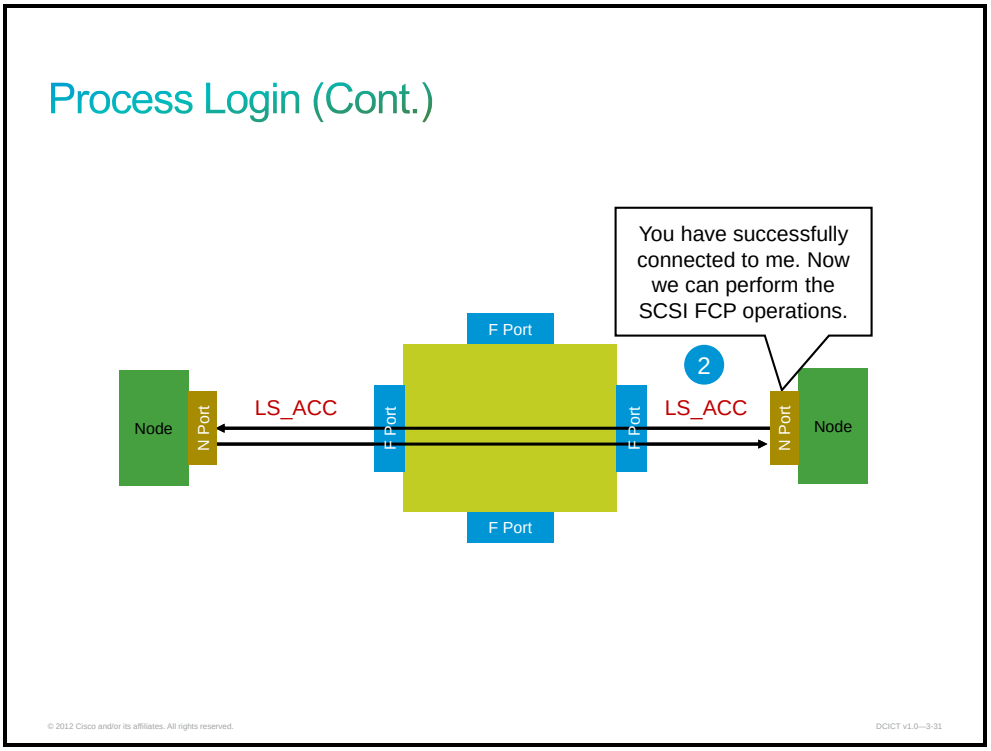
PRLI is carried out by using the ELS **PRLI** command. PRLI allows one or more images at one N Port to be related to corresponding images at another N Port, creating an image pair.

After completing the PLOGI protocol, each N Port knows about the Fibre Channel operating parameter capabilities of the other. At this point, the driver for the initiator port can use the PRLI protocol to open a channel with the driver for the target port. The PRLI protocol is used to establish a session between two FC-4 level logical processes.

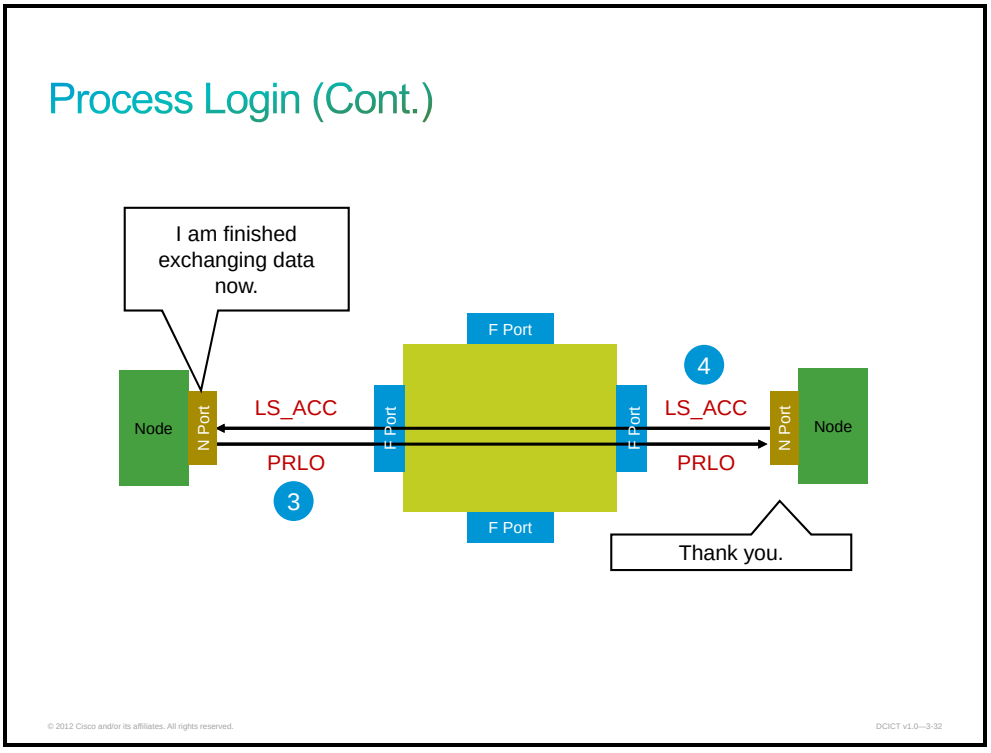
The PRLI protocol follows this process:

1. The initiator sends a PRLI frame that contains information about its ULP support.

https://t.me/learningnets



2. The target port responds with an ACC frame that contains details about its ULP support. At this point, a channel has been successfully opened and communication can take place. The relationship between the initiator process and the target process is known as an image pair.

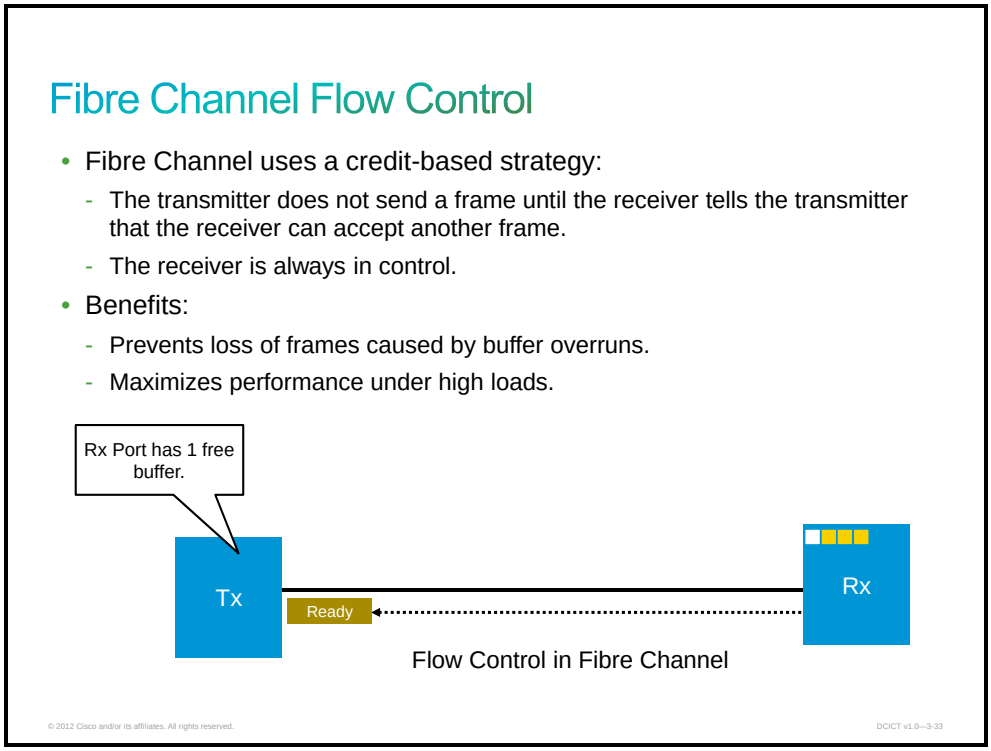


3. When the initiator has finished exchanging data with the target, the initiator sends a PRLO frame.
4. The target responds with an ACC frame, and the image pair is then terminated.

At this point, the image pair must be established again before further communication can take place.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



In a Fibre Channel environment, it is important to ensure that frames are not lost between the initiator and target. Therefore, a lossless environment must be provided. In Fibre Channel, the receiver is in control rather than the sender, as in an Ethernet network.

To improve performance under high traffic loads, Fibre Channel uses a credit-based flow-control strategy. The receiver must issue a credit for each frame that the transmitter sends, before that frame can be sent.

A credit-based strategy ensures that the Rx Port is always in control. The Rx Port must issue a credit for each frame that the transmitter sends. This strategy prevents frames from being lost when the Rx Port runs out of free buffers. Preventing lost frames maximizes performance under high-traffic load conditions because the Tx Port does not need to resend frames.

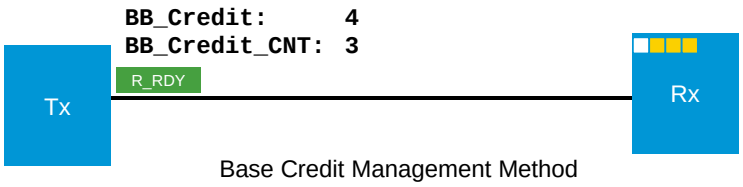
The figure shows a credit-based flow-control process:

- The Tx Port counts the number of free buffers at the Rx Port.
- Before the Tx Port can send a frame, the Rx Port must notify the Tx Port that the Rx Port has a free buffer and is ready to accept a frame. When the Tx Port receives the notification, it increments its count of the number of free buffers at the Rx Port.
- The Tx Port sends frames only when it knows that the Rx Port can accept them.

https://t.me/learningnets

Fibre Channel Flow Control (Cont.)

1. At login, Rx tells Tx how many buffers Rx has (BB_Credit)
2. Tx sets BB_Credit_CNT = 0 at login.
3. Tx increments BB_Credit_CNT when it sends a frame.
4. Rx sends an R_RDY when it processes the frame.
5. Tx decrements BB_Credit_CNT when it receives R_RDY. Tx sends only when BB_Credit_CNT < BB_Credit.



Base Credit Management Method

© 2012 Cisco and/or its affiliates. All rights reserved.

The base credit management method works in this way:

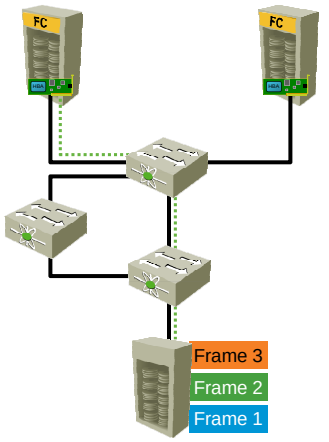
- When the Tx Port sends a PLOGI request, the Rx Port responds with an ACC frame. The ACC frame includes information on the size and numbers of frame buffers that it has (BB_Credit). The Tx Port stores the BB_Credit value in a table.
- The Tx Port also stores another value called buffer-to-buffer credit count (BB_Credit_CNT), which represents the number of used buffer credits. BB_Credit_CNT is set to zero after the ports complete the login process.
- Each time the Tx Port sends a frame, it increments BB_Credit_CNT.
- Upon receiving the frame, the Rx Port processes the frame and moves it to ULP buffer space. The Rx Port then sends a receiver ready (R_RDY) acknowledgment signal back to the Tx Port, informing it that a buffer is available.
- When the Tx Port receives the R_RDY signal, it then decrements its BB_Credit_CNT.

The Tx Port never allows BB_Credit_CNT (the count of frames that have not yet been acknowledged) to exceed BB_Credit (the total number of buffers in the Rx Port). In other words, if the Tx Port cannot confirm that the Rx Port has a free buffer, it stops sending frames.

https://t.me/learningnets

Fabric Shortest Path First

- Computes the least-cost path through the fabric, based on:
  - Link speed
  - Number of hops
- Avoids looping of frames.
- All frames follow the same path.
- Ensures in-order delivery in a stable SAN.



Single Path
In-Order Delivery

© 2012 Cisco and/or its affiliates. All rights reserved.

The Fabric Shortest Path First (FSPF) routing protocol is used on Fibre Channel SAN fabrics.

The figure shows that FSPF selects a single path for a given I/O transaction, avoiding looping and ensuring in-order delivery.

The FSPF algorithm is a cost-based routing algorithm that computes the most efficient path between two connected nodes. The cost of a given path is based on two factors:

- The speed of each ISL along the path
- The number of hops on the path

Routing using a single fixed path prevents looping of frames and, in a stable SAN, ensures in-order delivery. In other words, if routes are stable, frames always follow the same path. However, if the least-cost route changes while a session is in progress, then frames that are sent after the route change might take the new route.

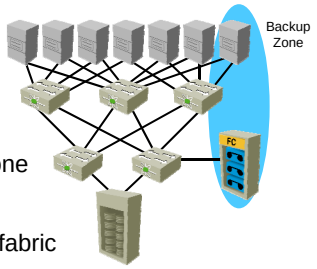
https://t.me/learningnets

Fibre Channel Zoning and LUN Masking

VSANs are used to segment the physical fabric into multiple logical fabrics. Zoning provides security within a single fabric, whether physical or logical, to restrict access between initiators and targets. LUN masking is used to provide additional security to LUNs after an initiator has reached the target device.

VSANs and Zoning

- Fabric zoning provides a means of restricting visibility and connectivity among devices that share the same SAN fabric.
- The primary goal is to prevent certain devices from accessing other fabric devices.
- Benefits of zoning:
  - Provides basic device security
  - Permits overlapping zones for shared devices
  - RSCN suppression limits messages to within a zone
- Limitations of zoning:
  - **Scalability:** Limited devices per zone, zones per fabric
  - **Fabric availability:** Shared Fibre Channel services within fabric
  - **No traffic management:** Single FSPF routing table
  - **Manageability:** Common administration for all devices within fabric
  - **Zone management security:** Multiple administrators of distributed zone set
  - **Accounting:** Difficult to account for bandwidth usage per device



© 2012 Cisco and/or its affiliates. All rights reserved.

The zoning service within a Fibre Channel fabric is designed to provide security between devices that share the same fabric. The primary goal is to prevent certain devices from accessing other devices within the fabric. With many different types of servers and storage devices on the network, the need for security is crucial. For example, if a host were to access a disk that another host is using, potentially with a different operating system, then the data on the disk could become corrupted. To avoid any compromise of crucial data within the SAN, zoning allows the user to overlay a security map. This process dictates which devices, namely hosts, can see which targets, reducing the risk of data loss.

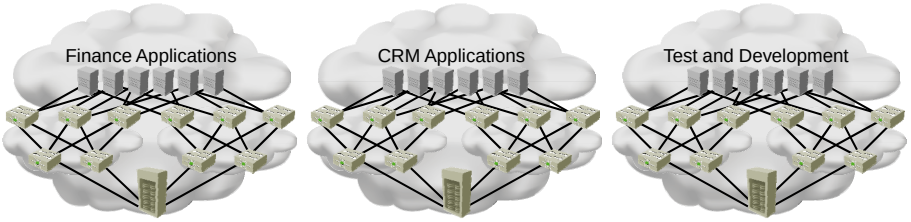
Zoning does have its limitations. Zoning was designed to do nothing more than prevent devices from communicating with other unauthorized devices. It is a distributed service that is common throughout the fabric. Any installed changes to a zoning configuration are therefore disruptive to the entire connected fabric. Zoning was also not designed to address availability or scalability of a Fibre Channel infrastructure.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

VSANs and Zoning (Cont.)

- Enterprise applications are often organized in a siloed architecture.
  - Physically isolated redundant SANs (SAN islands)
  - Provides security and isolation between applications
  - Easier to scale and manage
- Limitations
  - Difficult to share data and resources
  - Poor utilization of resources
  - Too many devices taking up space, power, and cooling



© 2012 Cisco and/or its affiliates. All rights reserved.

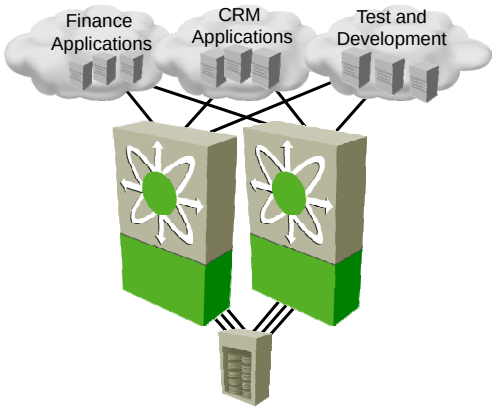
A SAN island refers to a completely physically isolated switch or group of switches that are used to connect hosts to storage devices. The reasons for building SAN islands include to isolate different applications into their own fabric, or to raise availability by minimizing the impact of fabric-wide disruptive events.

Physically separate SAN islands offer a higher degree of security because each physical infrastructure contains a separate set of Cisco Fabric Services and management access. Unfortunately, in practice this situation can become costly and wasteful in terms of fabric ports and resources.

https://t.me/learningnets

VSANs and Zoning (Cont.)

- VSANs allow SAN consolidation while maintaining logical isolation between application silos and departments:
  - Maintain departmental security and control
  - Reduce space, power, and cooling
  - Increase disk and tape utilization
  - Improve port utilization
- Each VSAN provides its own fabric services:
  - Name server
  - Zone server
  - Domain controller
  - Alias server
  - Login server
  - FSPF routing
  - Management



© 2012 Cisco and/or its affiliates. All rights reserved.

VSANs increase the efficiency of a SAN fabric by alleviating the need to build multiple physically isolated fabrics to meet organizational or application needs. Instead, fewer less-costly redundant fabrics can be built, each housing multiple applications and still providing island-like isolation.

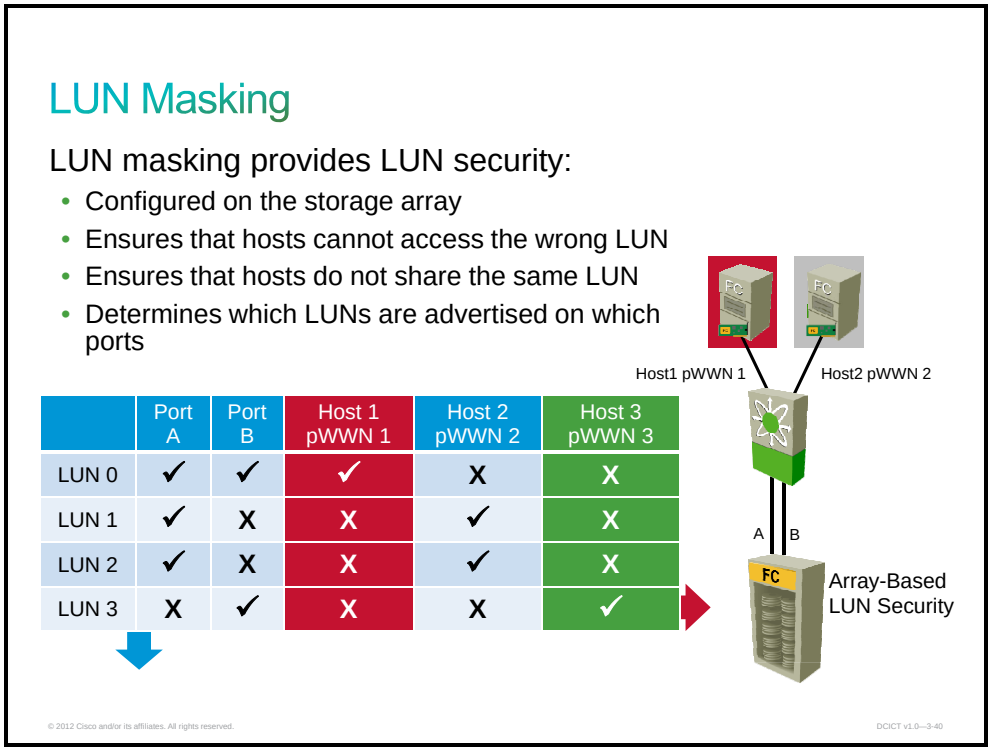
Spare ports within the fabric can be quickly and nondisruptively assigned to existing VSANs, providing a clean method of virtually growing application-specific SAN islands.

VSANs provide not only hardware-based isolation, but also a complete replicated set of Fibre Channel services for each VSAN. Therefore, when a VSAN is created, a completely separate set of Cisco Fabric Services, configuration management capability, and policies are created within the new VSAN.

Each separate virtual fabric is isolated by using a hardware-based frame tagging mechanism on VSAN member ports and EISL links. The EISL link type has been created and includes added tagging information for each frame within the fabric. The EISL link is supported between Cisco MDS and Nexus switch products. Membership to a VSAN is based on physical port, and no physical port may belong to more than one VSAN.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



LUN masking is the most common method of ensuring LUN security. Each SCSI host uses a number of primary SCSI commands to identify its target, discover LUNs, and obtain their size. These are the commands:

- **Identify:** Which device are you?
- **Report LUNs:** How many LUNs are behind this storage array port?
- **Report capacity:** What is the capacity of each LUN?
- **Request sense:** Is a LUN online and available?

It is important to ensure that each LUN on the storage array is accessed by only one host at a time, unless the hosts are configured in a cluster. As the host mounts each LUN volume, it writes a signature at the start of the LUN to claim exclusive access. If a second host should discover and try to mount the same LUN volume, it overwrites the previous signature.

LUN masking is one method that is used to ensure that only one host can access a LUN; all other hosts are masked out.

LUN masking is essentially a mapping table inside the front-end array controllers. LUN masking determines which LUNs are to be advertised through which storage array ports, and which host is allowed to own which LUNs.

https://t.me/learningnets

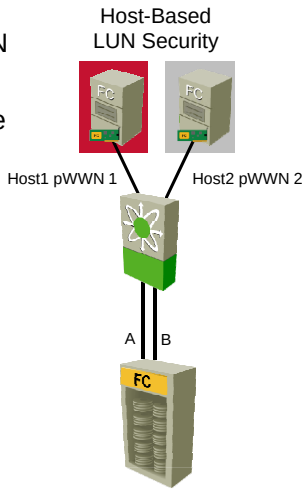
LUN Mapping

LUN mapping also provides LUN security:

- Configured on each host HBA
- Ensures that hosts do not share the same LUN
- Each host selects the LUNs to which it maps
- Used when LUN masking is unavailable on the array
- HBA persistent binding saves mapping across reboots

	Host 1 pWWN 1
LUN 0	✓
LUN 1	X
LUN 2	X
LUN 3	X

	Host 2 pWWN 2
LUN 0	X
LUN 1	✓
LUN 2	✓
LUN 3	X



© 2012 Cisco and/or its affiliates. All rights reserved.

An alternative method of LUN security is LUN mapping. If LUN masking is unavailable in the storage array, then LUN mapping can be used, although both methods can be used concurrently.

LUN mapping is configured in the HBA of each host, to ensure that only one host at a time can access each LUN.

LUNs can be advertised on many storage ports and discovered by several hosts at the same time. Many LUNs are to be visible, but it is the responsibility of the administrator to configure each HBA so that each host has exclusive access to its LUNs. When there are many hosts, mistake are more likely to be made and more than one host might access the same LUN by accident.

https://t.me/learningnets

Summary

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

Summary

- Initiator-to-target connectivity can be achieved by using one of the three Fibre Channel SAN topologies: DAS, FC-AL, or a SAN using Fibre Channel switches.
- Fibre Channel uses different port types to identify the type of devices that are connecting to the SAN.
- Fibre Channel addressing makes use of WWNs and FCIDs.
- The Fibre Channel protocol stack makes use of a layered model comprising five layers, from FC-0 through FC-4.
- For an initiator to communicate with its target device, both devices must log in and register in the fabric by using the FLOGI and PLOGI processes.
- To ensure that the correct devices have access to their storage resource, features such as zoning and LUN masking or mapping are performed in the SAN and on the edge devices.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 3

Verifying Fibre Channel Communications on Cisco MDS 9000 Series Multilayer Switches

Overview

This lesson describes how to configure Cisco MDS 9000 Series Multilayer Switches with the setup script. It also describes how to verify proper zoning, virtual storage area networks (VSANs), and the fabric login for initiators and targets.

Objectives

Upon completing this lesson, you will be able to perform an initial configuration of a Cisco MDS 9000 Series Multilayer switch and use various **show** commands. This ability includes being able to meet these objectives:

- Configure a Cisco MDS 9000 Series Multilayer switch from the CLI setup script
- Update the Cisco NX-OS on a Cisco MDS 9000 Series Multilayer switch
- Update licensed features on a Cisco MDS 9000 Series Multilayer switch
- Verify initiator and target fabric login
- Verify Fibre Channel zoning on a Cisco MDS 9000 Series Multilayer switch
- Collect technical support data on a Cisco MDS 9000 Series Multilayer switch

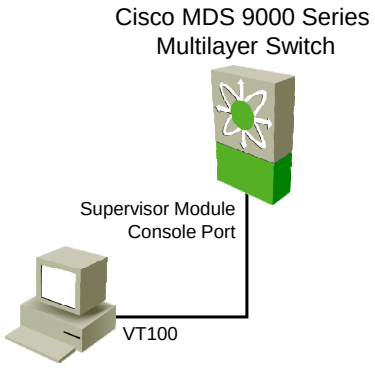
https://t.me/learningnets

Configure a Cisco MDS 9000 Series Multilayer Switch from the CLI Setup Script

Explain how to configure the initial settings of a Cisco MDS 9000 Series Multilayer switch and Secure Shell (SSH) in to the management CLI.

Terminal Setup and Configuration

- Verify that switch is powered off.
- Cable the VT100 to the switch console port.
- Terminal setup:
  - 9600 b/s
  - 8 data bits
  - No parity
  - 1 stop bit
  - No flow control
- Power on switch.
- Switch boots automatically.
- Switch setup script appears.

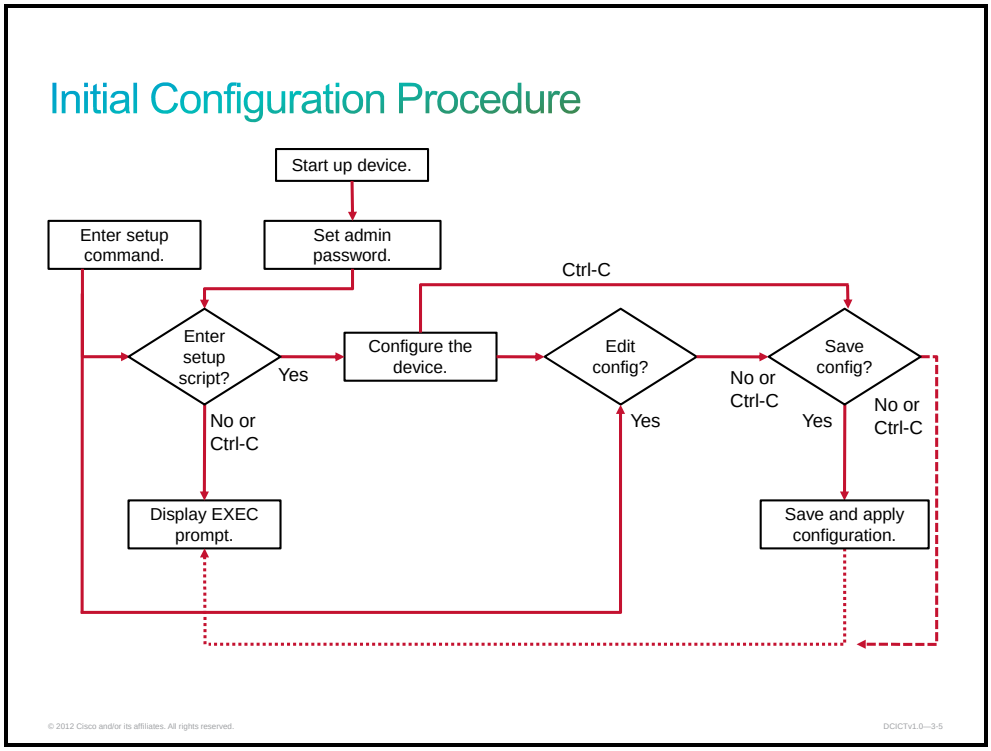


© 2012 Cisco and/or its affiliates. All rights reserved.

The console connection of a Cisco MDS 9000 Series Multilayer switch uses a rollover RJ-45 cable. To configure the switch by using the setup script at the CLI, use the rollover RJ-45 cable to connect your computer to the console connection on the switch. Power up the switch, and using the terminal application on your computer, wait for the switch to boot up and display the initial setup script.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



The Cisco Nexus Operating System (NX-OS) setup utility is an interactive CLI mode that guides you through a basic system configuration, which is called a startup. The setup utility allows you to configure enough connectivity for system management and to build an initial configuration by using the System Configuration dialog.

The setup utility is used mainly for initially configuring the system when no configuration exists. However, the utility can be used at any time for basic device configuration. Any configured values are kept when you skip steps in the script. For example, if there is already a configured mgmt0 interface address, the setup utility does not change that value if you skip that step. However, if there is a default value for the step, skipping a step causes the setup utility to change the configuration to that default rather than the configured value.

Note	Ensure that you configure the IP version 4 (IPv4) route, the default network IPv4 address, and the default gateway IPv4 address to enable remote access.

https://t.me/learningnets

Example: Initial Configuration Procedure

```
         ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: <enter>
 Enter the password for "admin": 1234QWer
  Confirm the password for "admin": 1234QWer

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

Please register Cisco MDS 9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. MDS devices must be registered to receive entitled
support services.

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

Would you like to enter the basic configuration dialog (yes/no): yes
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The setup utility has these prerequisites:

- Have a password strategy for the network environment.
- Connect the console port on the switch to the computer or a terminal server.
- Connect the Ethernet management port on the switch to the network.
- Enable the licensing grace period, if required.

When the initial configuration has been created, the switch becomes accessible via remote management terminal and GUI connections.

Note Secure passwords are enabled by default for the Cisco NX-OS. A secure password should contain characters from at least four classes: lowercase letters, uppercase letters, digits, and special characters. Passwords are case sensitive.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Example: Initial Configuration Procedure (Cont.)

```
 Create another login account (yes/no) [n]: <enter>
  Configure read-only SNMP community string (yes/no) [n]: <enter>
  Configure read-write SNMP community string (yes/no) [n]: <enter>
  Enter the switch name : MDS-A
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
<enter>
    Mgmt0 IPv4 address : 172.16.1.51
    Mgmt0 IPv4 netmask : 255.255.255.0
  Configure the default gateway? (yes/no) [y]: <enter>
    IPv4 address of the default gateway : 172.16.1.1
  Configure advanced IP options? (yes/no) [n]: <enter>
  Enable the ssh service? (yes/no) [y]: <enter>
    Type of ssh key you would like to generate (dsa/rsa) [rsa]: <enter>
    Number of rsa key bits <1024-2048> [1024]: <enter>
  Enable the telnet service? (yes/no) [n]: <enter>
  Enable the http-server? (yes/no) [y]: n
```

© 2012 Cisco and/or its affiliates. All rights reserved.

After the admin password has been created and you have entered the initial configuration dialog, you are requested to enter the name of your switch along with the IP address and subnet mask of the out-of-band (OOB) Ethernet management port (mgmt 0) interface. Without this information, management access to the switch through the OOB Ethernet port would not be possible.

When there are options to select with each dialog, you can either press Return, which will accept the choice indicated between the square brackets, or you can select an alternative.

A Network Time Protocol (NTP) server provides a precise time source to synchronize the system clocks of network devices. NTP is transported over UDP/IP. All NTP communications use Coordinated Universal Time (UTC). An NTP server receives its time from a reference time source such as a radio clock or atomic clock. NTP distributes this time across the network. Using NTP is optional but recommended.

The SSH service is enabled by default, to enable remote access to the switch. The Telnet service is disabled by default. SSH provides secure management access to the switch and is preferred over Telnet, which sends usernames and passwords across the network in a cleartext format.

https://t.me/learningnets

Example: Initial Configuration Procedure (Cont.)

```
Configure clock? (yes/no) [n]: <enter>

Configure timezone? (yes/no) [n]: <enter>

Configure summertime? (yes/no) [n]: <enter>

 Configure the ntp server? (yes/no) [n]: <enter>

 Configure default switchport interface state (shut/noshut) [shut]: <enter>

 Configure default switchport trunk mode (on/off/auto) [on]: <enter>

 Configure default switchport port mode F (yes/no) [n]: <enter>

 Configure default zone policy (permit/deny) [deny]: <enter>

 Enable full zoneset distribution? (yes/no) [n]: <enter>

 Configure default zone mode (basic/enhanced) [basic]: <enter>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The final section of the initial configuration procedure sets the default states of the interfaces, the default modes for the interfaces, and the default zone policy.

Example: Initial Configuration Procedure (Cont.)

The system prints a summary of the configuration for your review.

```
The following configuration will be applied:
  password strength-check
  switchname MDS-A
  interface mgmt0
    ip address 172.16.1.51 255.255.255.0
    no shutdown
  ip default-gateway 172.16.1.1
  ssh key rsa 1024 force
  feature ssh
  no feature telnet
  no feature http-server
  system default switchport shutdown
  system default switchport trunk mode on
  no system default zone default-zone permit
  no system default zone distribute full
  no system default zone mode enhanced

Would you like to edit the configuration? (yes/no) [n]: <enter>

Use this configuration and save it? (yes/no) [y]: <enter>

[########################################] 100%
Copy complete, now saving to disk (please wait)...
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The system prints a summary of the initial configuration for your review. The printed configuration is exactly what you entered. Compare it with the information that you obtained in the initial setup requirements to verify that there are no typing errors.

https://t.me/learningnets

The system asks whether you want to edit the configuration that just printed out. Any configuration changes that you make to a switch are immediately enforced but are not saved. If you have verified that everything was entered correctly, you do not need to edit the configuration. If no edits are needed, then are asked if you want to use and save this configuration. Because [y] (for yes) is the default selection, pressing Return activates this command, and the configuration becomes part of the running configuration and is copied to the startup configuration.

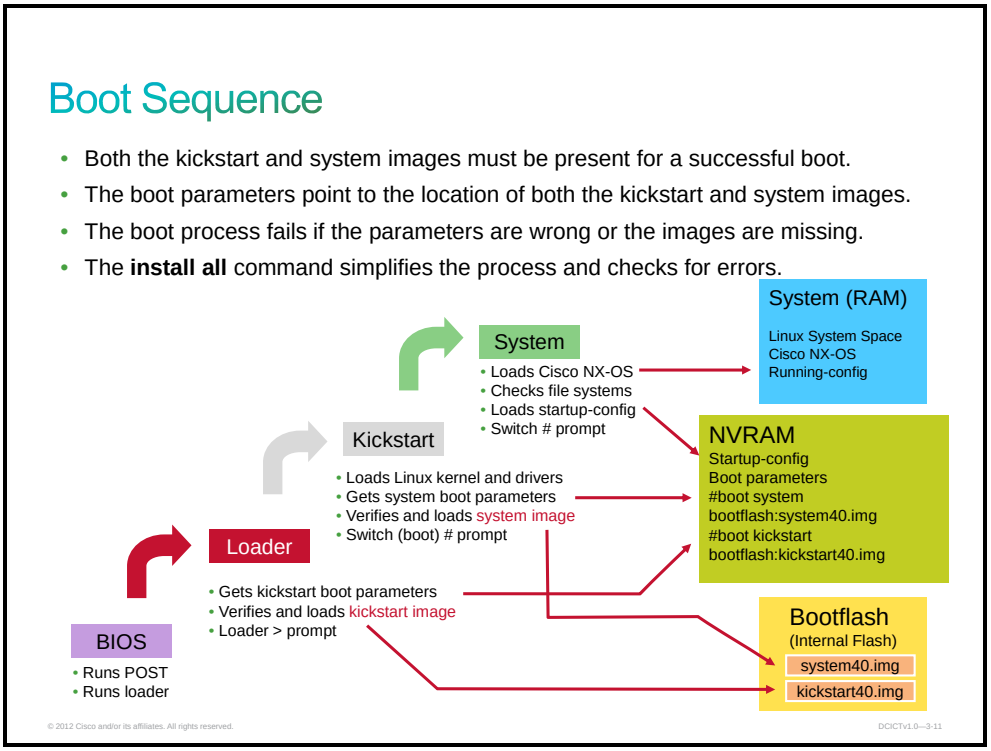
This process also ensures that the kickstart and system boot image locations are automatically configured. If you do not save the configuration at this point, none of your changes are updated on the switch.

Now that the switch has a management IP address, you can connect to the switch remotely via SSH by opening SSH on your computer and entering the IP address of the switch. On connecting to the switch, you are asked to enter the username and password, after which you gain access to the CLI of the switch.

https://t.me/learningnets

Update the Cisco NX-OS on a Cisco MDS 9000 Series Multilayer Switch

The Cisco NX-OS software can be upgraded as required on the Cisco MDS 9000 Series Multilayer switch by using the **install all** command. This topic explains the procedure for upgrading the switch.



When the Cisco MDS 9000 Series Multilayer switch is first switched on or during reboot, the system BIOS on the supervisor module first runs power-on self-test (POST) diagnostics. The BIOS then runs the loader bootstrap function.

The boot parameters are held in NVRAM and point to the location and name of both the kickstart and system images. The loader obtains the location of the kickstart file, usually on bootflash, and verifies the kickstart image before loading it.

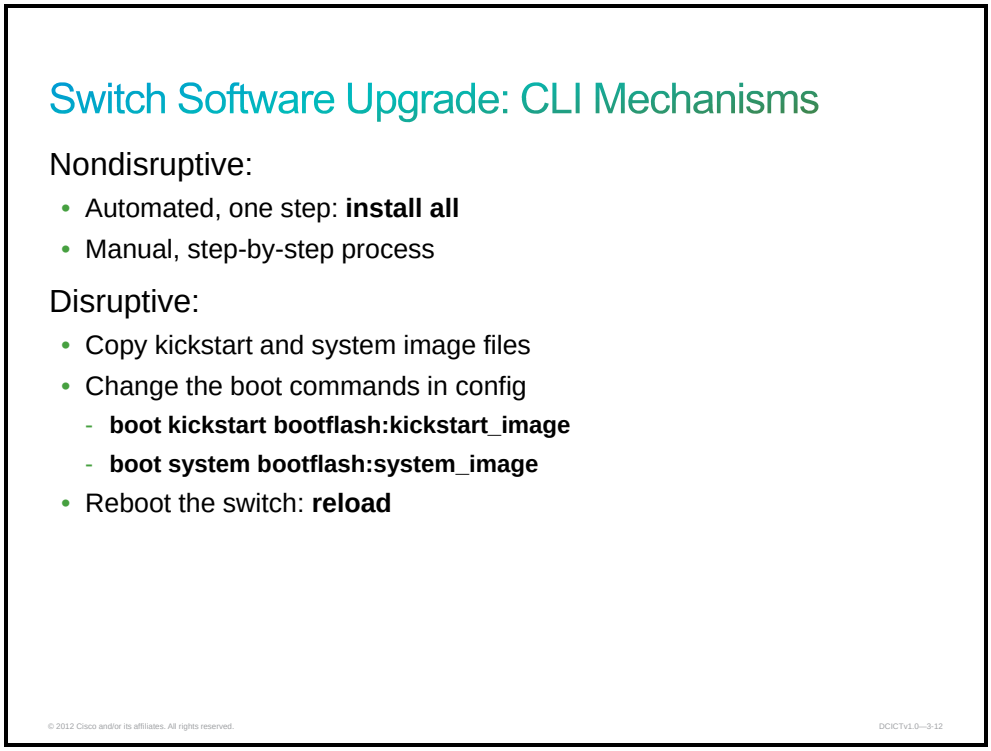
The kickstart loads the Linux kernel and device drivers and then needs to load the system image. Again, the boot parameters in NVRAM should point to the location and name of the system image, usually on bootflash. The kickstart then verifies and loads the system image.

Finally, the system image loads the Cisco NX-OS Software, checks the file systems, and proceeds to load the startup configuration that contains the switch configuration from NVRAM.

If the boot parameters are missing or have an incorrect name or location, the boot process fails at the last stage. If this error happens, the administrator must recover from the error and reload the switch. The **install all** command launches a script that greatly simplifies the boot procedure and checks for errors and the upgrade impact before proceeding.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Cisco NX-OS Software, which is designed for mission-critical high-availability environments, allows the upgrade of software without any disruptions. You can upgrade any Cisco MDS switch by using one of three mechanisms:

- An automated, one-step upgrade that uses the **install all** command
- A manual, step-by-step upgrade
- A quick one-step upgrade that uses the **reload** command

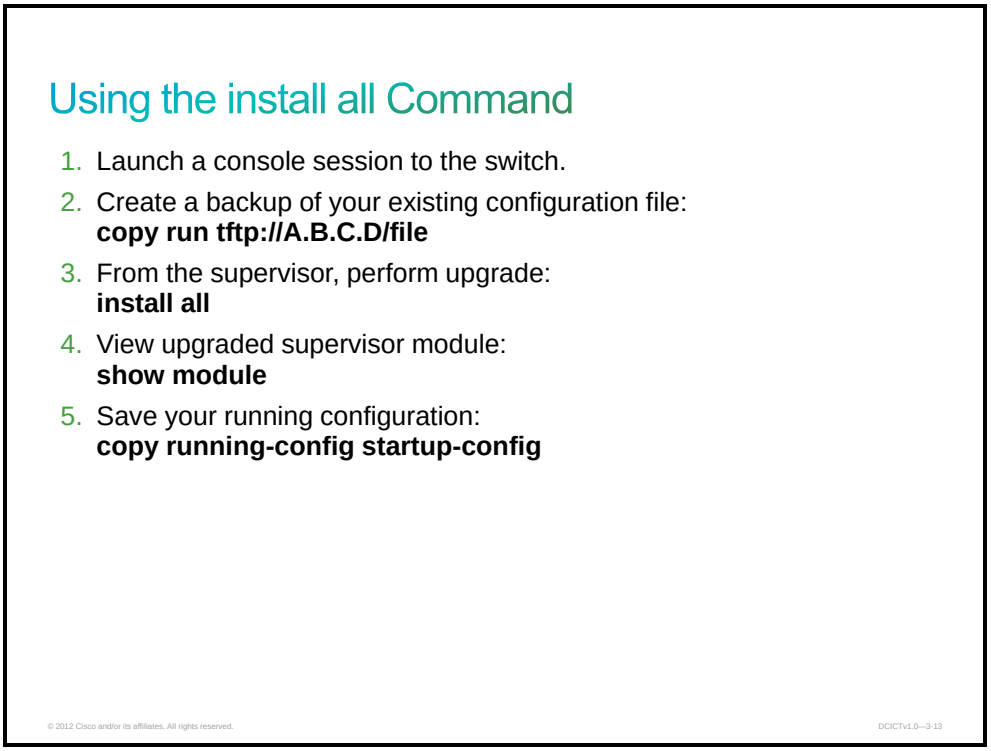
Before running the **reload** command, copy the correct kickstart and system images to the correct location and change the **boot** commands in your configuration. The quick upgrade is disruptive. For nondisruptive upgrades, use the automated one-step upgrade or the manual step-by-step upgrade.

In some cases, any software upgrade might be disruptive. These situations can occur under these conditions:

- A single supervisor system has incompatible system software images
- A dual supervisor system has incompatible system software images

Note	The **install all** command compares and presents the results of the compatibility before proceeding with the installation. You can exit if you do not want to proceed with these changes. To determine version compatibility between switch images, use the **show install all impact** command. This command displays the impact of using the **install all** command.

https://t.me/learningnets

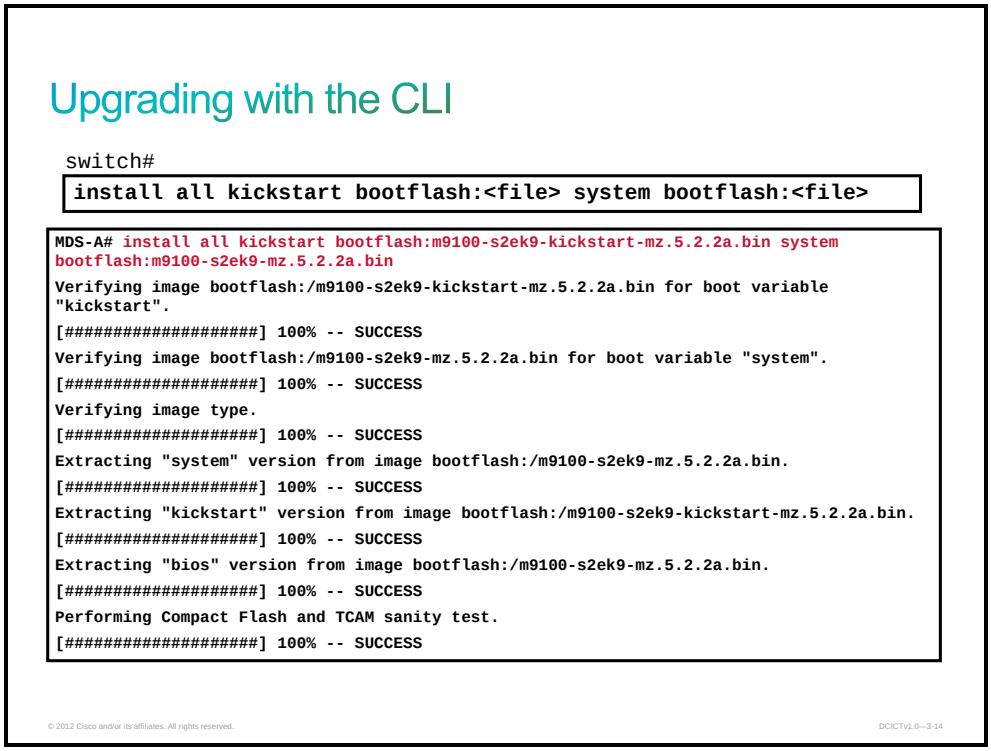


The one-step **install all** command is recommended to upgrade your system software. This command upgrades the Cisco MDS 9000 Series Multilayer switch. Only one **install all** command can run on a switch at any time, and no other command can be issued while running that command.

To save the configuration after an upgrade, enter the **copy running-config startup-config** command in EXEC mode to save the configuration into nonvolatile storage. After you issue this command, the running and startup copies of the configuration are identical. This process is necessary because the boot variables are not automatically instantiated to the startup config file after an upgrade (whether manual, install all, or GUI). Without issuing this command immediately after a successful upgrade, you run the risk of the switch rebooting to a previous release.

Note	You should issue the **install all** command from the console terminal of the switch. If you issue the **install all** command from a Telnet or SSH session, you cannot view the full set of messages because the connection is lost when the control plane reloads.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Images can be retrieved in one of two ways:

- **Locally:** Images are locally available on the switch. The **install all** command uses the specified local images.
- **Remotely:** Images are in a remote location. The user specifies the destination by using the remote server parameters and the filename to be used locally.

To upgrade the switch to a new image, you must specify the variables that direct the switch to the images. To select the kickstart image, use the kickstart variable. To select the system image, use the system variable. The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch.

https://t.me/learningnets

Upgrading with the CLI (Cont.)

Compatibility check and image upgrade table

```
Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive          none


Images will be upgraded according to following table:
Module       Image                 Running-Version(pri:alt)           New-Version  Upg-Required
------  ----------  ------------------------------------------  -------------------  ------------
     1      system                                     5.2(2a)               5.2(2a)            no
     1   kickstart                                     5.2(2a)               5.2(2a)            no
     1        bios     v1.0.19(02/01/10):  v1.0.19(02/01/10)     v1.0.19(02/01/10)            no
```

© 2012 Cisco and/or its affiliates. All rights reserved.

When you issue the **install all** command, the switch displays a summary of changes that will be made to your configuration. The switch then waits for your authorization to continue the command process.

A compatibility check is conducted and the impact of an upgrade and the install type are displayed.

Modules and specific images to be upgraded, based on the files that are specified in the previous step, are displayed in an upgrade table that also shows the running and new versions.

These are the compatibility check terms:

- **Bootable:** The ability of the module to boot or not boot, based on image compatibility
- **Impact:** The type of software upgrade mechanism—disruptive or nondisruptive

These are the install type terms:

- **Reset:** Resets the module
- **sw-reset:** Resets the module immediately after switchover
- **Rolling:** Upgrades each module in sequence
- **Copy-only:** Updates the software for BIOS, loader, or BootROM

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Upgrading with the CLI (Cont.)

Continuing the installation

```
Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Performing runtime checks.
[####################] 100% -- SUCCESS

Setting boot variables.
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[####################] 100% -- SUCCESS

Install has been successful.
MDS-A#
```

© 2012 Cisco and/or its affiliates. All rights reserved.

If there are no compatibility issues and you want to continue installing the identified images, enter **y** (yes) when prompted.

After you confirm to continue the installation, a display of the installation progress appears. If the install succeeds, you receive a notification that the install has been successful and the switch prompt is displayed.

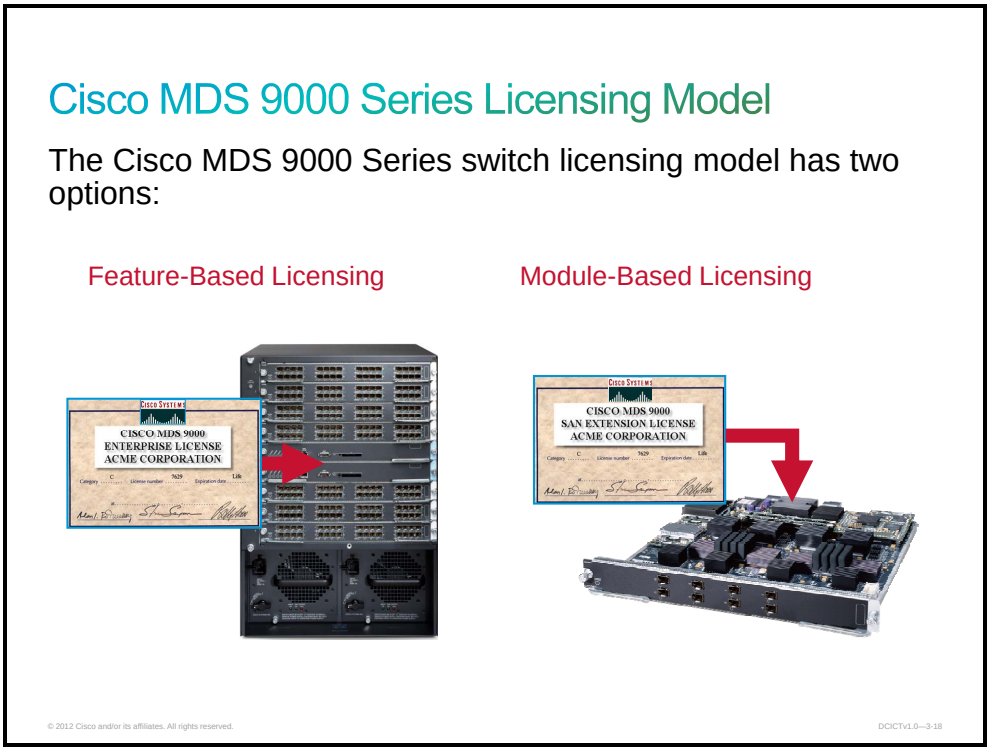
These are some benefits of using the **install all** command to upgrade system software:

- The streamlined process ensures using the best possible procedures to upgrade the software in the least disruptive manner.
- All the software is upgraded in one easy step.
- The process verifies all the images before installation and detects incompatibilities.
- The process checks configuration compatibility.
- Information is provided about the impact of the upgrade before it takes place.

https://t.me/learningnets

Update Licensed Features on the Cisco MDS 9000 Series Multilayer Switch

When purchasing a Cisco MDS 9000 Series Multilayer switch, you have the option of purchasing the base license or additional licenses for the switch. If you purchase a base license model, you might want to upgrade your license for additional features and ports. This topic explains how to upgrade the licenses on the Cisco MDS 9000 Series Multilayer switch.

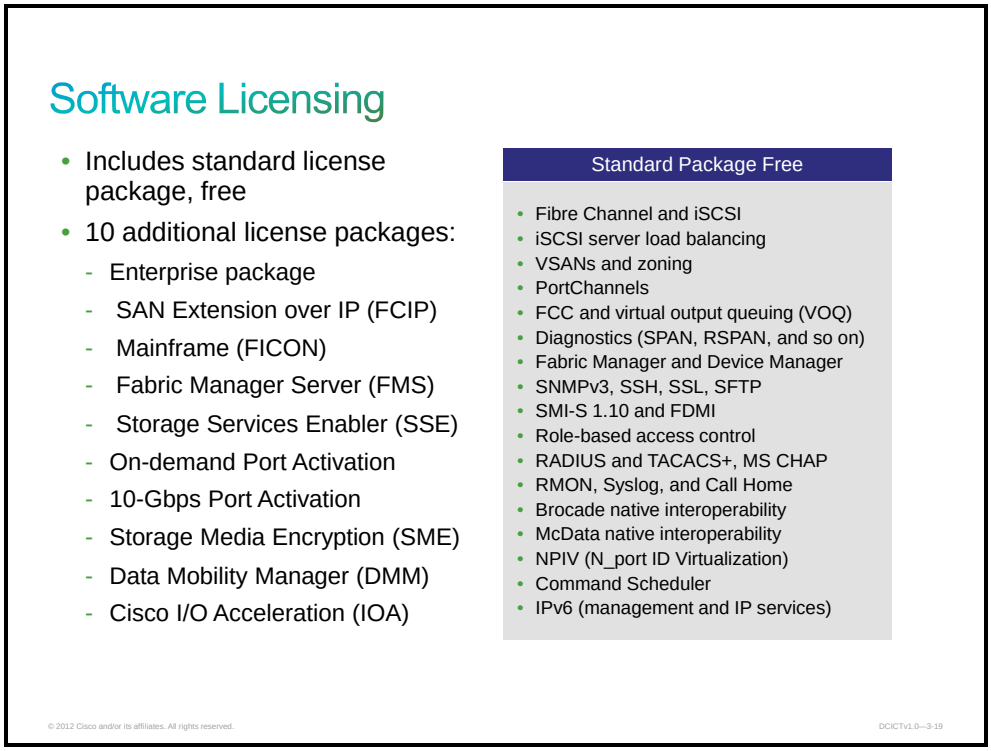


Licensing allows access to specified premium features on the switch after you install the appropriate licenses. Licenses have been sold, supported, and enforced since the release of Cisco MDS NX-OS Release 1.3(1).

The licensing model that is defined for the Cisco MDS 9000 Series Multilayer switch product line has two options:

- Feature-based licensing covers features that are applicable to the entire switch. The cost varies based on a per-switch usage.
- Module-based licensing covers features that require additional hardware modules. The Cisco MDS 9000 Series Multilayer switch does not use module-based licensing.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Cisco NX-OS Software is the underlying system software that powers Cisco MDS 9000 Series Multilayer Switches. In the best traditions of Cisco IOS Software, Cisco NX-OS Software is designed to create a strategic SAN platform of superior reliability, performance, scalability, and features.

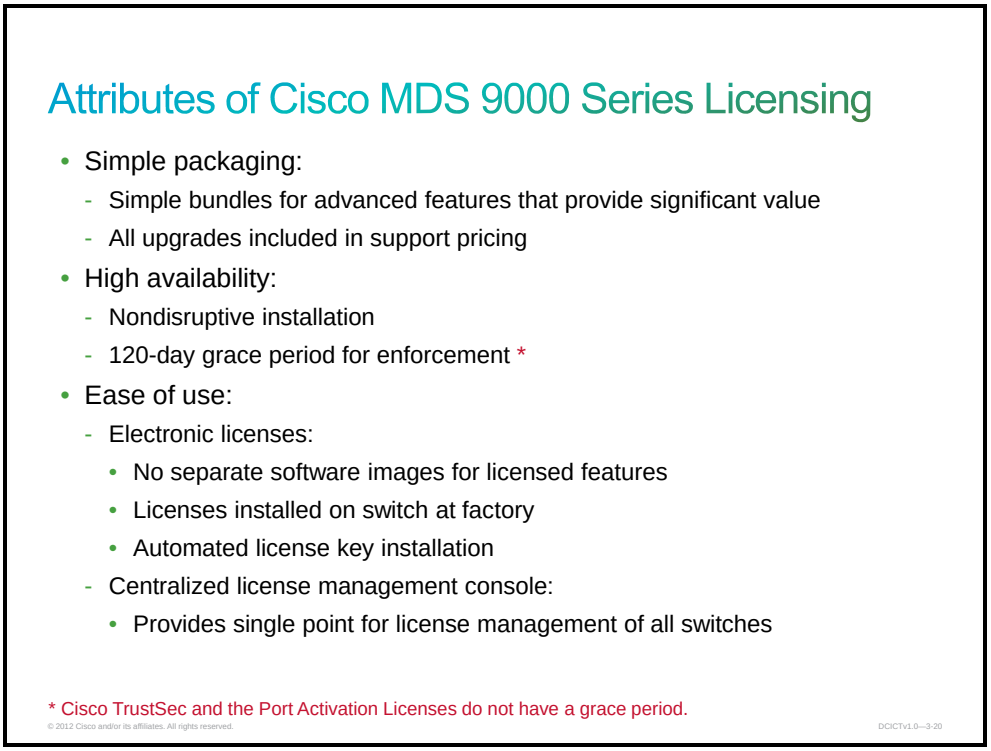
In addition to providing all the features that the market expects of a storage-network switch, Cisco NX-OS Software provides many unique features that help Cisco MDS 9000 Series Multilayer Switches deliver low total cost of ownership (TCO) and a quick return on investment (ROI).

Common Software Across All Platforms

Cisco NX-OS Software runs on all Cisco MDS 9000 Series Multilayer Switches, from multilayer fabric switches to multilayer directors. Using the same base system software across the entire product line enables Cisco to provide an extensive, consistent, and compatible feature set on the Cisco MDS 9000 Series Multilayer switch.

Most Cisco MDS 9000 Series software features are included in the base switch configuration. The standard software package includes the base set of features that Cisco believes are required by most customers for building a SAN. However, some features are logically grouped into add-on packages that must be licensed separately.

https://t.me/learningnets



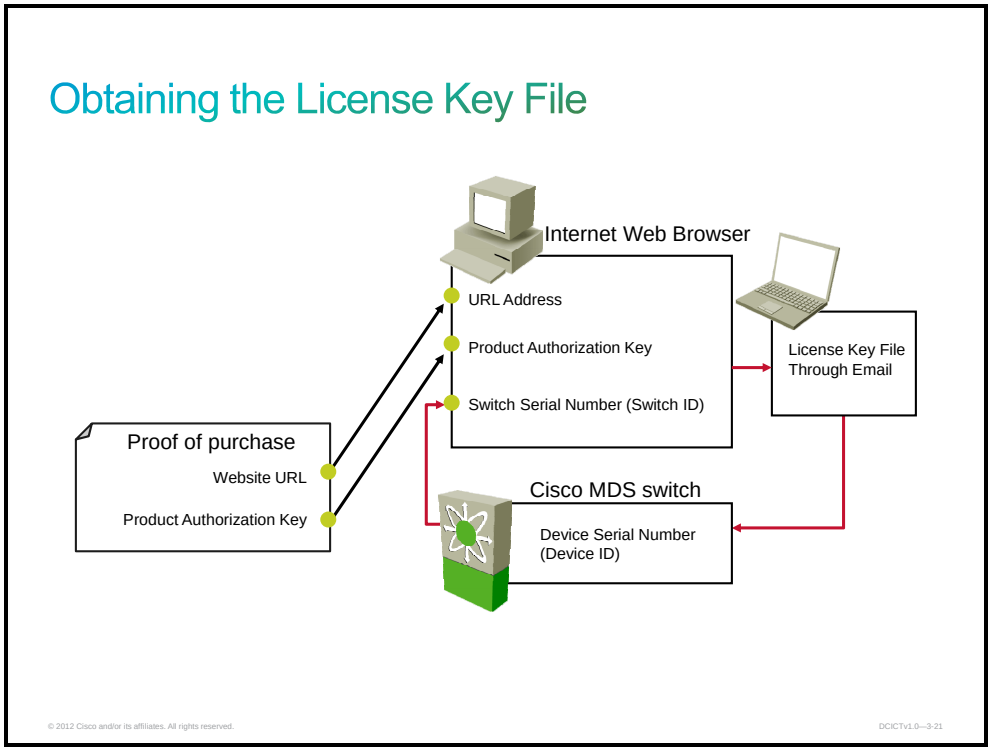
License usability can be challenging with existing products. Customers have concerns about compromising availability with disruptive software installations for licensed features. License management is a notorious problem.

Cisco license packages require a simple installation of an electronic license. No software installation or upgrade is required. Licenses can also be installed on the switch in the factory. Cisco MDS 9000 Series Multilayer Switches store license keys on the chassis serial PROM (SPROM), so license keys are never lost, even during a switch-software reinstall.

Cisco Fabric Manager includes a centralized license-management console that provides a single interface for managing licenses across all Cisco MDS 9000 Series Multilayer Switches in the fabric. This console reduces management overhead and prevents problems caused by improperly maintained licensing. If an administrative error occurs with licensing, the switch provides a grace period before the unlicensed features are disabled. This delay allows plenty of time to correct the licensing issue.

All licensed features may be evaluated for 120 days before a license is required.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Manual Installation

If you have existing switches, or if you want to install the licenses on your own, you must first obtain the license key file and install that file in the switch.

Obtaining the License Key File

See the *Cisco DCNM Fundamentals Guide, Release 5.x*, at http://www.cisco.com/en/US/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/fund/DCNM-SAN-LAN_5_2/DCNM_Fundamentals.html, for details about using the Cisco DCNM for SAN GUI to install automated licenses. To use the CLI to obtain new or updated license key files, follow these steps.

Step 1 Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.

```
switch# show license host-id
License hostid: VDH=FOX064317SQ
```

Step 2 Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco MDS 9000 Series Multilayer switch.

Step 3 Get the Product Authorization Key (PAK) from either the claim certificate or the proof-of-purchase document.

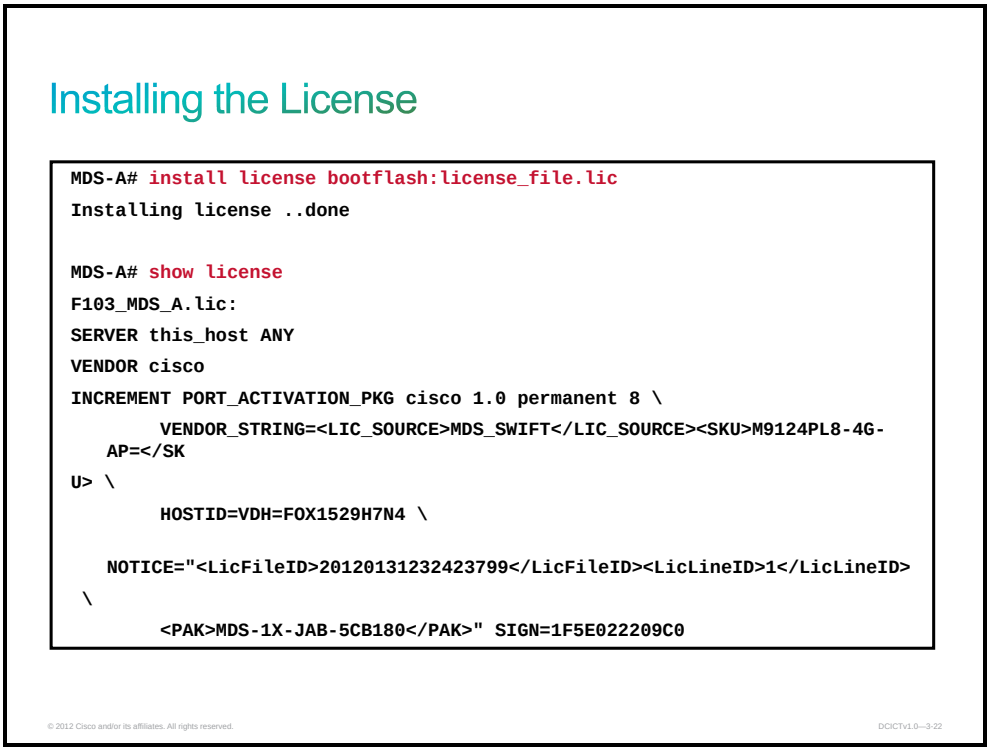
Step 4 Locate the website URL from either the claim certificate or the proof-of-purchase document.

Step 5 Access the specified URL that applies to your switch and enter the switch serial number and the PAK. The license key file is sent to you by email. The license key file is digitally signed to authorize use only on the requested switch. The requested features are enabled after the Cisco NX-OS Software on the specified switch accesses the license key file.

https://t.me/learningnets

Note Install the license key file in the specified Cisco MDS 9000 Series Multilayer switch without making any modifications. A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for use of a licensed feature begins the first time that you use the feature that the license covers.

Step 6 Use the **copy licenses** command in EXEC mode to save your license file to the bootflash directory.



Installing the License Key File

If you need to install multiple licenses in any Cisco MDS 9000 Series switch, be sure to provide unique filenames for each license key file. To install a license key file in any switch, follow these steps:

Step 1 From a console session on the switch, invoke the **install license** command:

```
switch# install license bootflash:license_file.lic
Installing license . . . done
```

Note If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the filename that is specified in the license key file is used to install the license.

Step 2 Exit the switch console and open a new terminal session to view all license files that are installed on the switch, by using the **show license** command:

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
HOSTID=VDH=FOX0646S017 \
```

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

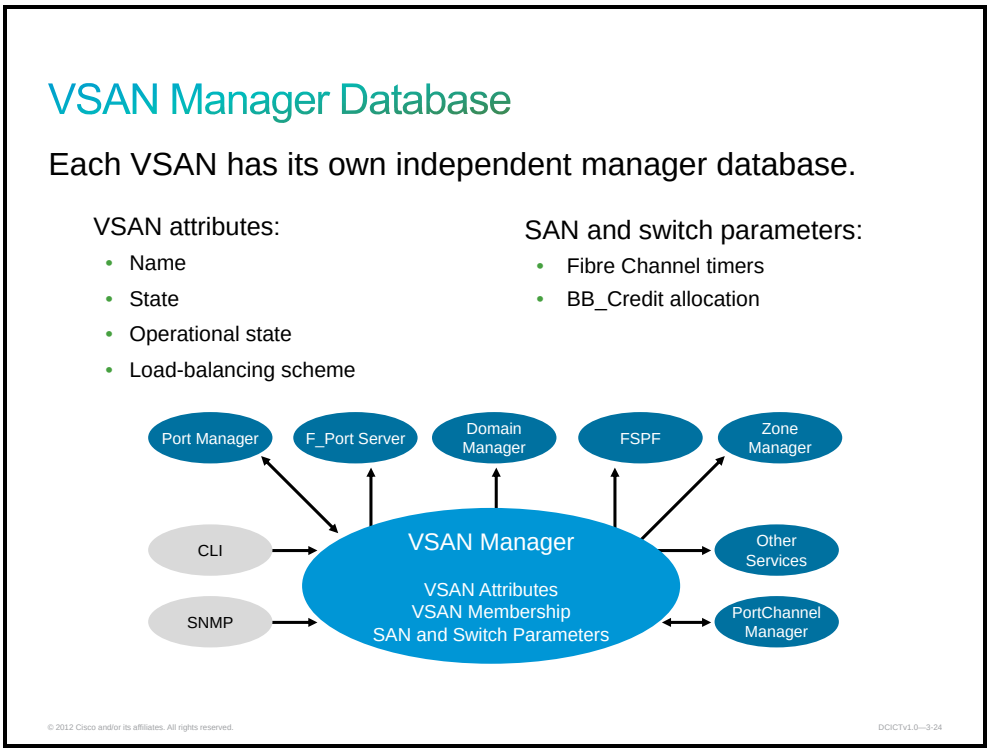
```
NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
<PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Note If the license meets all guidelines when the **install license** command is issued, all features and modules continue functioning as configured. This policy is true for any Cisco MDS 9000 Series switch.

https://t.me/learningnets

Verify Initiator and Target Fabric Login

Use **show** commands to verify that the initiators and targets that connect to the fabric have successfully logged in to the SAN. This topic describes how to use these commands.



The VSAN Manager keeps a list of all the members of each VSAN, as well as the allowed VSAN list for trunking expansion ports (TE Ports) and port channels. The VSAN Manager also maintains the VSAN shared parameters. The VSAN Manager coordinates the Fibre Channel services for each VSAN and allows configuration and other management activities through the CLI, XML, and Simple Network Management Protocol (SNMP).

Each VSAN has its own independent manager database. Furthermore, each VSAN uses an independent Cisco Fabric Services database such as Fabric Shortest Path First (FSPF), zoning, name server, and Fabric Configuration Server (FCS), thus increasing security and reliability. The VSAN Manager keeps track of additions, modifications, deletions, creations, removals, updates, and membership.

The VSAN Manager maintains a database that contains information about each VSAN:

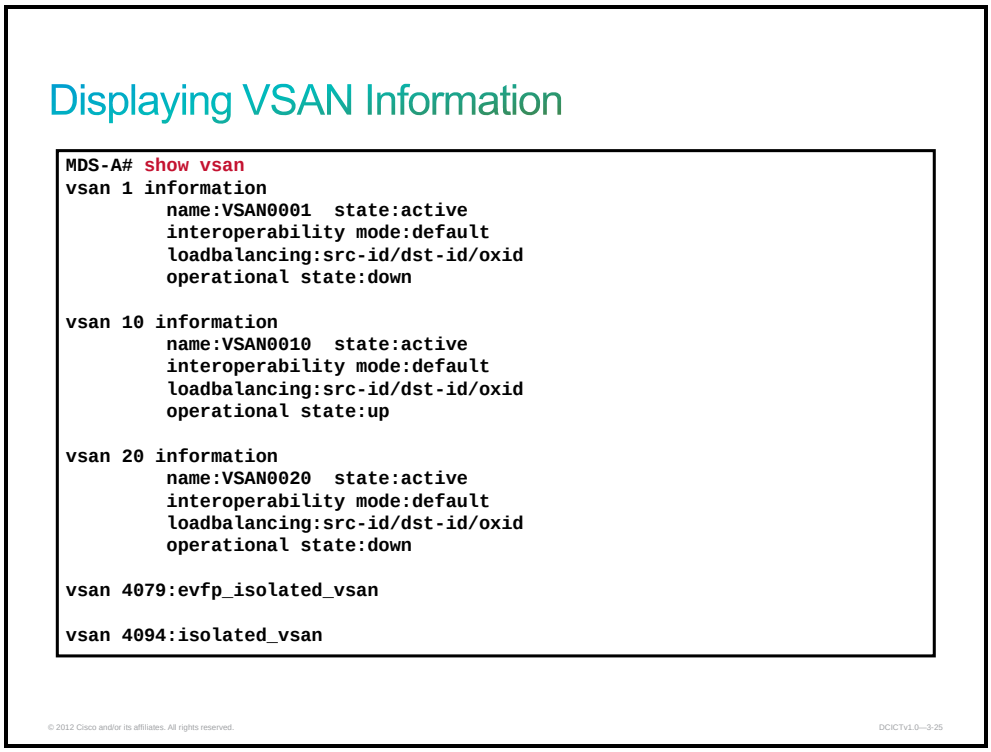
- Name of the VSAN (must be unique for each VSAN)
- State of the VSAN (suspended or active)
- Operational state of the VSAN (up or down); up interfaces in the VSAN cause the VSAN operational state to be up
- Load-balancing scheme that is used for each VSAN

These switch-specific parameters are not per VSAN:

- Fibre Channel timers, such as error detect timeout value (E_D_TOV), distributed services timeout value (D_S_TOV), and resource allocation timeout value (R_A_TOV)
- Buffer-to-buffer credit (BB_Credit) allocation

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

VSAN 1, the default VSAN, is configured on the switch by default. You can suspend this VSAN, but you cannot delete it.



To verify the current VSAN configuration and status, use the **show vsan** command in EXEC mode. Use this command to display this information:

- VSANs created
- VSAN name
- Administrative state (active or suspended)
- Interoperability setting
- Load-balancing scheme (default is S_ID/D_ID/OX_ID; S_ID/D_ID is the alternative)
- Operational state (up or down)

To report the status for a specific VSAN, add the VSAN number to the **show vsan** command; for example **switch# show vsan 20**.

https://t.me/learningnets

Displaying VSAN Information (Cont.)

Use the **show vsan membership** command to verify VSAN membership.

```
MDS-A# show vsan membership
vsan 1 interfaces:
    fc1/1           fc1/2           fc1/3           fc1/4
    fc1/5           fc1/6           fc1/7           fc1/8
    fc1/10          fc1/11          fc1/12          fc1/13
    fc1/14          fc1/15          fc1/16          fc1/17
    fc1/18          fc1/19          fc1/20          fc1/21
    fc1/22          fc1/23          fc1/24

vsan 10 interfaces:
    fc1/9

vsan 20 interfaces:


vsan 4079(evfp_isolated_vsan) interfaces:


vsan 4094(isolated_vsan) interfaces:
```

© 2012 Cisco and/or its affiliates. All rights reserved.

To verify port VSAN membership, use the **show vsan membership** command in EXEC mode. The report displays port VSAN assignments, including the isolated VSAN (4094).

To report the membership for a specific VSAN, add the VSAN number to the **show vsan membership** command:

```
switch# show vsan 10 membership
vsan 10 interfaces:
 fc1/5       fc1/6
```

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Displaying the FLOGI Database

```
MDS-A# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/9            10    0x6f0000  50:0a:09:81:98:ba:78:eb 50:0a:09:80:88:ba:78:eb

Total number of flogi = 1.
```

- A device with the pWWN 50:0a:09:81:98:ba:78:eb is attached to switch interface fc1/9.
- Interface fc1/9 is a member of VSAN 10. The runtime domain ID of VSAN 10 on this switch is 6f (111).
- The VSAN 10 domain manager has assigned the FCID 6f 00 00 to the device with pWWN 50:0a:09:81:98:ba:78:eb during FLOGI.
- The assigned FCID and device WWN information are registered in the FCNS and attain fabric-wide visibility.

© 2012 Cisco and/or its affiliates. All rights reserved.

In a Fibre Channel fabric, each host or disk requires a Fibre Channel ID (FCID). Use the **show flogi** command to verify whether a storage device is displayed in the fabric login (FLOGI) table.

If the required device is displayed in the FLOGI table, then the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host bus adapter (HBA) of the host.

https://t.me/learningnets

Displaying FCNS Entries

- Use the **show fcns** command to display the FCNS and statistical information for a specific VSAN or for all VSANs.
- Extended information is displayed with the **show fcns detail** command.

```
MDS-A# show fcns database vsan 10

VSAN 10:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x6f0000    N     50:0a:09:81:98:ba:78:eb (NetApp)        scsi-fcp:target

Total number of entries = 1
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The Fibre Channel Name Server (FCNS) stores name entries for all hosts in the FCNS database. The name server permits an Nx Port to register attributes during a FLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

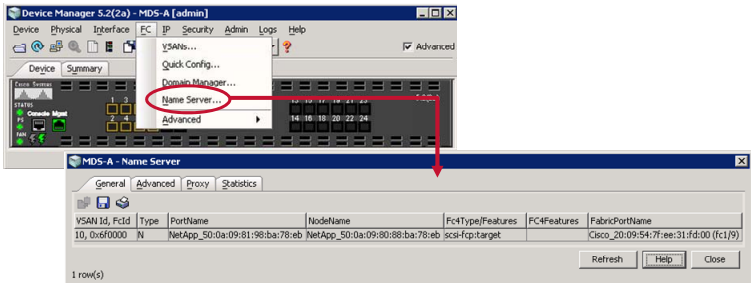
In a multiswitch fabric configuration, the name server instances that are running on each switch share information in a distributed database. One instance of the name server process runs on each switch, with separate databases for each VSAN.

Use the **show fcns** command to display the name server database and statistical information for a specific VSAN or for all VSANs.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Displaying FCNS Entries in Device Manager

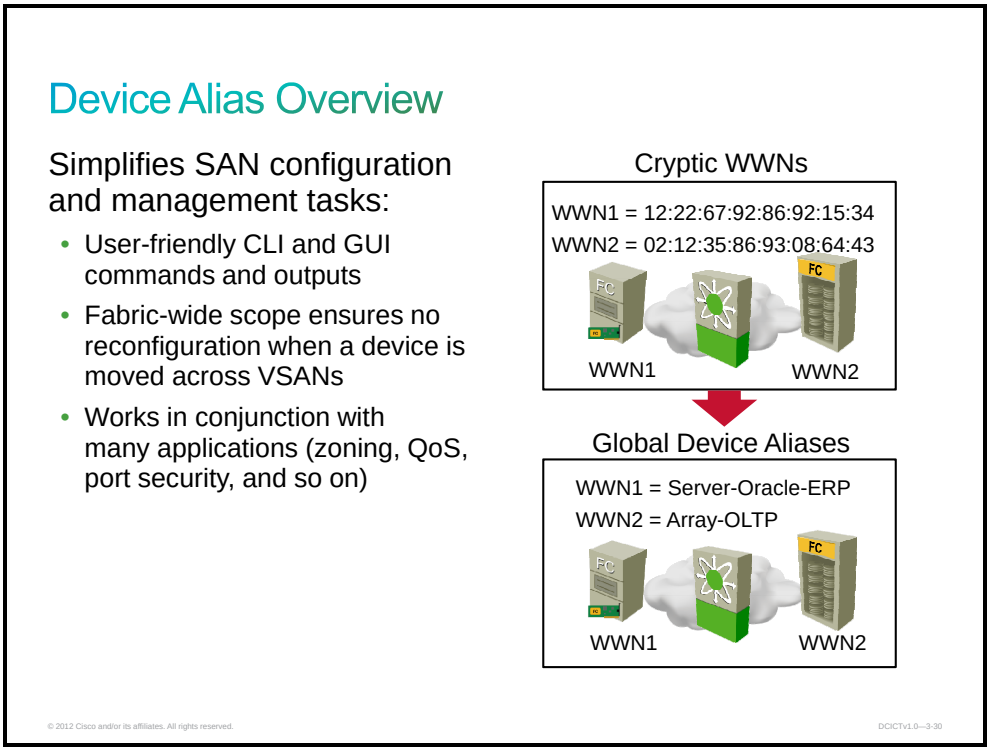
Choose the **FC > Name Server** menu to view the FCNS database.



© 2012 Cisco and/or its affiliates. All rights reserved.

The menu bar at the top of the Cisco Device Manager main window provides options for managing and troubleshooting a single switch.

https://t.me/learningnets



In some instances, the port world wide name (pWWN) of a device must be specified to configure features such as zoning, quality of service (QoS), and port security in a Cisco MDS 9000 Series Multilayer switch. You must assign the correct device name each time you configure these features, or unexpected results will occur. This problem can be avoided by using a user-friendly name for the pWWN for all configuration commands. These user-friendly names are referred to as device aliases on the Cisco MDS 9000 Series Multilayer switch.

Device aliases support two modes: basic and enhanced.

- When a device alias runs in basic mode, all applications function like the applications on the Cisco MDS NX-OS 3.0 switches. When you configure the basic mode using device aliases, the application immediately expands to pWWN. This operation continues until the mode is changed to enhanced.
- When a device alias runs in enhanced mode, all applications accept the device-alias configuration in the native format. The applications store the device alias in the configuration and distribute it in the device-alias format instead of expanding to pWWN. The applications track the device alias database changes and take actions to enforce it.

A native device-alias configuration is not accepted in the interop mode VSAN. Inter-VSAN Routing (IVR) zone set activation fails in interop mode VSANs if the corresponding twilight zones that are being injected are native device-alias members.

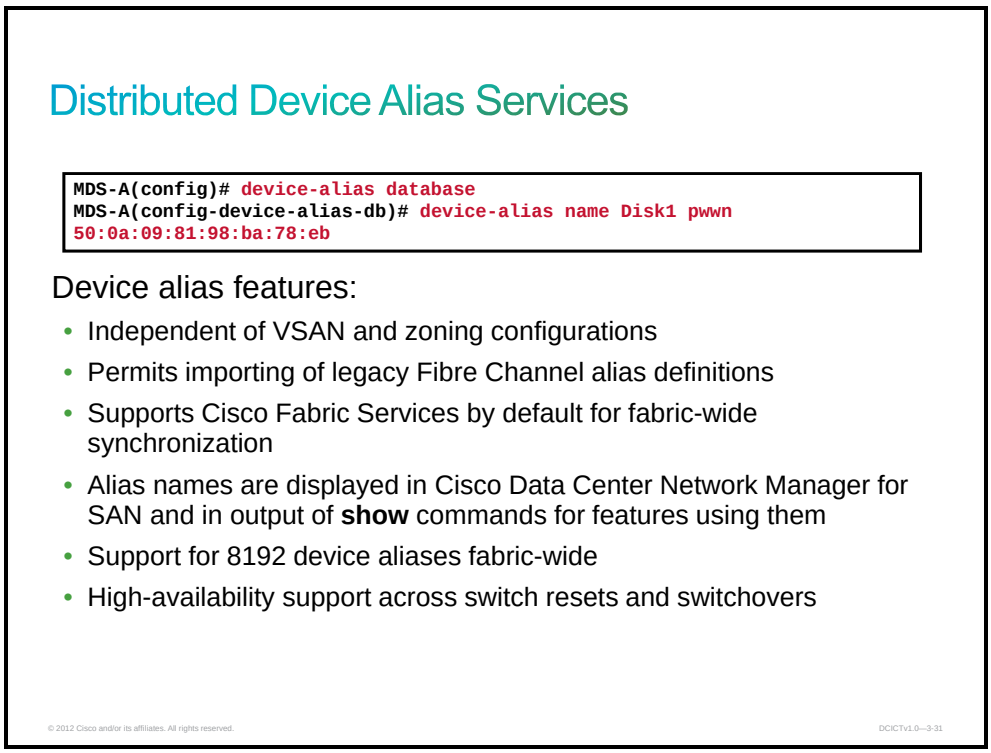
When the device-alias mode is changed from basic to enhanced, the applications are informed about the change and start accepting the device alias-based configuration in the native format.

If the device-alias mode is changed from enhanced to basic, then the applications check for an existing device-alias configuration in the native format. If the device alias is in the native format, the applications reject the request and the device-alias mode cannot be changed.

To change the mode back to basic, you must explicitly remove all native device-alias configurations, both on local and remote switches. Alternatively, you must replace all device-alias members with the corresponding pWWN before changing the mode back to basic.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

You can automate the process by using the force option. Use the **no device-alias mode enhanced force** command to enable applications to automatically replace all device-alias members with the corresponding pWWN. If a device-alias member does not have a corresponding pWWN mapping in the device-alias database, the configuration is removed.



Device aliases have these features:

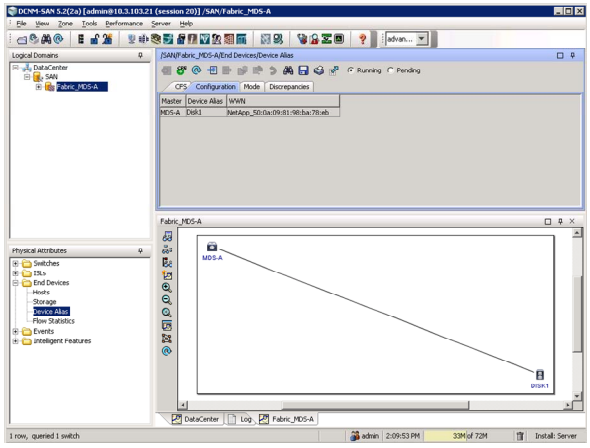
- Device-alias information is independent of your VSAN configuration. Aliases need to be defined only once and can then be used regardless of their VSAN.
- Device-alias configuration and distribution are independent of the zone server and the zone server database.
- You can import existing zone alias configurations without losing data.
- The device-alias application uses the Cisco Fabric Services infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the physical distribution scope. When you use device aliases to configure zones, IVR zones, or QoS features and display the configuration, you automatically see that the device aliases are displayed along with their respective pWWN.
- As many as 8192 device aliases can be configured fabric-wide.
- Distributed Device Alias Services is a highly available process. The device-alias databases are preserved across switch restarts and switchovers, and the system can be restarted or switched over during Cisco Fabric Services distributions or merges.

https://t.me/learningnets

Device Alias Configuration Verification

```
MDS-A# show device-alias database
device-alias name Disk1 pwwn 50:0a:09:81:98:ba:78:eb

Total number of entries = 1
```



© 2012 Cisco and/or its affiliates. All rights reserved.

Device-alias information can be viewed by using these **show** commands:

- Display all configured device aliases from the effective database

```
switch# show device-alias database
```

- Display the pending database with no modifications

```
switch# show device-alias database pending
There are no pending changes
```

- Display the pending database with modifications

```
switch# show device-alias database pending
```

- Display the difference between the pending and effective database

```
switch# show device-alias database pending-diff
```


https://t.me/learningnets

Device Alias Configuration Verification (Cont.)

```
MDS-A# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/9            10    0x6f0000  50:0a:09:81:98:ba:78:eb 50:0a:09:80:88:ba:78:eb
                         [Disk1]

Total number of flogi = 1.
```

```
MDS-A# show fcns database

VSAN 10:
-----------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
-----------------------------------------------------------------------
0x6f0000    N     50:0a:09:81:98:ba:78:eb (NetApp)        scsi-fcp:target
                  [Disk1]

Total number of entries = 1
```

© 2012 Cisco and/or its affiliates. All rights reserved.

- Display the device alias in the FLOGI database

  `switch# show flogi database`

- Display the device alias in the FCNS database

  `switch# show fcns database`

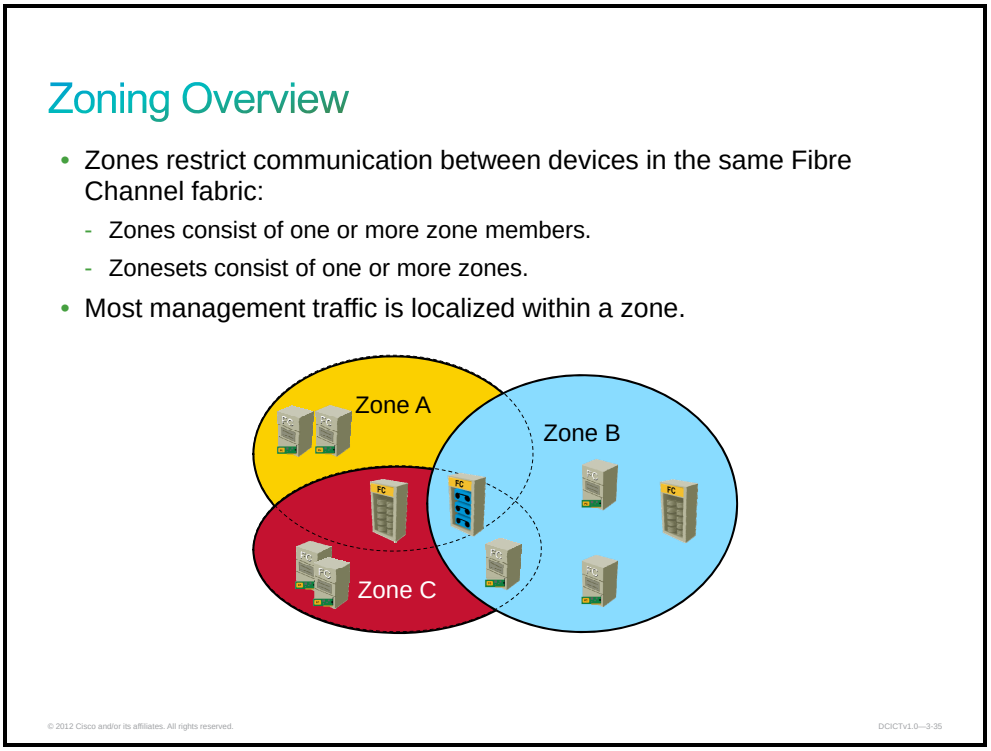
- Display statistics for the device alias application

```
switch# show device-alias statistics
      Device Alias Statistics
============================================
Lock requests sent: 2
Database update requests sent: 1
Unlock requests sent: 1
Lock requests received: 1
Database update requests received: 1
Unlock requests received: 1
Lock rejects sent: 0
Database update rejects sent: 0
Unlock rejects sent: 0
Lock rejects received: 0
Database update rejects received: 0
Unlock rejects received: 0
Merge requests received: 0
Merge request rejects sent: 0
Merge responses received: 2
Merge response rejects sent: 0
Activation requests received: 0
Activation request rejects sent: 0
Activation requests sent: 2
Activation request rejects received: 0
```


https://t.me/learningnets

Verify Fibre Channel Zoning on a Cisco MDS 9000 Series Multilayer Switch

Zoning is used to provide security and to restrict access between initiators and targets on the SAN fabric. This topic discusses the **show** commands that are used to verify the zone configuration.



With many types of servers and storage devices on the network, security is crucial. For example, if a host were to gain access to a disk that was being used by another host, potentially with a different operating system, the data on this disk could become corrupted. To avoid any compromise of crucial data within the SAN, zoning allows the user to overlay a security map. This map dictates which devices, namely hosts, can see which targets, thus reducing the risk of data loss.

As the figure shows, a zone set consists of one or more zones:

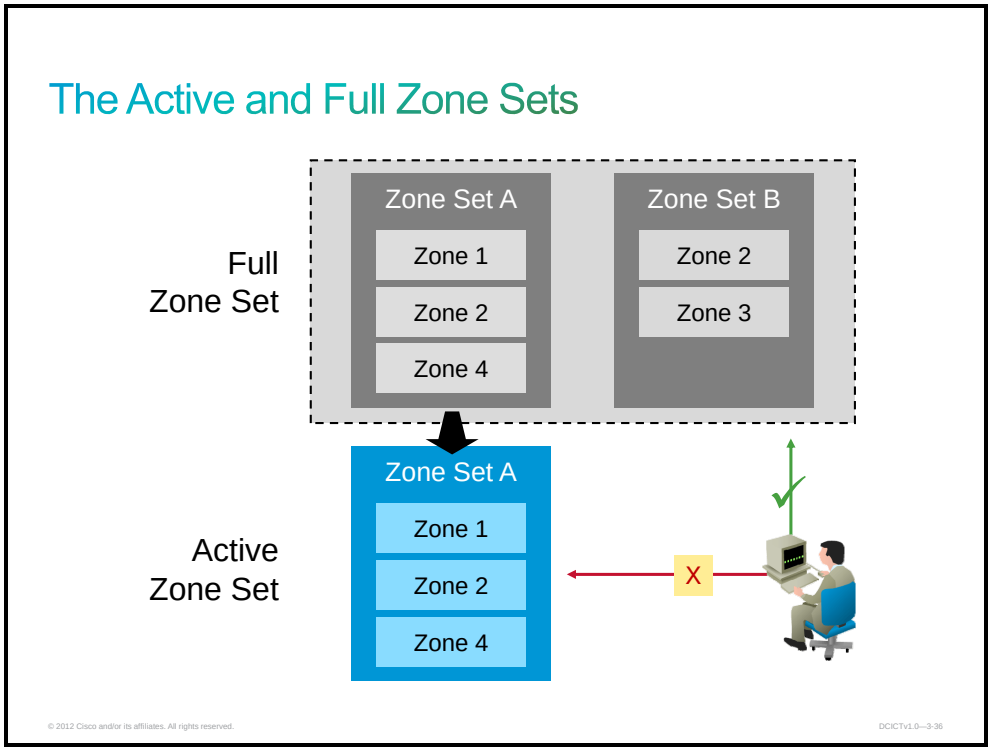
- A zone set can be activated or deactivated as a single entity across all switches in the fabric.
- Only one zone set can be activated at any time.
- A zone can be a member of more than one zone set.
- A zone consists of multiple zone members. Members in a zone can access one another. Members in different zones cannot access one another.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Uses for Zoning

There are several typical uses for zoning:

- Separate initiators from their targets. Frequently, each initiator port belongs in a separate zone with its targets.
- Separate devices that use different operating systems. This practice is used to protect some operating systems from treating disks that are formatted by other operating systems as blank disks and potentially taking over and overwriting their storage.
- Separate devices that have no need to communicate with other devices in the fabric or that have classified data.
- Separate devices into departmental, administrative, or some other functional groupings.
- Localize management traffic. Most management traffic does not cross zone boundaries. Zones help to reduce the impact of management traffic on the fabric.

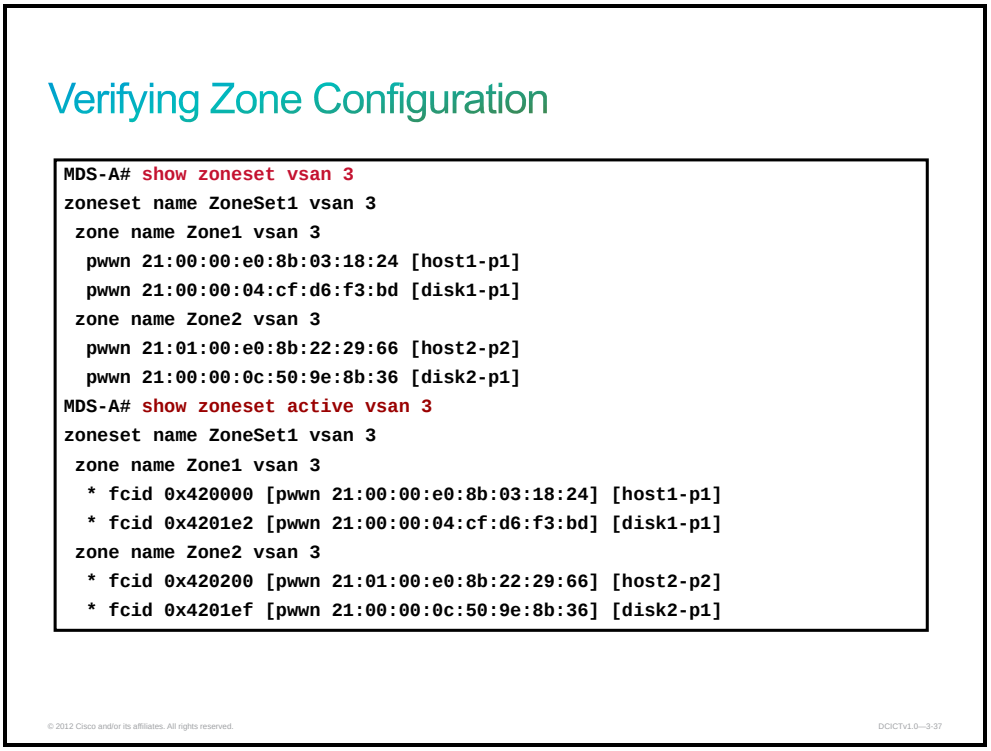


When you create a zone set, that zone set becomes a part of the full zone set, which is a database of all zone sets. Only one zone set can be active at a time. When you activate a zone set, the switch makes a copy of that zone set. The copy becomes the active zone set.

The active zone set cannot be modified. When you need to modify the active zone, you modify the original copy of the zone set in the full zone set database. When you want the changes to take effect, you then reactivate the zone set.

When the activation is finished, the active zone set is automatically stored in persistent configuration. This process enables the switch to preserve the active zone set information across switch resets. The active zone set is also propagated to all other switches in the fabric so that they can enforce zoning in their respective switches.

https://t.me/learningnets



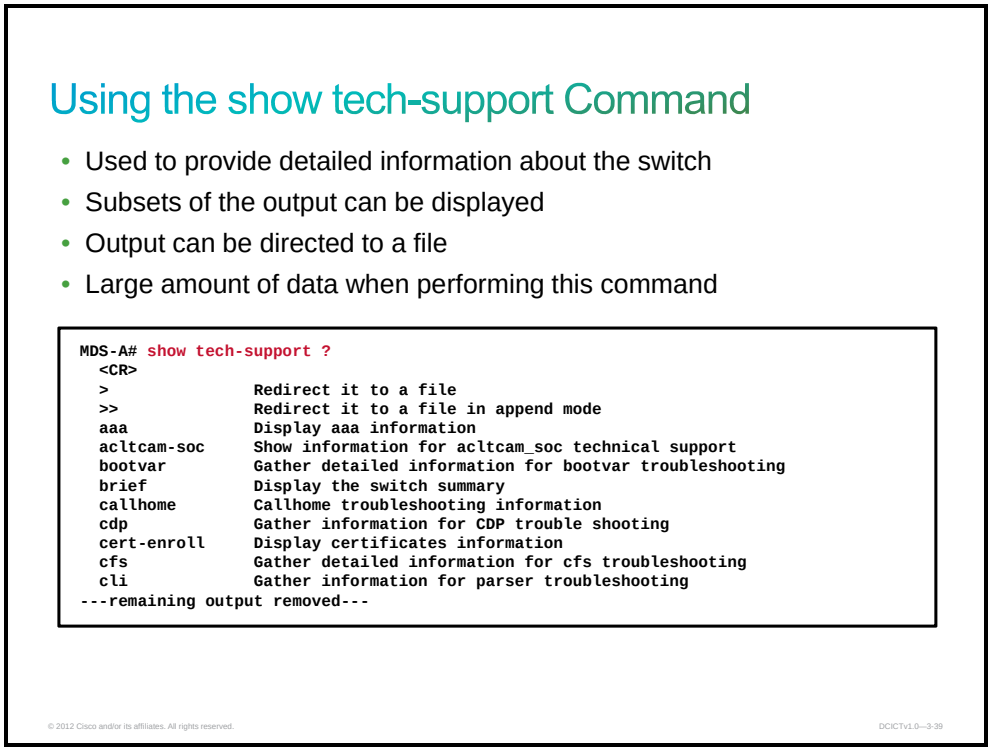
To display the current zoning configuration on the local switch, use the command **show zone set** in EXEC mode.

To verify the current active zone set, use the **show zoneset active** command in EXEC mode. The asterisks (*) indicate that a device is visible (online). A missing asterisk may indicate an offline device or an incorrectly configured zone, possibly a mistyped pWWN.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Collect Technical Support Data on a Cisco MDS 9000 Series Multilayer Switch

Use the **show tech-support** command when you require detailed information relating to the switch, including configuration, hardware, and software details, for analysis purposes. If the full details are required, then you should direct the output to a file so that it can be copied off the switch and analyzed offline.



If a specific component is needed for troubleshooting purposes, then you can use this command with additional parameters to be more specific.

https://t.me/learningnets

Summary

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

Summary

- The Cisco MDS 9000 Series Multilayer switch can be configured initially by using the CLI Setup Script.
- Cisco NX-OS can be upgrade from the CLI by using the **install all** command.
- Cisco NX-OS uses licensed features that can be upgraded nondisruptively so that additional features can be used.
- Devices log in to the fabric by using the FLOGI and PLOGI processes. These processes can be confirmed at the CLI by using the **show flogi database** and **show fcns database** commands.

© 2012 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- To verify that devices have accessibility to storage, use the **show zone** and **show zoneset** commands.
- The **show tech-support** command provides detailed output regarding the software and hardware configuration of the switch.

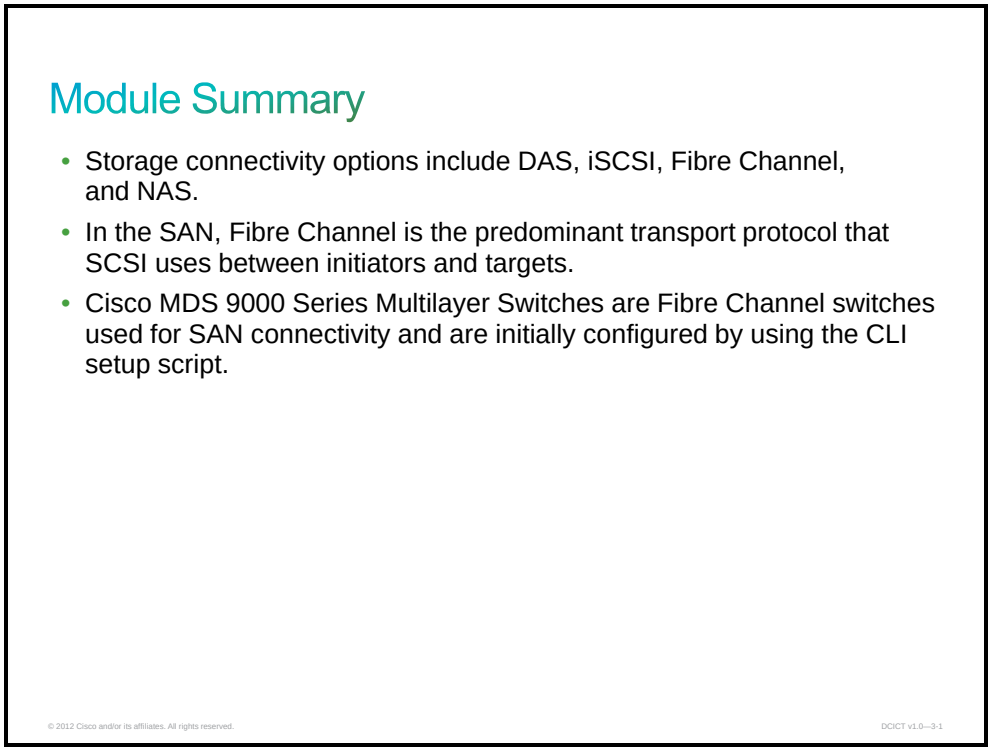
© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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



This module discussed remote storage-connectivity options for hosts and their relevant protocols. When identifying which remote storage to use, you must first compare block- and file-based storage. File-based storage generally uses the Network File System (NFS) and Common Internet File System (CIFS) protocols, whereas block-based storage uses the SCSI protocol.

File-based storage that uses NFS often uses a network-attached storage (NAS) infrastructure. Block-based storage uses either iSCSI as the transport protocol or a SAN infrastructure that uses the Fibre Channel Protocol (FCP).

Before network-based infrastructures, SCSI used direct-attached storage (DAS). DAS has scalability limitations, so the Fibre Channel Arbitrated Loop (FC-AL), followed by SAN, was developed to improve scalability and to provide better utilization of storage devices.

On a SAN, Fibre Channel is the predominant transport protocol. Fibre Channel is a layered model that uses Fibre Channel Generic Services (FC-GS) at FC-3 to provide the services that are required to support host and storage SAN connectivity. Each device in the SAN has a world wide name (WWN) that is embedded into an interface that is used to connect to the fabric. There is also a Fibre Channel ID (FCID) to provide a routable address for reaching devices that are remotely connected. To reach remote storage devices, each device must go through a login process that allows it to identify the remote storage devices that it may use. For security purposes, zoning is used in the fabric to restrict access between hosts and storage devices.

Cisco MDS 9100 Series Multilayer Switches in the Cisco MDS 9000 Series product range provide Fibre Channel switching capabilities. This switch can be configured by using a CLI setup script for basic functionality and the CLI for verification of registration and login of devices in the fabric.

https://t.me/learningnets

References

For additional information, refer to these resources:

- Kembel, R. W. *Fibre Channel: A Comprehensive Introduction.* Hingham, Massachusetts: Northwest Learning Associates, inc.; 2009.
- Cisco Systems, Inc. Cisco DCNM Fundamentals Guide, Release 5.x. http://www.cisco.com/en/US/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/fund/DCNM-SAN-LAN_5_2/DCNM_Fundamentals.html.
- Cisco Systems, Inc. Cisco MDS 9000 NX-OS and SAN-OS Software Configuration Guides. http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html.

 © 2012 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 two protocols are file-based? (Choose two.) (Source: Comparing Storage-Connectivity Options in the Data Center)

A) CIFS
B) Fibre Channel
C) SCSI
D) NFS
E) iSCSI

Q2) Which protocol is used primarily on UNIX or Linux operating systems to provide access to remote storage devices? (Source: Comparing Storage-Connectivity Options in the Data Center)

A) CIFS
B) Fibre Channel
C) TCP/IP
D) NFS

Q3) At which layer of the OSI protocol stack does NFS reside? (Source: Comparing Storage-Connectivity Options in the Data Center)

A) Layer 7
B) Layer 6
C) Layer 5
D) Layer 4

Q4) How many devices can be daisy-chained on a parallel SCSI bus? (Source: Comparing Storage-Connectivity Options in the Data Center)

A) 4
B) 8
C) 16
D) 32

Q5) A word is the smallest unit of data that Fibre Channel uses. How many bits does a word consist of? (Source: Comparing Storage-Connectivity Options in the Data Center)

A) 16
B) 32
C) 64
D) 128

Q6) How many sequences make up a single exchange in a SCSI-FCP write operation? (Source: Comparing Storage-Connectivity Options in the Data Center)

A) 1
B) 2
C) 3
D) 4

https://t.me/learningnets

Q7) How many usable addresses are on an FC-AL for servers and storage devices? (Source: Describing Fibre Channel Storage Networking)

A) 32
B) 64
C) 126
D) 127

Q8) Which type of fabric port supports one or more node loop ports? (Source: Describing Fibre Channel Storage Networking)

A) F port
B) E port
C) NP port
D) FL port

Q9) Which type of port is used to create an ISL on a Fibre Channel SAN? (Source: Describing Fibre Channel Storage Networking)

A) E
B) TE
C) TN
D) NP

Q10) Which process allows an N Port to exchange information about ULP support with its target N Port, to ensure that the initiator and target process can communicate? (Source: Describing Fibre Channel Storage Networking)

A) FLOGI
B) PLOGI
C) PRLI
D) PRLO

Q11) Which feature is used to ensure that the initiator (server) has access to the correct storage device and cannot corrupt data from other initiators? (Source: Describing Fibre Channel Storage Networking)

A) zoning
B) VSANs
C) ACLs
D) VLANs

Q12) Which mapping table is inside the front-end array controller and determines which LUNs are advertised through which storage array ports? (Source: Describing Fibre Channel Storage Networking)

A) LUN mapping
B) LUN masking
C) zoning
D) LUN zoning

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Q13) For how many days may a licensed feature be evaluated before a license must be installed to ensure continual operation of that feature? (Source: Verifying Fibre Channel Communications on Cisco MDS 9000 Series Multilayer Switches)

A) 180
B) 120
C) 60
D) 30

Q14) Which database stores name entries for all hosts in the fabric? (Source: Verifying Fibre Channel Communications on Cisco MDS 9000 Series Multilayer Switches)

A) FLOGI
B) distributed device
C) FCNS
D) zoning

https://t.me/learningnets

Answers

Q1) A, D
Q2) D
Q3) A
Q4) C
Q5) B
Q6) D
Q7) C
Q8) D
Q9) A
Q10) C
Q11) A
Q12) B
Q13) B
Q14) C

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Module 4

Cisco Data Center Unified Fabric

Overview

This module describes the function and operation of Cisco data center unified fabric technologies.

Module Objectives

Upon completing this module, you will be able to describe the function and verify correct operation of Cisco data center unified fabric technologies. This ability includes being able to meet these objectives:

- Describe the four IEEE standards that allow a lossless Fibre Channel fabric to be transported over Ethernet
- Make physical connections from the host CNA to a Cisco Nexus 5000 Series switch and from the Cisco Nexus 5000 Series switch to northbound LAN and SAN domains
- Describe the role of NIV with a Cisco Nexus 2232 fabric extender in FCoE deployments

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 1

Describing DCB

Overview

The purpose of this lesson is to describe Data Center Bridging (DCB) and the IEEE standards that allow a lossless Fibre Channel fabric to be transported over Ethernet.

Objectives

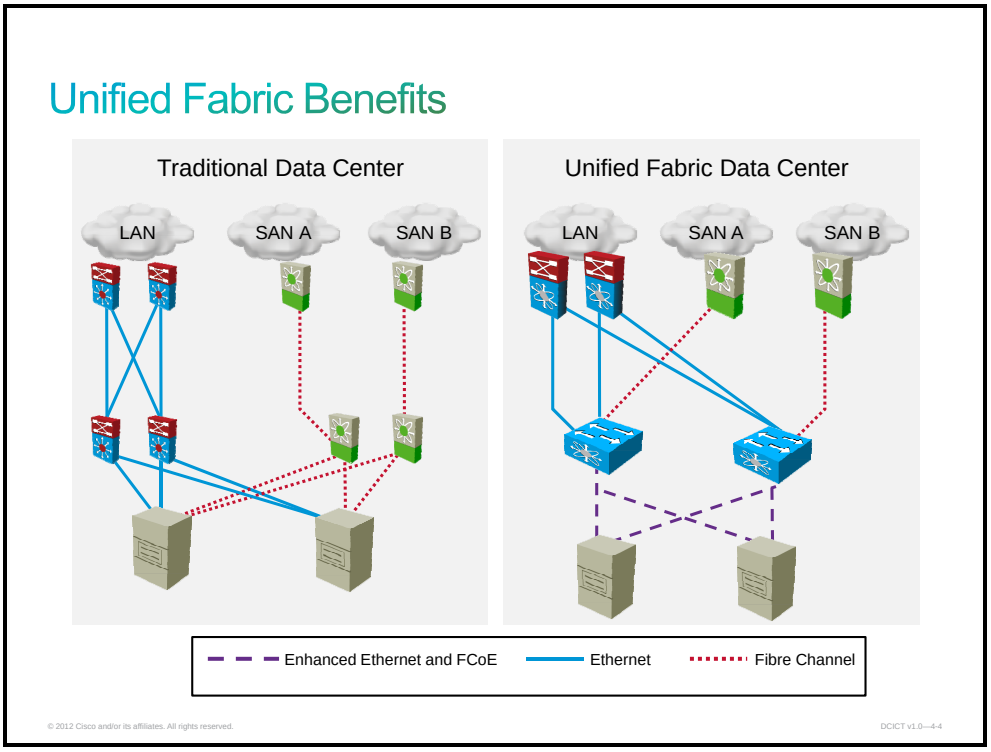
Upon completing this lesson, you will be able to describe the four IEEE standards that allow a lossless Fibre Channel fabric to be transported over Ethernet. You will be able to meet these objectives:

- Describe the benefits of unified fabric and FCoE
- Describe the IEEE standards that enable FCoE and FIP
- Describe PFC
- Describe ETS
- Describe DCBX

https://t.me/learningnets

Unified Fabric Benefits

A unified fabric comprises technologies and features to enhance the capabilities and reduce the complexity of the data center. One such technology is the ability to provide I/O consolidation of both network and SAN traffic over 10-Gb/s Ethernet links. This topic describes how Cisco Unified Fabric reduces capital costs and management complexity in the data center.



One feature of Cisco Unified Fabric is the ability to provide I/O consolidation over 10-Gb/s Ethernet links.

The left-hand side of the figure shows the traditional data center infrastructure. This traditional infrastructure has separate LAN and SAN networks in the data center.

The right-hand side of the figure shows the unified fabric data center infrastructure, in which LAN and SAN traffic is consolidated over 10-Gb/s Ethernet connections. Consolidating I/O at various layers of the data center provides these benefits:

- **Reduction of server adapters:** Host bus adapters (HBAs) and network interface cards (NICs) can be replaced with converged network adapters (CNAs).
- **Simplification of cabling:** Fewer adapters means fewer cables.
- **Layer 2 multipath access:** You can take advantage of features such as virtual port channel (vPC) and Cisco FabricPath
- **Lower total cost of ownership (TCO)**
- **Investment protection**
- **Consistent operational model:** The operational model remains consistent between LAN and SAN operations.
- **Centralization:** The consolidated architecture is centralized, which enables easier and more effective centralized management.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

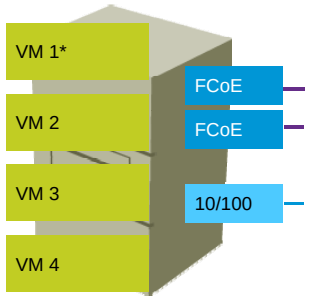
Unified Fabric I/O Requirements

LAN/IP

- Ethernet predominant
- Losing some frames is okay for TCP

Storage

- Fibre Channel model predominant
- Losing frames is not an option



FCoE = Fibre Channel over Ethernet

*VM = Virtual Machine

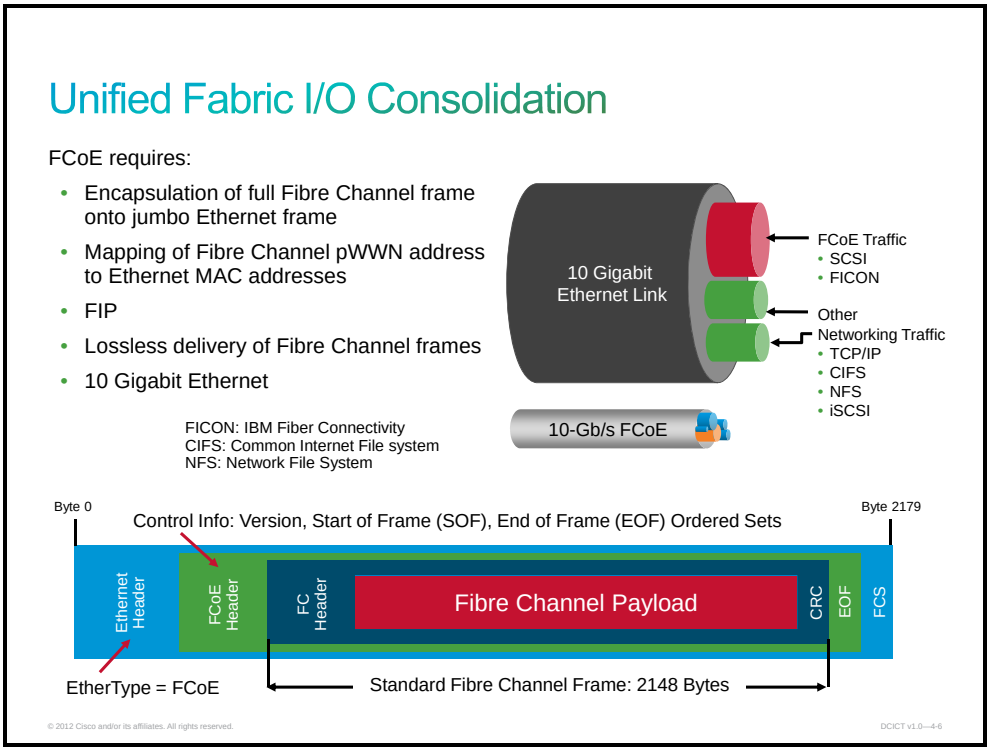
© 2012 Cisco and/or its affiliates. All rights reserved.

When performing I/O consolidation in the data center, you should take certain requirements into consideration:

- **LAN traffic:** Traditional Ethernet applications predominate in this environment, with TCP/IP being the accepted standard. In the LAN environment, dropped packets are expected and required to resolve the state of congestion. Any dropped or lost packets are retransmitted with minimal or no impact to application operation.
- **SAN traffic:** Fibre Channel has predominated in this environment. Dropping of frames is not allowed. A dropped frame introduces a disruption for storage applications. Most storage applications expect all Fibre Channel frames and upper-layer protocol data units (PDUs) to arrive in order.

Because of the different traffic characteristics, different parallel quality of service (QoS) congestion management and avoidance mechanisms are required. When you perform I/O consolidation, data center bridging features are required.

https://t.me/learningnets



Fibre Channel over Ethernet (FCoE) is a protocol that is based on the Fibre Channel layers that are defined by the ANSI T11 committee. FCoE replaces the lower layers of Fibre Channel with unified fabric I/O consolidation over Ethernet.

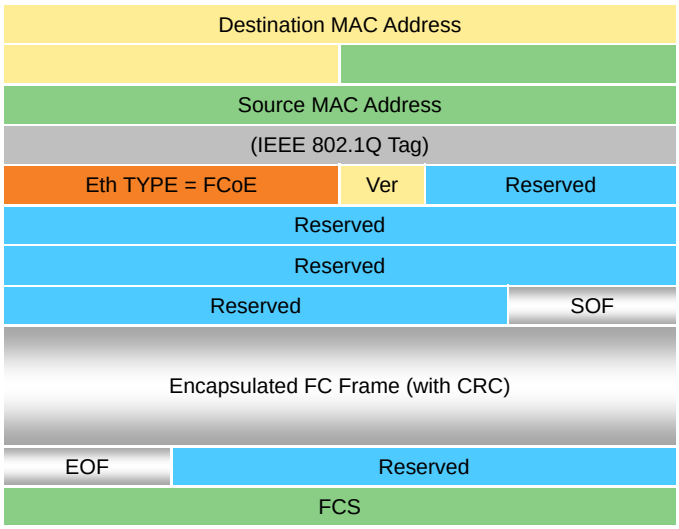
FCoE has several minimum requirements:

- The entire Fibre Channel frame must be carried in the payload of a single Ethernet frame (jumbo frame).
- Fibre Channel port world wide name (pWWN) addresses must be mapped to Ethernet MAC addresses.
- An FCoE Initialization Protocol (FIP) is needed to provide login for Fibre Channel devices across a unified fabric that supports I/O consolidation.
- Lossless delivery of Fibre Channel frames is necessary.
- A minimum 10-Gb/s Ethernet platform is required.

FCoE traffic consists of a Fibre Channel frame that is encapsulated within an Ethernet frame. The Fibre Channel payload can in turn carry Small Computer Systems Interface (SCSI) messages and data.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

FCoE Architecture



© 2012 Cisco and/or its affiliates. All rights reserved.

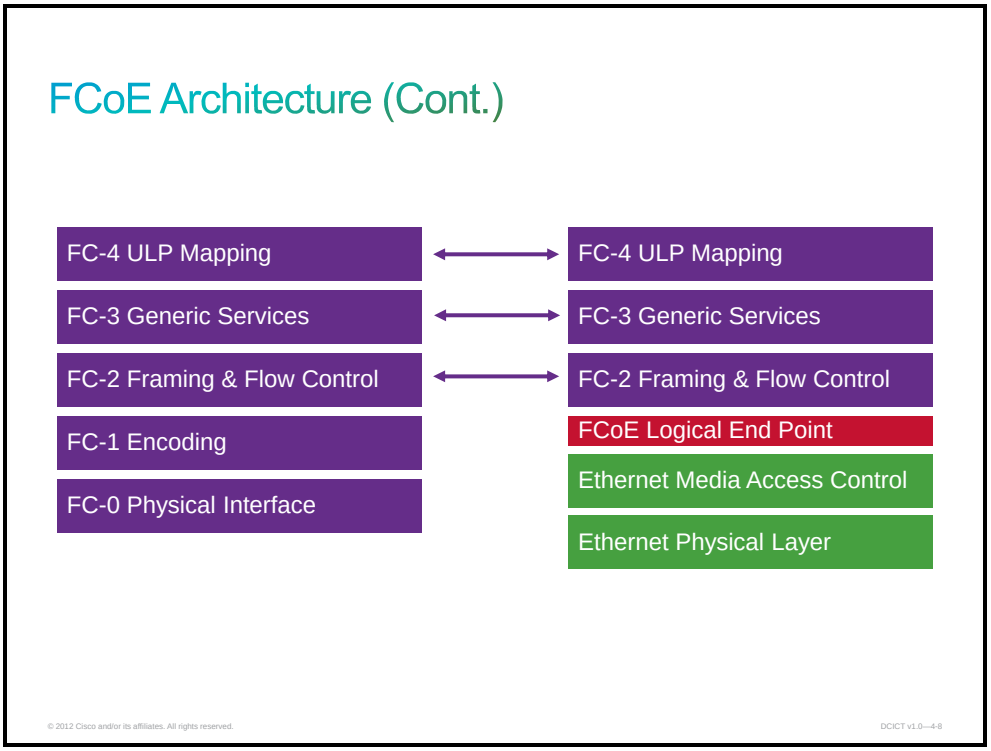
FCoE is encapsulated in an Ethernet frame with a dedicated EtherType of 0x8906. The packet has a 4-bit version field. All other header fields in the frame (the source and destination MAC addresses, VLAN tags, and frame markers) are standard Ethernet fields. The reserved bits pad the FCoE frame to the IEEE 802.3 minimum frame length of 64 bytes.

A Fibre Channel frame consists of 36 bytes of headers and as many as 2112 bytes of data for a total maximum size of 2148 bytes. The encapsulated Fibre Channel frame has all the standard headers, allowing it to be passed to the storage network without further modification. To accommodate the maximum Fibre Channel frame in an FCoE frame, the class-fcoe is defined with a default maximum transmission unit (MTU) of 2240 bytes.

The Ethernet frames that are sent by the switch to the adapter may include the IEEE 802.1Q tag. This tag includes a field for the class of service (CoS) value that priority flow control (PFC) uses. The IEEE 802.1Q tag also includes a VLAN field.

The Cisco Nexus switch expects frames from a FIP T11-compliant CNA to be tagged with the VLAN tag for the FCoE VLAN. Frames that are incorrectly tagged are discarded.

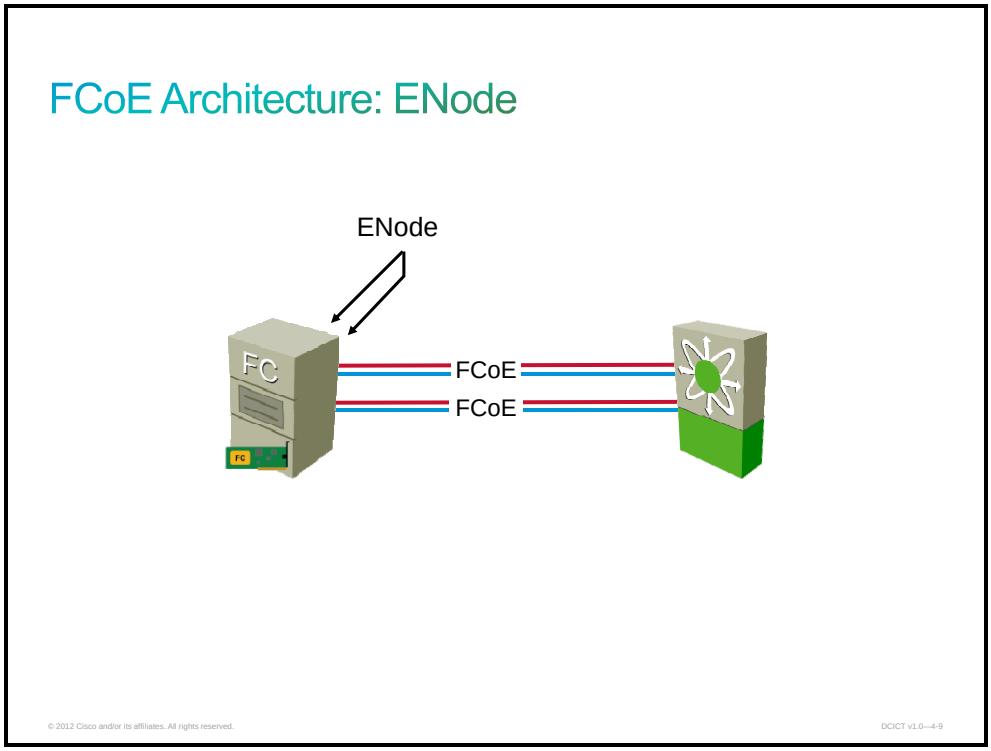
https://t.me/learningnets



FCoE Logical End Point

The FCoE Logical End Point (FCoE_LEP) is responsible for the encapsulation and de-encapsulation that is necessary to transport Fibre Channel frames over Ethernet. The figure shows that the FCoE_LEP has the standard Fibre Channel layers, starting with FC-2 and continuing up the Fibre Channel Protocol (FCP) stack. To higher-level system functions, this feature gives the appearance that the FCoE network is a standard Fibre Channel network. This compatibility allows all the tools that are used in native Fibre Channel to be used in an FCoE environment. Below the FCoE_LEP is the standard Ethernet media and physical layers for 10 Gigabit Ethernet, with enhancements that allow lossless Ethernet. Using the Ethernet standards allows FCoE to leverage a significant amount of existing technology.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



FCoE Node

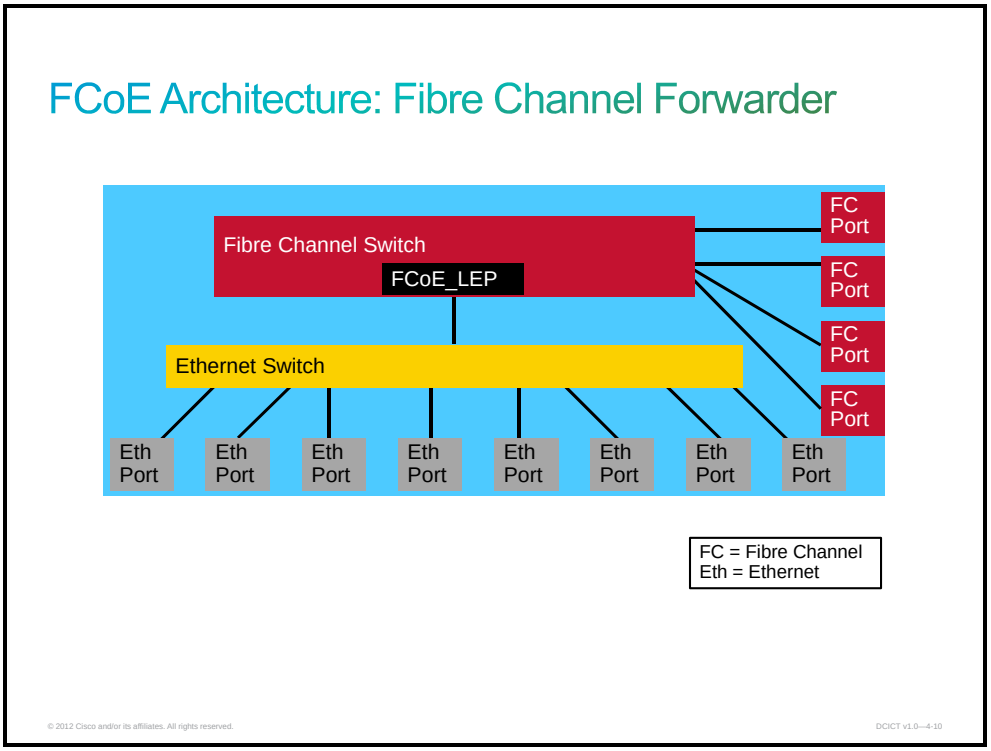
An FCoE node (ENode) is an Ethernet NIC that can operate as a Fibre Channel node port (N Port). An ENode consists of a Fibre Channel node with an FCoE_LEP attached to an Ethernet port. Each Ethernet port has a unique MAC address that the FCoE_LEP can use to transmit and receive FCoE frames. A CNA is an example of an ENode, as is a 10 Gigabit Ethernet NIC using the FCoE software stack.

The ENode creates at least one virtual N Port (VN Port) to the FCoE-enabled lossless Ethernet link. Each VN Port has an ENode MAC. The VN Port represents an FCoE controller.

The FCoE controller that associates with the ENode MAC supports the VN Port and FCoE_LEP pairing.

The VN Port-to-virtual F Port (VF Port) communication is part of the Fibre Channel Backbone Ethernet (FC-BB-E) standard, which is part of the FCoE Fibre Channel Backbone 5 (FC-BB-5) standard.

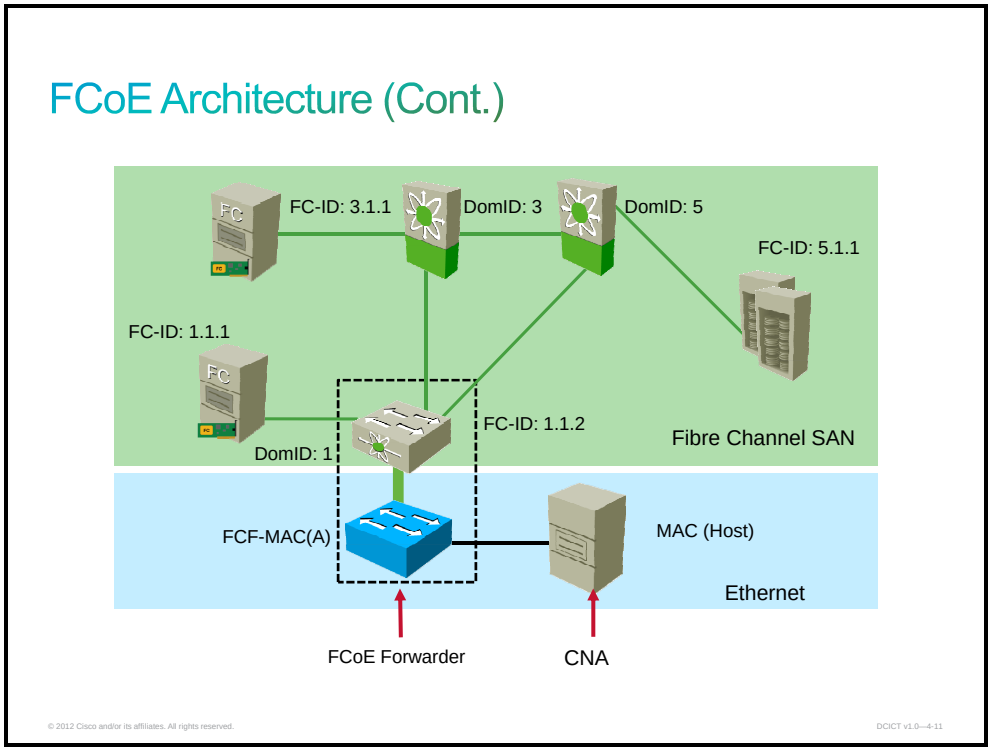
https://t.me/learningnets



Fibre Channel Forwarder

A Fibre Channel Forwarder (FCF) is a switch that contains both a Fibre Channel switch and an Ethernet switch. One or more FCoE_LEPs are used to attach the Fibre Channel switch internally to the Ethernet switch. The FCoE_LEP defines the logical connection between the VN Port of the end node and the VF Port of the switch. There is one universal MAC address per FCoE_LEP. The FCoE_LEP is responsible for encapsulating and de-encapsulating Fibre Channel frames for FCoE transportation. Fibre Channel traffic flows in to one end of the FCoE_LEP, and FCoE flows into the other end.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



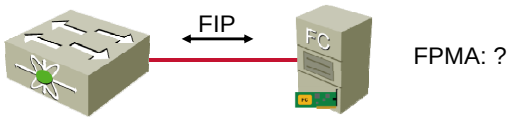
The FCoE standard supports the use of FCoE between an ENode and a directly attached FCF that is deployed as an access switch.

The figure shows a Fibre Channel network (in green) that has been extended over an Ethernet network (in blue) by using FCoE. The host with an FCoE_ENode is attached directly to a switch that acts as an FCF. The FCF is shown in the figure as a separate Ethernet and Fibre Channel switch that is connected by the FCoE_LEP. Note that the FCF and the ENode both have Ethernet and Fibre Channel characteristics, which can be seen through the Fibre Channel addressing that is associated with each device.

https://t.me/learningnets

FCoE Initialization Protocol

- FIP discovers FCoE-capable devices through DCB-capable switches.
- FIP discovers FCoE FCFs on a VLAN to establish a VN-VF link.
- The EtherType for FIP frames is different from those used for FCoE frames.



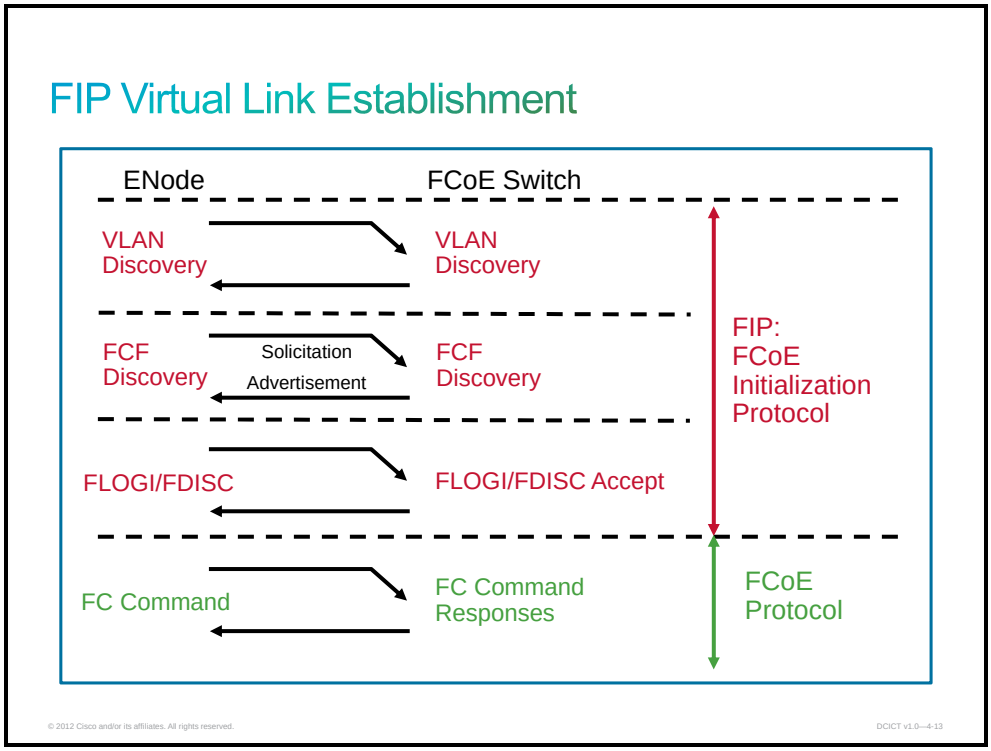
© 2012 Cisco and/or its affiliates. All rights reserved.

FIP is the FCoE control protocol that is responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices (ENodes or FCFs). During the virtual-link establishment phase, FIP first discovers the FCoE VLANs and remote virtual Fibre Channel interfaces. FIP then performs virtual-link initialization functions (such as fabric login) and fabric discovery or exchange link parameters (ELPs) similar to their native Fibre Channel equivalents. After the virtual link is established, Fibre Channel payloads can be exchanged on the virtual link. FIP remains in the background to perform virtual-link maintenance functions. FIP continuously verifies reachability between the two virtual Fibre Channel interfaces on the Ethernet network. FIP offers primitives to delete the virtual link in response to administrative actions.

FIP establishes virtual Fibre Channel links between VN Ports and VF Ports (ENode to FCF), as well as between virtual expansion ports (VE Ports) (FCF to FCF).

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Discovery and Virtual Link Establishment

FIP defines two discovery protocols as well as a protocol to establish virtual links between VN Ports and VF Ports. The figure shows a typical FIP protocol exchange. The exchange results in the establishment of a virtual link between the VN Port of an ENode and the VF Port of an FCF.

All the protocols are usually initiated by ENodes, although FCFs can generate unsolicited FIP advertisements.

Note these points:

- The FIP frames at the top and the FCoE frames at the bottom of the figure use different EtherTypes and encapsulations.
- FCoE frames encapsulate native Fibre Channel payloads.
- FIP frames describe a new set of protocols that have no reason to exist in native Fibre Channel definitions.
- FIP packets are built by using a globally unique MAC address, called the ENode MAC address, which is assigned to the CNA at manufacturing.
- FCoE packets are encapsulated by using a locally unique MAC address that is unique within the boundaries of the local Ethernet subnet. The FCF dynamically assigns this address to the ENode as part of the FIP virtual-link establishment process.

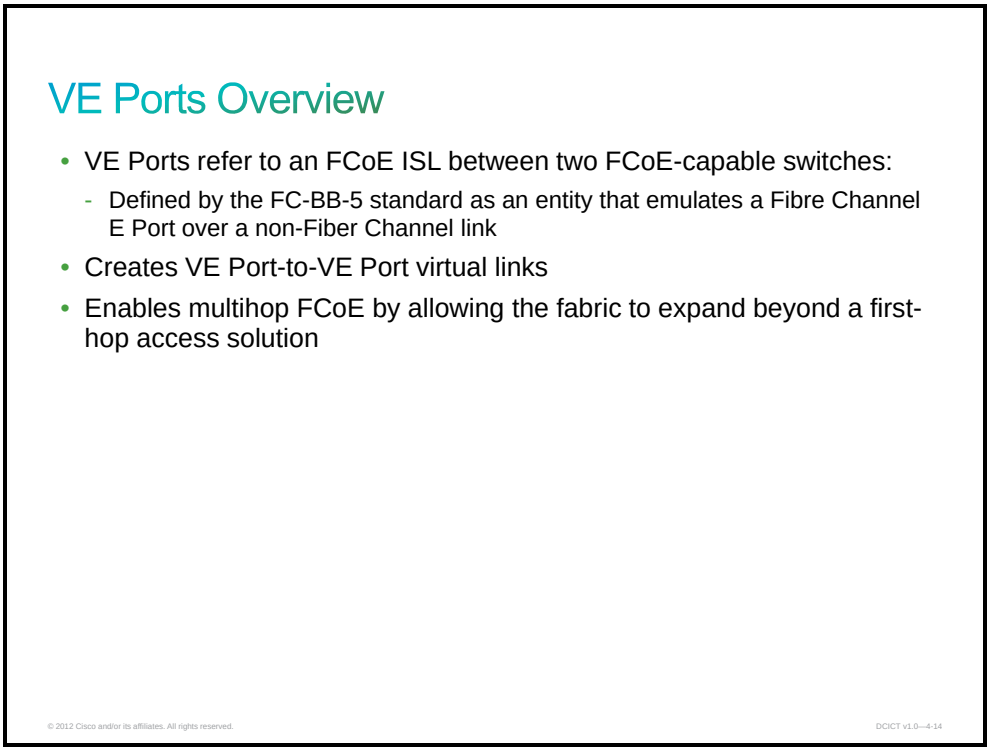
FIP VLAN discovery discovers the FCoE VLAN that the FIP protocols will use. FCoE encapsulation also uses this FCoE VLAN for Fibre Channel payloads on the established virtual link. FIP VLAN discovery is performed on the native VLAN. All other FIP protocols run on the discovered FCoE VLAN.

https://t.me/learningnets

The ENode sends a FIP VLAN discovery request to a multicast MAC address called All-FCF-MACs, which is a multicast MAC address to which all FCFs listen. The ENode discovers all FCFs and selects one for login. The final step is to inform the selected FCF of the intention to create a virtual link with its VF Port. After this step has been performed, Fibre Channel payloads can be exchanged on the newly established virtual link.

The FIP FLOGI and Discover Fabric (FDISC) unicast frames are almost identical to the native Fibre Channel FLOGI and FDISC frames that they replace. The VN Port sends a FLOGI or FDISC request, followed by the corresponding FLOGI or FDISC accept payload from the FCF.

The FDISC protocol is used when discovering which Cisco Fabric Services a device can use.

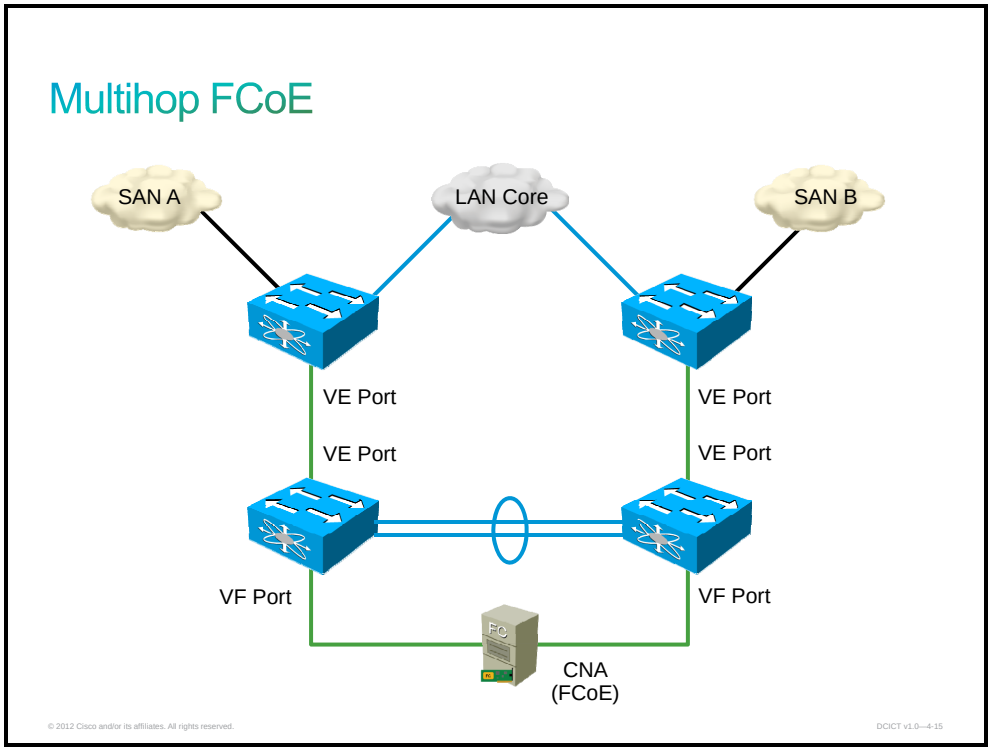


A VE Port refers to an FCoE interswitch link (ISL) between two FCoE-capable switches. The FC-BB-5 standard defines this specification as an entity that emulates a Fibre Channel E Port over a non-native Fibre Channel link. This emulation allows the creation of a VE Port-to-VE Port virtual link between directly connected FCFs.

VE Ports enable multihop FCoE by allowing the fabric to expand beyond a first hop access solution.

Whether using single-hop FCoE with VF Ports or multihop FCoE with VE Ports, the supported distance for FCoE remains at 984 feet (300 m).

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Multihop FCoE is supported across VE Ports that are established between two Cisco Nexus switches or between the Cisco Nexus switch and a Cisco MDS 9500 switch.

A VE Port emulates an E Port over a non-Fibre Channel link. VE Ports are supported over point-to-point links between FCFs. These links can be Ethernet interfaces or members of an Ethernet port channel interface. For each of the FCF-connected Ethernet interfaces, you must create and bind a virtual Fibre Channel (vFC) interface to the Ethernet interface.

VE Ports have these guidelines:

- Auto mode on the vFC interface is not supported.
- VE Port trunking is supported over FCoE-enabled VLANs.
- VE Port interface binding to MAC addresses is not supported.
- VE Ports are enabled for trunk mode by default.

The Spanning Tree Protocol (STP) is disabled on the FCoE VLANs on any interface to which a vFC is bound, including the interfaces to which the VE Ports are bound.

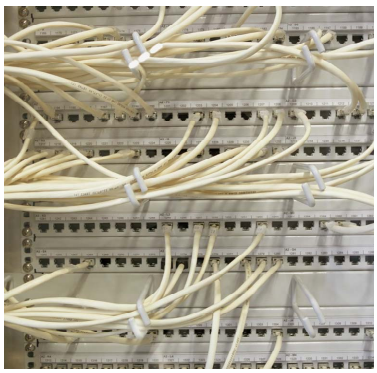
https://t.me/learningnets

IEEE Standards That Enable FCoE

To support I/O consolidation in the form of FCoE, certain data center Ethernet bridging enhancements are required. This topic describes the high-level definitions of IEEE standards that enable FCoE.

IEEE Standards

- 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)
  - PFC by traffic class
  - End-to-end congestion management and notification
  - Layer 2 multipathing
- Benefits of enhancements:
  - Eliminates transient and persistent congestion
  - Lossless fabric (no-drop storage links)
  - Deterministic latency for HPC clusters
  - Enables a converged Ethernet fabric for reduced cost and complexity

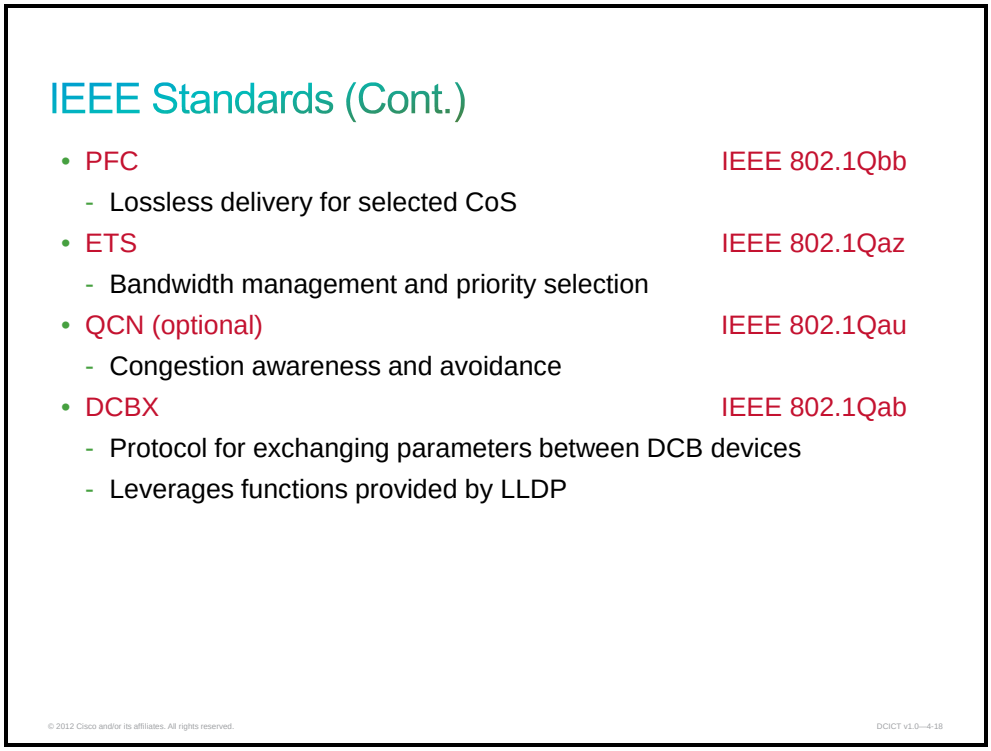
© 2012 Cisco and/or its affiliates. All rights reserved.

DCB Ethernet enhancements are required to achieve unified fabric I/O consolidation. These enhancements are being proposed and implemented:

- PFC
- Enhanced Transmission Selection (ETS)
- End-to-end congestion management (backward congestion notification/quantized congestion notification [BCN/QCN], an optional standard that Cisco does not implement)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



DCB is a collection of IEEE standards that provide the foundation for a unified fabric that supports I/O consolidation.

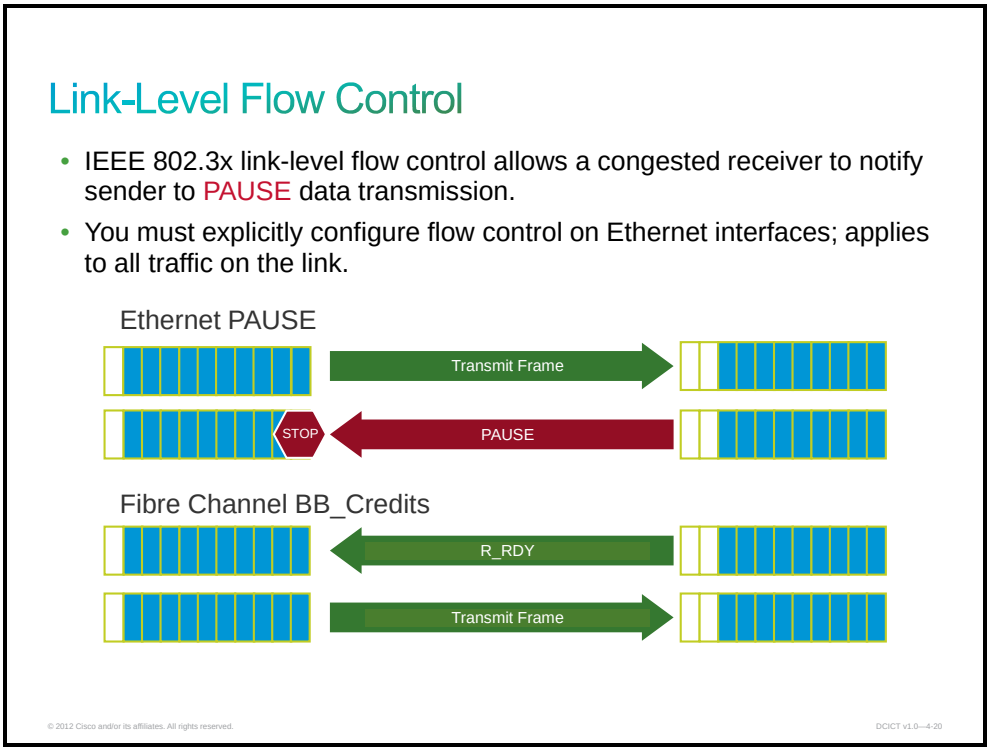
These standards are being proposed and implemented:

- **PFC:** The IEEE 802.1Qbb standard provides lossless delivery for selected CoS.
- **ETS:** The IEEE 802.1Qaz standard provides bandwidth management and priority selection.
- **QCN:** The 802.1Qau standard provides congestion awareness and avoidance. This standard is an optional standard that vendors can choose to implement. Cisco does not implement QCN.
- **Data Center Bridging Exchange (DCBX):** The IEEE 802.1Qab standard provides a protocol for exchanging parameters between DCB-capable devices and for leveraging functions that the Link Layer Discovery Protocol (LLDP) provides.

https://t.me/learningnets

Priority Flow Control

PFC is a mechanism for multiplexing lossless classes of traffic across the 10-Gb/s DCB Ethernet link. PFC allows priority to be provided to the FCoE channel, to ensure lossless transport of packets across DCB links and no loss of frames on the DCB links during congestion. This topic describes how PFC allows lossless Ethernet fabric links.



Link-Level Flow Control

Link-level flow control is required for a lossless fabric. Both Ethernet and Fibre Channel already contain mechanisms for link-level flow control, through the Ethernet pause feature and the Fibre Channel buffer-to-buffer credit (BB_Credit). Both technique stops traffic on the entire link, limiting their usefulness in a unified fabric. Also note that not all upper-level protocols require or desire a lossless fabric. TCP, for example, requires packet loss for congestion management.

The IEEE 802.3x link-level flow-control capability allows a congested receiver to communicate to the far end, asking it to pause its data transmission for a short period. The link-level flow control feature applies to all traffic on the link.

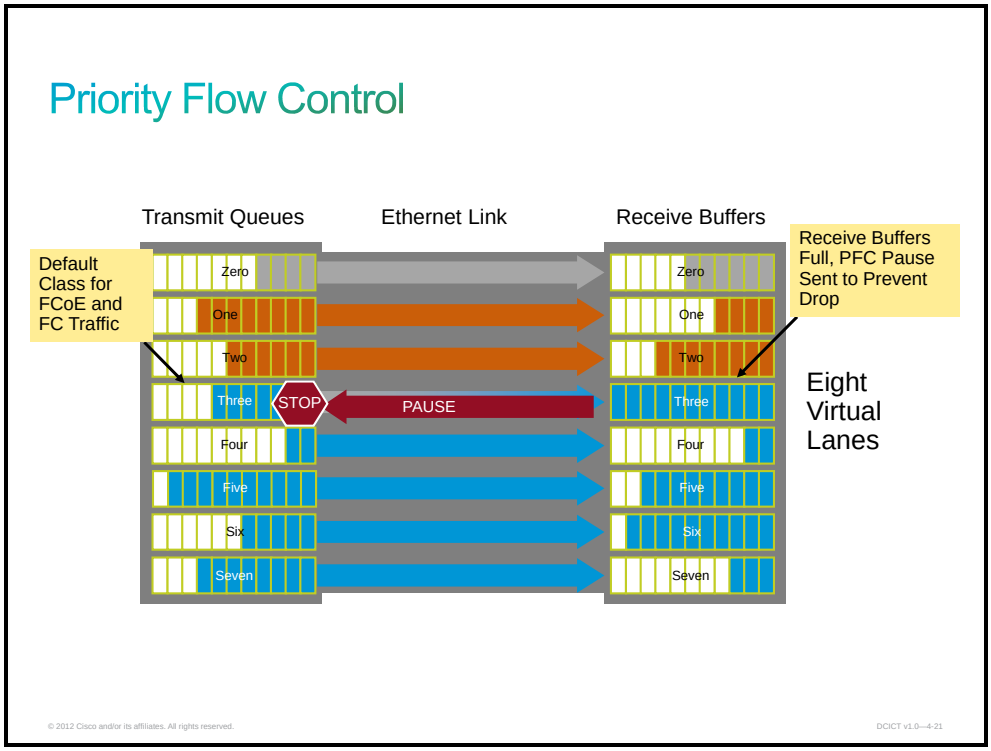
The transmit and receive directions are configurable separately. By default, link-level flow control is disabled for both directions.

On the Cisco Nexus switch, Ethernet interfaces do not automatically detect the link-level flow-control capability. You must explicitly configure the capability on the Ethernet interfaces.

On each Ethernet interface, the switch can either negotiate and enable PFC on DCB ports, or an administrator can configure link-level flow control on traditional Ethernet links. Both methods cannot be implemented at the same time.

© 2012 Cisco Systems, Inc.

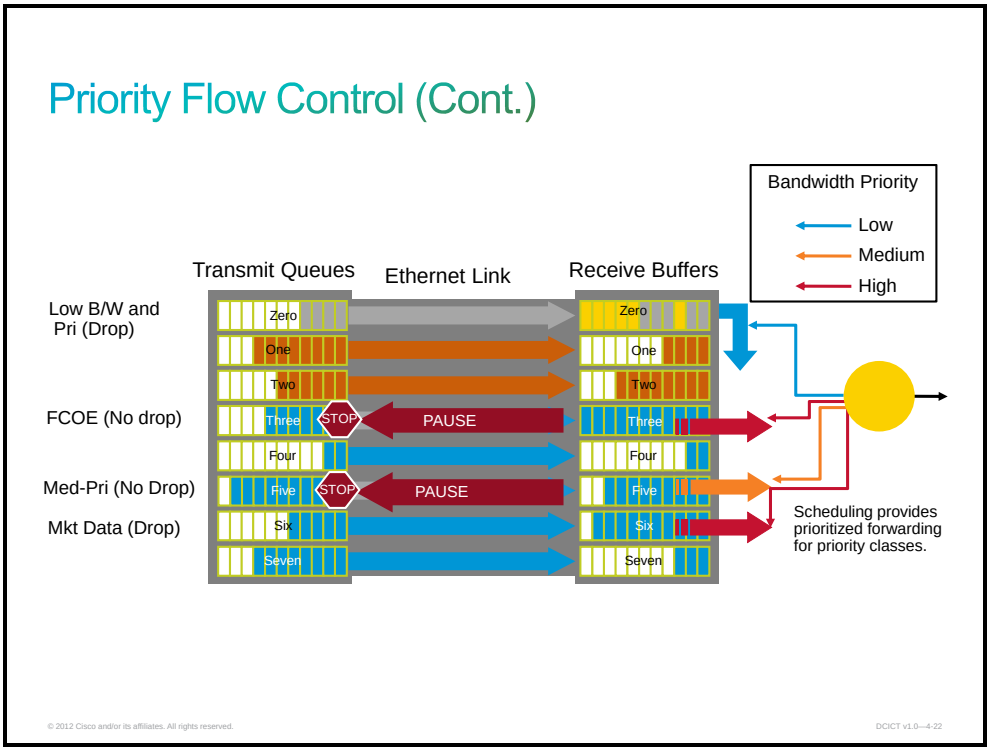
https://t.me/learningnets



PFC introduces a new pause frame format that differs from the traditional pause frame format. This new frame allows you to apply the pause functionality to specific classes of traffic, rather than all traffic, on a link. PFC applies pause functionality that is based on the IEEE 802.1p CoS value. When the switch enables PFC, it communicates to the remote switch port or to the CNA adapter which CoS values to apply the pause.

DCB interfaces use PFC to provide a lossless server to no-drop system classes. PFC implements pause frames on a per-class basis and uses the IEEE 802.1p CoS value to identify the classes that require lossless services.

https://t.me/learningnets

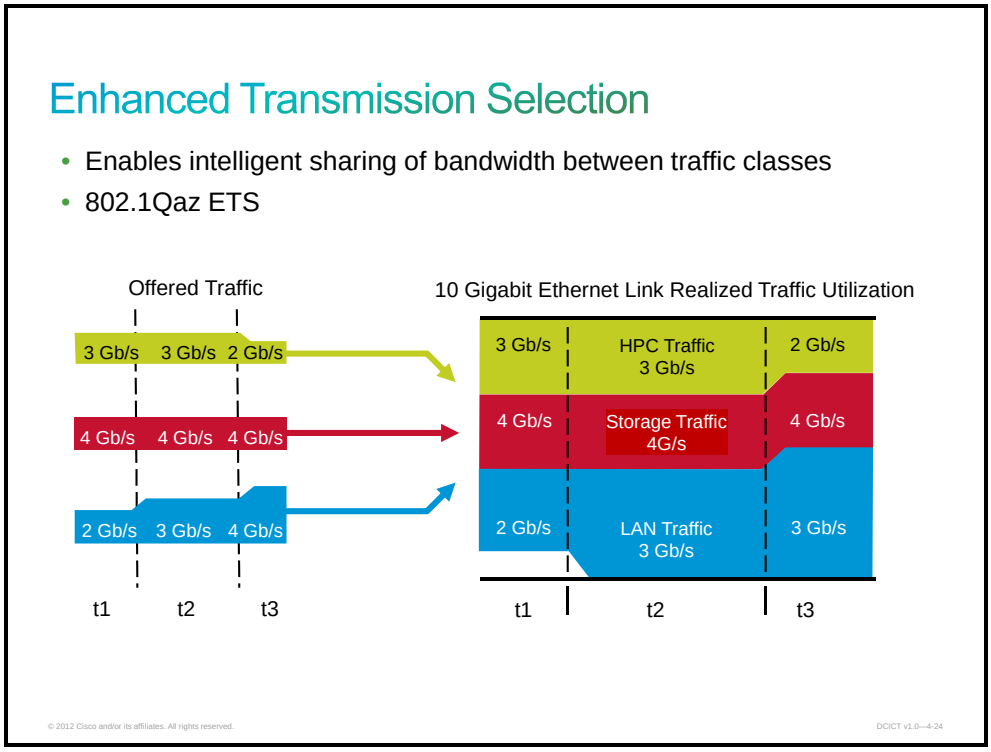


The figure shows how multiple classes of traffic can use the PFC mechanism. Low-priority traffic, which is also low-bandwidth, can be dropped if buffers fill, relying on TCP retransmissions for recovery. FCoE and Internet SCSI (iSCSI) traffic can be set to the no-drop FCoE class to prevent retransmissions and to apply suitable, required flow control for storage traffic. Medium-priority data traffic can be set to no-drop to pause when buffers hit a maximum mark. When combined with low-medium bandwidth allocation, the traffic will be paused to allow more bandwidth for latency-sensitive applications, such as voice or market data (as shown in the figure). There is no point in pausing voice because these applications do not tolerate jitter or latency. Pausing other traffic streams when buffers fill benefits these latency-sensitive applications.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Enhanced Transmission Selection

ETS is an IEEE 802.1Qaz standard that enables optimal QoS strategy for DCB links and bandwidth management of virtual links. ETS allows differentiation among traffic of the same priority class, thus creating priority groups. This topic describes how ETS contributes to a lossless Ethernet fabric.



ETS is also called priority grouping. Eight distinct virtual-link types can be created by implementing PFC on DCB links. Defining different traffic classes within the different PFC types can be advantageous.

ETS enables these differentiated treatments within the same priority class of PFC. This treatment provides prioritized processing that is based on bandwidth allocation, low latency, or best effort.

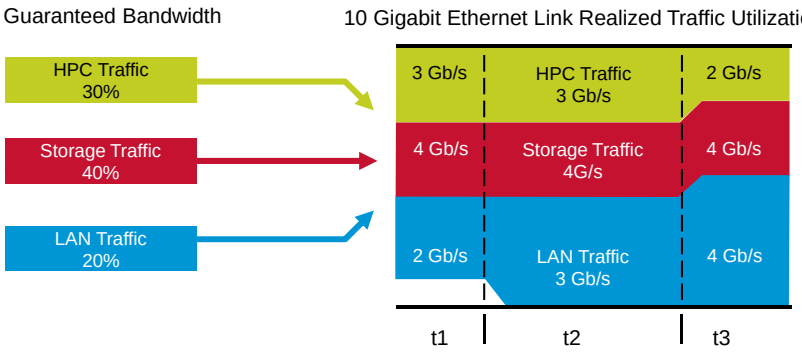
For example, an Ethernet class of traffic might have a high-priority designation and a best effort within that class. ETS allows differentiation between traffic of the same priority class, thus creating priority groups.

ETS provides the capability to group each type of data flow, such as storage or networking, and assigns an identification number to each traffic class group. The value of this new feature lies in the ability to manage bandwidth on the Ethernet link by allocating portions (percentages) of the available bandwidth to each group. Bandwidth allocation allows traffic from the different groups to receive a specific target service rate (such as 8 Gb/s for storage and 2 Gb/s for LAN). Bandwidth allocation provides QoS to applications.

https://t.me/learningnets

Enhanced Transmission Selection (Cont.)

- Bandwidth is guaranteed but can be used by other classes if not in use.
- ETS enables Intelligent sharing of bandwidth between traffic classes.
- Proposed standard is IEEE 802.1Qaz ETS.
- Bursty traffic in managed classes can exist alongside strict priority traffic classes.



© 2012 Cisco and/or its affiliates. All rights reserved.

Bandwidth management is an important requirement when consolidating I/O. The IEEE 802.1Qaz standard specifies ETS to support the allocation of bandwidth among traffic classes. When a given load in a traffic class does not fully utilize its allocated bandwidth, this standard allows other traffic classes to use the available bandwidth. This feature helps to accommodate the bursty nature of some traffic classes while maintaining bandwidth guarantees. Managed bandwidth classes exist alongside strict priority classes. The ability to share bandwidth between priorities that carry bursty, high-offered loads, rather than servicing them with strict priority, is desirable. However, so is the ability to allow strict priority for time-sensitive or management traffic that requires minimum latency.

In any switch fabric, traffic from multiple ingress ports can compete for the limited bandwidth in a single egress port. For example, bursts of data traffic can reduce the bandwidth that is available to storage traffic, causing congestion and increased latency. To solve this problem, ETS provides advanced traffic scheduling. Advanced traffic scheduling includes features such as guaranteed minimum bandwidth for certain traffic classes, such as storage or high-performance computing.

Consider an example in which three classes of traffic are offered to a 10-Gb/s link. The guaranteed bandwidth is 30 percent for high-performance computing traffic, 40 percent for storage traffic, and 20 percent for LAN traffic.

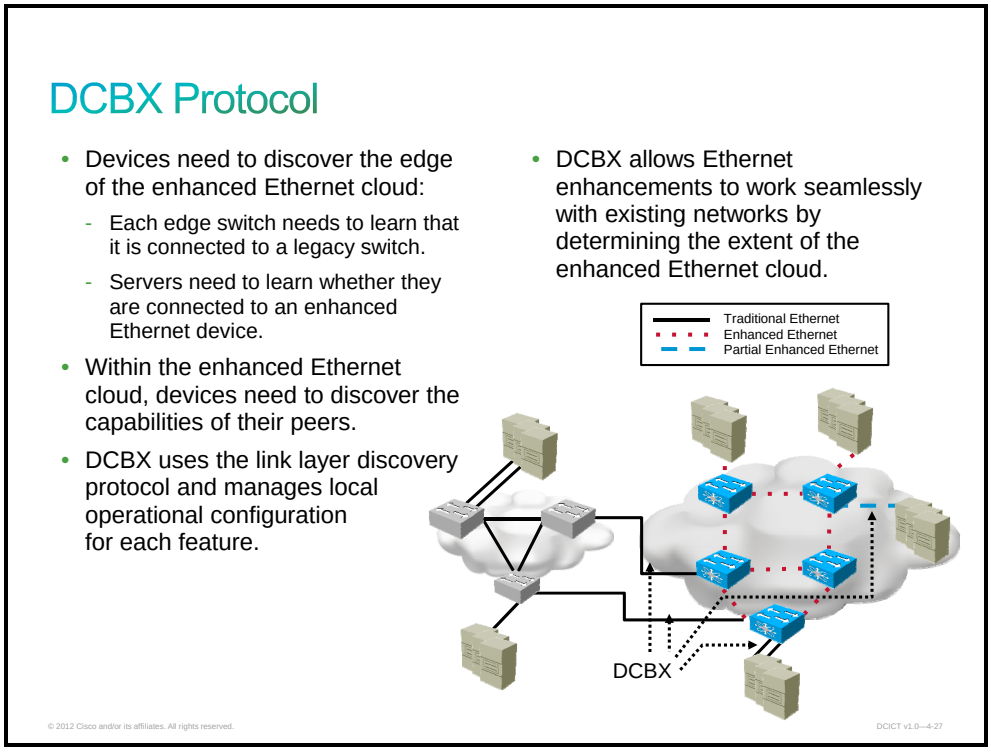
- At T1, all the traffic uses 3 Gb/s. At this point, each type of traffic uses its guaranteed bandwidth, and the link bandwidth is not yet fully utilized.
- At T2, high-performance computing and storage traffic remain the same, while LAN traffic increases to 3 Gb/s. At this point, link bandwidth is fully utilized.
- At T3, high-performance computing traffic drops to 2 Gb/s, storage remains at 4 Gb/s, and LAN traffic increases to 4 Gb/s. Because one class of traffic (high-performance computing) is not using all its reserved bandwidth, another class of traffic (LAN) can use the remaining available bandwidth.

The reserved bandwidth implementation is based on deficit weighted round robin (DWRR). Strict priority can also be configured for a CoS.

https://t.me/learningnets

DCB Exchange

DCBX is a discovery and capability exchange protocol that is used to discover peers and to exchange configuration information between DCBX-compliant bridges. This topic describes how DCBX allows capability queries between Fibre Channel over Ethernet Forwarders (FCoE Forwarders).



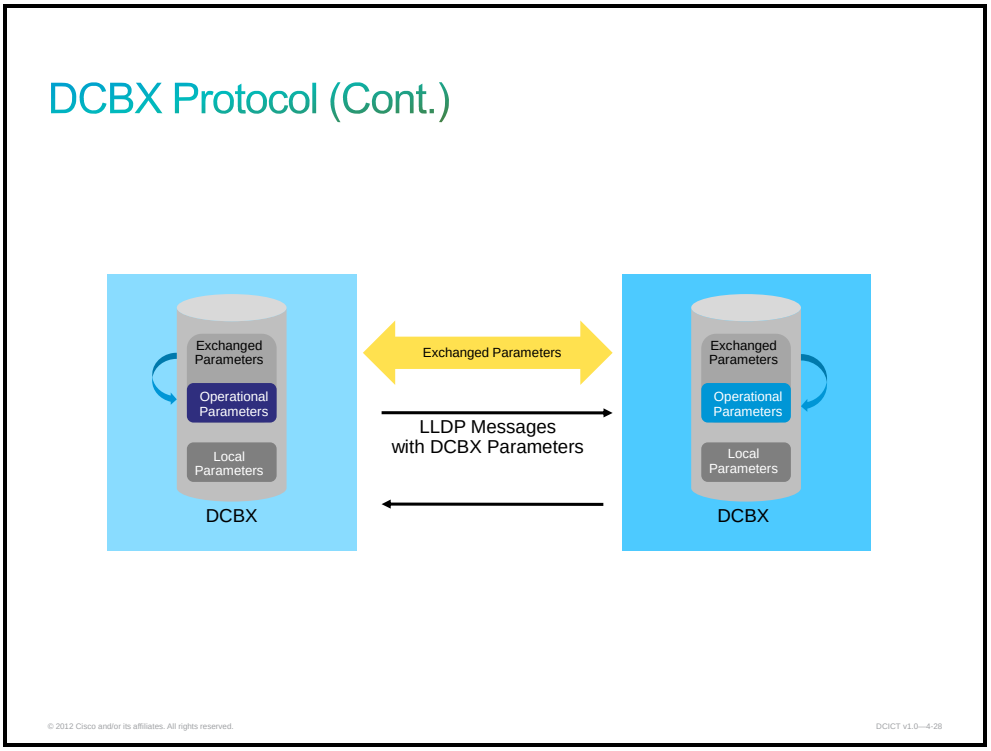
The enhancements that have been presented so far enable I/O consolidation on an Ethernet network. For these enhancements to work seamlessly with existing Ethernet deployments, a management protocol that dynamically discovers the capabilities of its peers is needed. The DCBX protocol provides the needed configuration and discovery functionality.

Consider the typical Ethernet network that the figure shows. Multiple devices of varying capabilities are connected to one another. In such a network, dynamic configuration management is required in two distinct scenarios. First, devices need to discover the edge of the enhanced Ethernet cloud. Each edge switch needs to learn that it is connected to a traditional switch, or whether it is connected to an enhanced Ethernet switch. Servers need to learn whether they are connected to an enhanced Ethernet device. Second, within the enhanced Ethernet cloud, devices need to discover the capabilities of their peers and the supported DCB features.

ETS incorporates DCBX, a discovery and initialization protocol that discovers the resources that are connected to the DCB cloud and establishes its limits. DCBX distributes the local configuration and detects the misconfiguration of ETS and PFC between peers. DCBX also provides the capability for configuring a remote peer with PFC, ETS, and other application parameters. The application parameter is used for informing the end station which priority to use for a given application type (such as FCoE or iSCSI). The DCBX protocol leverages the capabilities of IEEE 802.1AB LLDP.

DCBX uses LLDP 802.1AB-2005 and defines new type, length, values (TLVs) for capability-exchange parameters. The protocol state machines manage local operational configuration for each feature by comparing and synchronizing with the peer feature settings. The protocol provides the capabilities that the figure describes.

https://t.me/learningnets



Goals of DCBX

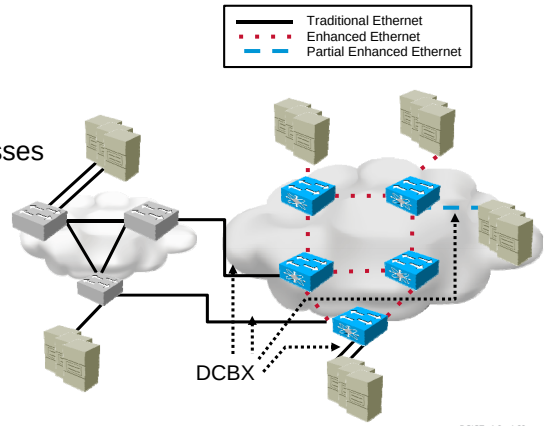
- **Discovery of DCB capabilities in a peer:** DCBX is used to identify the capabilities of peer devices. DCBX provides a mechanism by which a device can know if its peer device supports a particular feature, such as priority groups or PFC.
- **Misconfiguration detection:** DCBX can be used to detect misconfiguration of a feature between the peers on a link. The misconfiguration detection is feature-specific because some features permit asymmetric configuration.
- **Peer configuration:** A device can use DCBX to perform configuration of DCB features in its link peer. The initial purpose of DCBX was to provide basic peer-to-peer configuration. As DCBX evolves, higher-level applications can build on this basic purpose to provide more complex configuration-distribution mechanisms. Each DCB feature has a set of parameters. These parameters are classified into two broad categories:
  - **Exchanged parameters:** Exchange parameters are sent to the peer. Within these parameters, there are two subgroups:
    - **Administered:** These parameters are configured parameters.
    - **Operational:** These parameters provide the operational state of the related administered parameters. The operational state might be different from the administrative state, primarily as a result of the DCBX exchange with the peer. The operational parameters are included in the LLDP message for information purposes only.
  - **Local parameters:** Local parameters are not exchanged in LLDP messages.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

DCBX Protocol (Cont.)

LLDP PDUs carry values exchanged specific to the feature. DCBX provides hop-by-hop support for:

- Priority groups
- PFC
- ETS
- Logical link down
- Applications such as FCoE
- Information about traffic classes



© 2012 Cisco and/or its affiliates. All rights reserved.

The DCBX protocol is an extension of the LLDP. DCBX endpoints exchange request and acknowledgment messages. For flexibility, parameters are encoded in a TLV format.

Cisco Nexus switches support two versions of DCBX:

- **Converged Enhanced Ethernet DCBX (CEE-DCBX):** CEE-DCBX is supported on all T11-compliant second-generation CNAs.
- **Cisco, Intel, Nuova DCBX (CIN-DCBX):** CIN-DCBX is supported on first-generation CNAs and is used to perform link detection in addition to other functions.

DCBX runs on the physical Ethernet link between the Cisco Nexus switch DCB ports and between the Cisco Nexus switch DCB port and the CNA port. By default, DCBX is enabled on DCB Ethernet interfaces. When a DCB Ethernet interface is brought up, the switch automatically starts to communicate with the CNA.

During the normal operation of FCoE between the switch and the CNA, DCBX provides link-error detection to the CNA.

The CNAs that connect to a Cisco Nexus switch are programmed to accept the configuration values that the switch sends. This feature allows the switch to distribute configuration values to all attached CNAs, reducing the possibility of configuration errors and simplifying CNA administration.

https://t.me/learningnets

Summary

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

Summary

- A Unified Fabric has several features and benefits, such as I/O consolidation, in which a single physical link can be multiplexed to carry traffic from multiple fabrics such as the network and SAN.
- To support FCoE in a unified fabric environment, there are IEEE standards that should be implemented at a switch and adapter level to provide the correct levels of service to different types of traffic.
- PFC provides multiple channels through a single physical connection so that different levels of service can be applied to those channels to ensure a lossless fabric for FCoE.
- ETS provides a bandwidth-management mechanism so that each channel has access to a percentage of the bandwidth available on the physical pipe.
- DCBX is a low-layer protocol for discovering the unified fabric cloud and ensuring that both the switch and adapter are synchronized with regards to certain feature configurations.

© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Identifying Connectivity Options for FCoE on the Cisco Nexus 5000 Series Switch

Overview

The purpose of this lesson is to identify connectivity options for connecting the converged network adapter (CNA) to Cisco Nexus 5000 Series Switches and for connecting the switch to upstream LAN and SAN environments.

Objectives

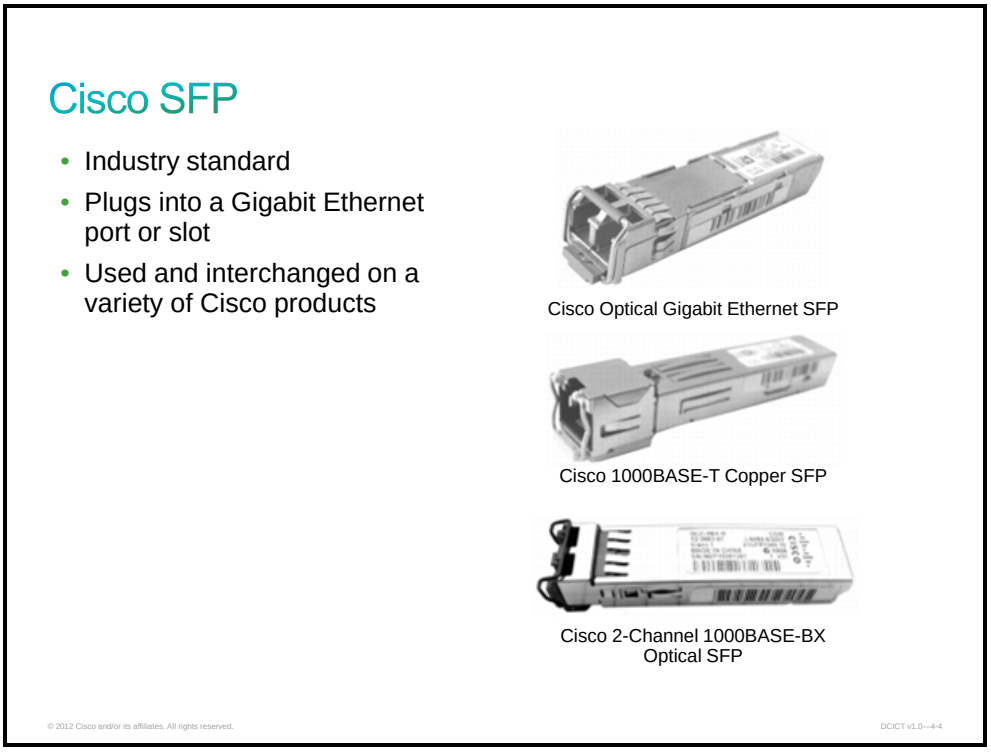
Upon completing this lesson, you will be able to identify options for connecting the CNA to a Cisco Nexus 5000 Series switch and for connecting the switch to upstream LAN and SAN environments. You will be able to meet these objectives:

- Describe SFP modules
- Describe cabling requirements and distance limitations for SFP and SFP+ transceivers
- Connect the Cisco UCS P81E Virtual Interface Card to the Cisco Nexus 5500UP unified fabric switch
- Connect the Cisco Nexus 5500UP unified fabric switch to northbound LAN and SAN fabrics

https://t.me/learningnets

SFP Modules

This topic describes the types of small form-factor pluggable (SFP) modules.



An SFP transceiver is a compact, hot-swappable, I/O transceiver that is used in data communication and telecommunications networks. SFP creates an interface between communication devices such as switches, routers and fiber optic cables and performs conversions between optical and electrical signals. SFP transceivers support communications standards such as SONET/SDH, Gigabit Ethernet, and Fiber Channel. The transceivers also allow the transport of Fast Ethernet and Gigabit Ethernet LAN packets over time-division multiplexing (TDM) WANs, as well as the transmission of E1 and T1 streams over packet-switched networks. SFP is also called a mini Gigabit Interface Converter (GBIC) because its function is similar to the GBIC transceiver, but with much smaller dimensions.

1000BASE-T SFP for Copper Networks

The 1000BASE-T SFP operates on standard Category 5 unshielded twisted-pair (UTP) copper cabling of up to 328 feet (100 m) link length. Cisco 1000BASE-T SFP modules support 10/100/1000 autonegotiation and automatic medium-dependent interface (auto-MDI) or MDI crossover (auto-MDIX).

1000BASE-SX SFP for Multimode Fiber

The 1000BASE-SX SFP, which is compatible with the IEEE 802.3z 1000BASE-SX standard, operates on traditional 50-micrometer multimode fiber links of up to 1800 feet (550 m). This SFP also operates on 62.5 micrometer FDDI-grade multimode fibers of up to 720 feet (220 m). The SFP can support up to 0.63 miles (1 km) of laser-optimized 50-micrometer multimode fiber cable.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

1000BASE-LX/LH SFP for Multimode and Single-Mode Fiber

The 1000BASE-LX/LH SFP is compatible with the IEEE 802.3z 1000BASE-LX standard. This SFP operates on standard single-mode fiber-optic link spans of up to 6.2 miles (10 km) and on multimode fibers of up to 1800 feet (550 m). When used over traditional multimode fiber, the transmitter should be coupled through a mode-conditioning patch cable. For details about this implementation, refer to *Use of Mode Conditioning Patch Cables in Gigabit Ethernet and 10 Gigabit Ethernet Laser-Based Transmissions* at http://www.cisco.com/en/US/prod/collateral/modules/ps5455/product_bulletin_c25-530836.html.

1000BASE-EX SFP for Long-Reach Single-Mode Fiber

The 1000BASE-EX SFP operates on standard single-mode fiber-optic link spans of up to 24.8 miles (40 km). Insert a 5-dB inline optical attenuator between the fiber-optic cable and the receiving port on the SFP at each end of the link, for back-to-back connectivity.

1000BASE-ZX SFP for Long-Reach Single-Mode Fiber

The 1000BASE-ZX SFP operates on standard single-mode fiber-optic link spans of up to approximately 43 miles (70 km). The SFP provides an optical link budget of 21 dB, but the precise link-span length depends on multiple factors, such as fiber quality, number of splices, and connectors.

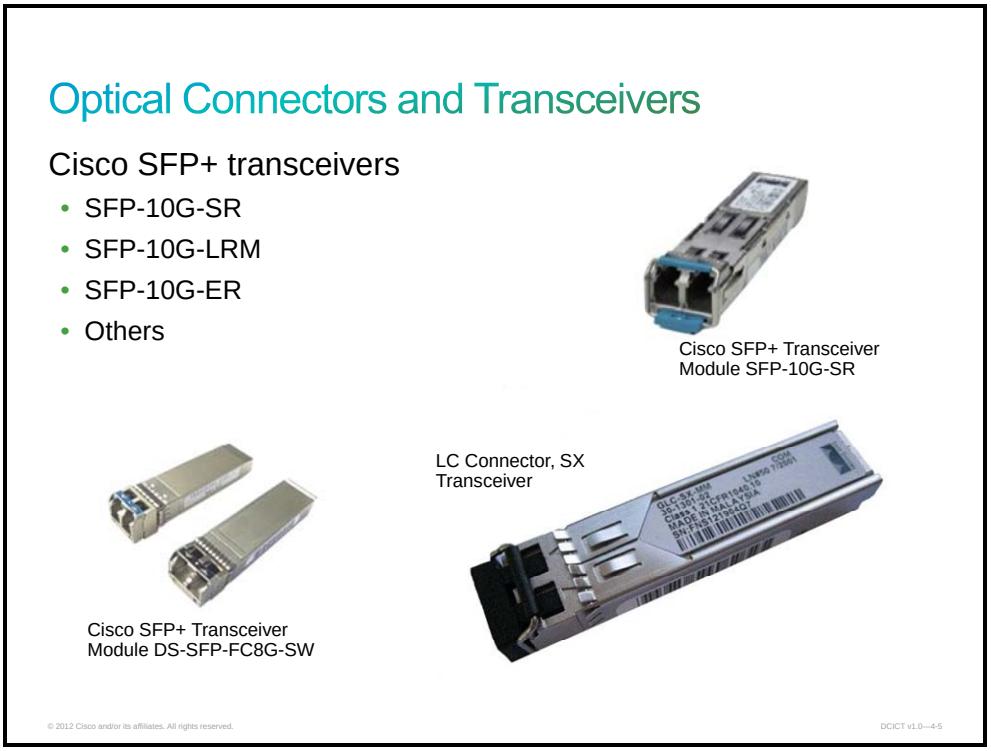
When shorter distances of single-mode fiber are used, you might need to insert an inline optical attenuator in the link, to avoid overloading the receiver. Insert a 10-dB inline optical attenuator between the fiber-optic cable plant and the receiving port on the SFP. Place these attenuators at each end of the link whenever the fiber-optic cable-span loss is less than 8 dB.

1000BASE-BX10-D and 1000BASE-BX10-U SFP for Single-Fiber Bidirectional Applications

The 1000BASE-BX-D and 1000BASE-BX-U SFPs, which are compatible with the IEEE 802.3ah 1000BASE-BX10-D and 1000BASE-BX10-U standards, operate on a single strand of standard single-mode fiber.

A 1000BASE-BX10-D device is always connected to a 1000BASE-BX10-U device with a single strand of standard single-mode fiber and has an operating transmission range of up to 6.2 miles (10 km).

https://t.me/learningnets



The SFP+ 10 Gigabit Ethernet transceiver module is a bidirectional device with a transmitter and receiver in the same physical package. This module has a 20-pin connector on the electrical interface and duplex Lucent Connector (LC) on the optical interface.

Cisco Nexus switches support these SFP+ optical transceivers:

- **10 GBASE-SR:** 98 to 980 feet (30 to 300 m) using OM1- or OM3-grade multimode fiber)
- **DCR, SR-Lite:** Shorter reach than short reach (SR)
- **10 GBASE-LR:** 6.2 miles (10 km) using single-mode fiber

In addition, Cisco provides LRM 10G SFP+ transceivers. However, as of Cisco NX-OS Software Release 5.0.2.N1.1, these transceivers are not supported on any Cisco Nexus 5000 Series, 5500 platform, or fabric extender (FEX) platforms. The Cisco Nexus 5596 Switch supports both SFP and SFP+ Ethernet transceivers, as well as SFP Fibre Channel transceivers.

Cisco Nexus switches support these SFP Fiber Channel transceivers:

- **DS-SFP-FC4G-SW:** 4-Gb/s Fibre Channel-SW SFP, LC
- **DS-SFP-FC4G-LW:** 4-Gb/s Fibre Channel-LW SFP, LC (6.2 miles [10 km] reach)

Cisco Nexus 2000 Series Fabric Extenders supports these SFP+ Ethernet transceivers:

- **SFP-10G-SR:** 10 Gigabit Ethernet (short-range SFP+ module)
- **SFP-10G-LR:** 10 Gigabit Ethernet (long-range SFP+ module)

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Twinax Cabling

- Passive cables in 1-, 3-, and 5-m lengths
- Active cables in 7- and 10-m lengths
  - Requires Cisco NX-OS Software Release 4.2.1
- Reduced cost and power use over optics
- Refer to the twinax cables certification matrix for Cisco Nexus Series switches



SFP+ Copper (Twinax) Cable



SFP+ Optic



© 2012 Cisco and/or its affiliates. All rights reserved.

Copper-based twinax cables are offered in 1-, 3-, 5-, 7- and 10-m lengths and are available for use with the 10 Gigabit Ethernet SFP+ module. Twinax cables are designed as top-of-rack (ToR) server interconnects on a per-rack basis, or as server connects between one to three racks in a middle-of-rack deployment. These cables can also be used between the FEX and the parent Cisco Nexus 5000 Series or Nexus 5500 platform switches.

Cisco NX-OS Software Release 4.2.1 for the Cisco Nexus 5000 Series and Nexus 5500 switches and in Cisco NX-OS Software Release 5.0.2 for the Cisco Nexus 5500 switches introduced support for these cables.

Twinax copper cables offer a reduced-cost alternative with reduced power consumption, as compared with fiber optic cabling.

For a complete and updated list, refer to the twinax cables certification matrix for Cisco Nexus 2000 Series Fabric Extenders and Cisco Nexus 5000 Series Switches in *Twinax Cables Certification Matrix for Cisco Nexus 2000 and Nexus 5000 Series Switches Data Sheet* at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps10110/data_sheet_c78-568589_ps9670_Products_Data_Sheet.html.

https://t.me/learningnets

Cisco FET-10G

- The Cisco FET-10G Fabric Extender Transceiver is supported only on fabric uplinks from a Cisco Nexus 2000 to the parent switch:
  - Not individually orderable
- Supports link lengths up to 100 m on laser-optimized OM3 multimode fiber, 25 m for OM2.
- Refer to the Cisco Nexus 2000 data sheet at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps10110/data_sheet_c78-507093.html.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco Fabric Extender Transceiver (FET-10G) is supported only on fabric uplinks from a Cisco Nexus 2000 Series fabric extender to the parent switch. This cable comes with the fabric extender and cannot be ordered individually. The FET-10G must be connected to another Cisco FET-10G.

For complete information about the Cisco FET-10G, refer to *Cisco Nexus 5548P, 5548UP, 5596UP, and 5596T Switches Data Sheet* at http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-618603.html.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cabling Requirements and Distance Limitations for Common SFP and SFP+ Transceivers

This topic describes transceiver requirements and distance limitations.

Transceiver Distances

Connector (Media)	Cable	Distance	Power (Each Side)	Transceiver Latency (Link)	Standard
SFP+ CU copper	Twinax	5m	Approx 0.1W	Approx 0 microsecond	SFF 8431
SFP+ ACU copper	Active Twinax	7m 10m	Approx 0.5W	Approx 0.1 microsecond	SFF 8461
SFP+ SR MMF and SR	MMF (OM2) MMF (OM3)	82m 300m	1W	Approx 0.1 microsecond	IEEE 802.3ae
SFP+ LR	SMF	10Km	1W	Approx 0.1 microsecond	IEEE 802.3ae
SFP+ ER	SMF	40Km	1.5W	Approx 0.1 microsecond	IEEE 802.3ae

The Cisco Nexus 5500 Platform supports SFP+ direct-attach 10 Gigabit Ethernet copper for in-rack cabling and optical solutions for longer connections (Ethernet only).

© 2012 Cisco and/or its affiliates. All rights reserved.

The high bandwidth of 10 Gigabit Ethernet poses challenges to transmissions. The transceiver and cabling options that the Cisco Nexus 5500 platform supports meet these challenges.

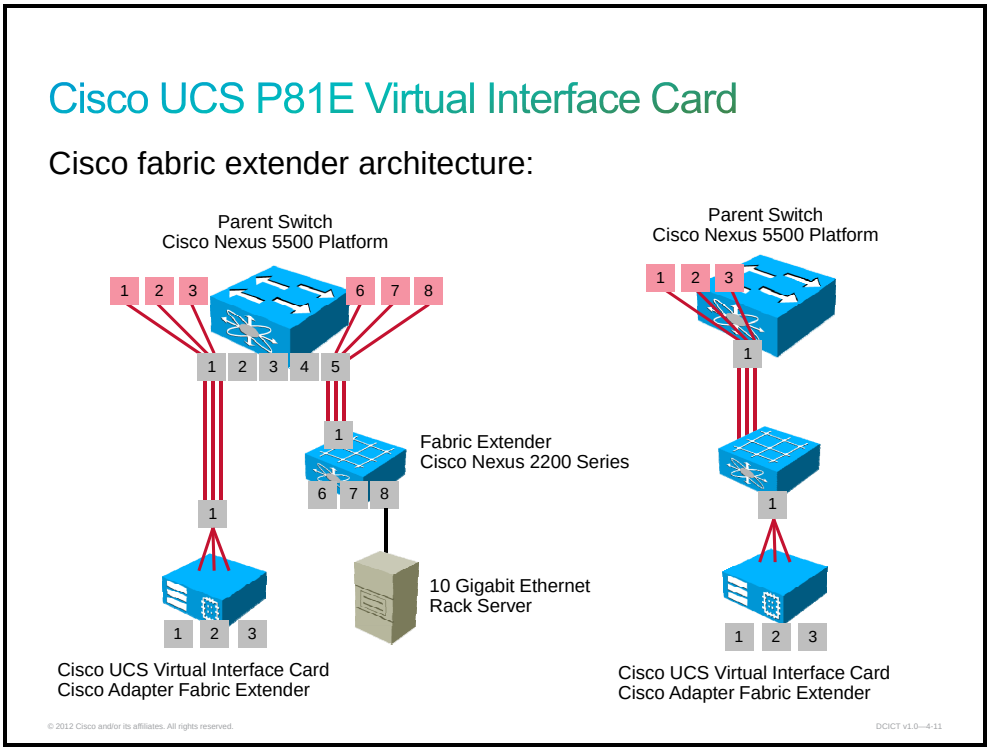
The platform supports an innovative twinax copper-cabling solution that connects to standard SFP+ connectors for in-rack use, as well as optical cabling for longer cable runs.

- For in-rack or adjacent-rack cabling, the Cisco Nexus 5500 platform supports SFP+ direct-attach 10 Gigabit Ethernet copper, a solution that integrates transceivers with twinax cables into an energy-efficient, low-cost, and low-latency solution. SFP+ direct-attach 10 Gigabit twinax copper cables use only 0.1 W of power per transceiver and have 0.25 ms of latency per link.
- For longer cable runs, the Cisco Nexus 5500 platform supports multimode, SR optical SFP+ transceivers. These optical transceivers use approximately 1 W per transceiver and have a latency of less than 0.1 ms.

https://t.me/learningnets

Connecting the Cisco UCS P81E Virtual Interface Card to Cisco Nexus 5500UP Unified Fabric Switches

This topic describes transceiver options for the Cisco Nexus P81E Virtual Interface Card.



The Cisco Fabric Extender Architecture (FEX Architecture) enables these capabilities:

- **Architecture flexibility:** The common, scalable, and adaptive architecture across data center racks and points of delivery (PoDs) is server and server-adapter agnostic. The architecture supports several connectivity options, physical topologies, and evolving needs.
- **Highly scalable server access:** Scalable Gigabit and 10 Gigabit Ethernet server access is offered, with no reliance on Spanning Tree Protocol (STP).
- **Simplified operations:** One point of management and policy enforcement, using upstream Cisco Nexus switches, eases the commissioning and decommissioning of server racks through zero-touch installation and automatic configuration of FEXs.
- **Increased business benefits:** The extremely cost-effective in-rack cabling solution offers these benefits, which contribute to reduced operating expenses (OpEx) and capital expenditures (CapEx):
  - Consolidation
  - Rack-space reduction
  - Reduced power and cooling needs
  - Investment protection through feature inheritance from the parent switch and the capability to add functions without the need for a major equipment upgrade of server-attached infrastructure

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

- **Open standards-based implementation:** The Cisco FEX Architecture supports the IEEE 802.1Qbh standard. The introduction of the Cisco Adapter FEX provides the same architectural benefits for the network interface card (NIC) as are provided for the physical access layer. Adapter FEX is logically an extension of the parent switch inside the server. FEX on the adapter is used with Adapter FEX and the Cisco Virtual Machine Fabric Extender (VM-FEX) feature.

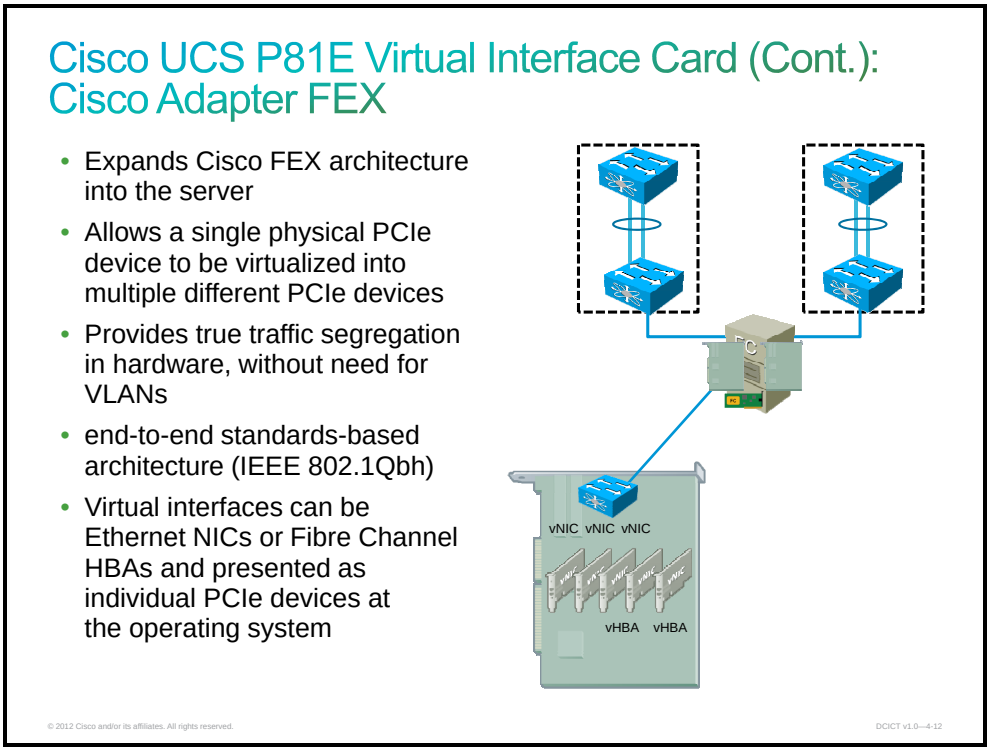
The host interfaces of FEX on the adapter are local logical ports on the management node or on the parent switch. FEX on the adapter uses server connectivity (I/O connectivity) technology that enables on-demand creation of virtual NICs (VNICs) or virtual host bus adapters (vHBAs) on one NIC. A Cisco virtual interface card (VIC) enables the physical Peripheral Component Interconnect (PCI) bus to be presented as multiple virtual PCI busses to the operating system. With Adapter FEX, a single physical adapter port is presented as multiple logical adapter ports (vHBAs and VNICs) to the server operating system and the network. It appears as if multiple physical adapter ports are available to the operating system. A dual-port 10 Gigabit Ethernet Adapter FEX can support hundreds of PCI Express (PCIe) standards-compliant virtual interfaces, which can be configured dynamically by the server administrator or by the Cisco process.

Each VNIC and vHBA that is created on the physical adapter automatically corresponds to a virtual Ethernet interface or virtual Fibre Channel interface on the parent switch to which it is mapped. The FEX on the adapter is connected to the Data Center Bridging (DCB) port. The network administrator then assigns network properties to each logical interface, to help guarantee advanced QoS and granular bandwidth allocation.

Adapter FEX and VM-FEX technology extends the current benefits of the Cisco Nexus 2000 Series fabric extender architecture to the server NICs, providing architecture flexibility, high scalability with 4000 logical interfaces per management node, and one point of management and policy enforcement, and resulting in increased business benefits.

In the Cisco Unified Computing System (UCS) B-Series Blade Servers, the Cisco UCS M81KR Virtual Interface Card is the first product that implements Adapter FEX technology.

https://t.me/learningnets



Static Interface and Static vNIC or vHBA

A static interface is an interface with parameters that the administrator configures manually. A static virtual adapter can be a vNIC or a vHBA. A static interface can be a virtual Ethernet (vEth) or virtual Fibre Channel (vFC) interface.

For a statically created fixed and floating vEth, a network administrator can associate configuration to the vEth before it is brought up. When creating a static vEth, the network administrator specifies which channel (for simplicity, you can consider this equivalent to a virtual networking tag [VNTag] number) a given vEth uses. The server administrator must define a vNIC on the adapter with the same channel number.

Dynamic Interface and Dynamic vNIC

A dynamic interface is a vEth interface that is configured automatically as a result of adapter and switch communications. The provisioning model of a dynamic interface consists in the configuration of a port profile on the Cisco Nexus 5500 switch, which appears on the network adapter, followed by the server administrator associating the port profile with the vNIC.

A dynamic virtual adapter can be a vNIC but not a vHBA. A dynamic virtual interface can be a vEth interface. Dynamic interfaces have support for hardware-based failover in the Cisco UCS P81E Virtual Interface Card.

Fixed vEth Interface

A fixed vEth interface is a virtual interface that does not support migration across physical interfaces. The scope of Adapter FEX is always on fixed vEth because Adapter FEX refers to the use of network virtualization by one nonvirtualized operating system.

For fixed vEth (static or dynamic), administrators can change configurations, including shut/no shut or create/delete, at any time. The vEth number-to-channel number binding is persistent unless the administrator changes it.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

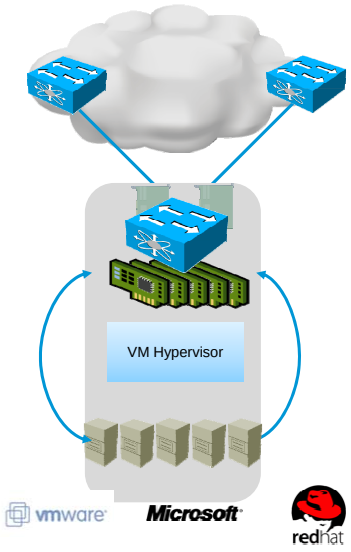
Floating vEth Interface

When the Cisco UCS P81E network adapter is used in a hypervisor environment, each vNIC on the adapter is associated with one virtual machine (VM). VMs can migrate from one physical server to another. A virtual interface that migrates along with a VM and virtual network link is called a floating virtual interface.

For a floating (static or dynamic) vEth, a configuration change (including shut/no shut) is allowed at any time, regardless of the attached state—except for a binding configuration. Changing the binding configuration is not allowed while a vEth is attached. Binding configurations can be changed only if no vEth is associated.

Cisco UCS P81E Virtual Interface Card (Cont.): Cisco VM-FEX

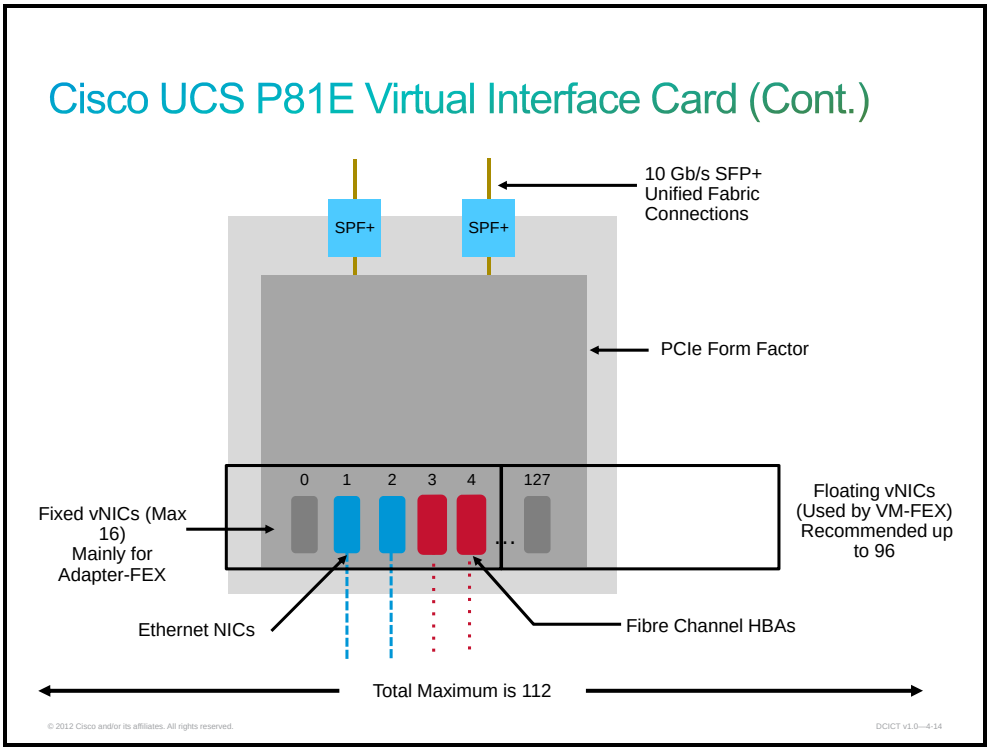
- Further extends Cisco Adapter FEX technology to the VM
- VM awareness, with policy enforcement and management
- Consolidated management of network, server, and virtual server interfaces
- Offload I/O processing from server CPU, with Cisco Nexus hardware performance



© 2012 Cisco and/or its affiliates. All rights reserved.

Adapter FEX can be extended to the VM level and can support offloading of the I/O processing from the server CPU with Cisco Nexus hardware performance.

https://t.me/learningnets



Expanding this capability outside the Cisco UCS B-Series Blade Servers, the Cisco Nexus 5500 platform can support adapters that implement Adapter FEX and VM-FEX technology. An ecosystem of adapter vendors will soon support this technology using the IEEE 802.1Qbh standard. Cisco is the first to introduce this protocol idea in the Cisco UCS P81E Virtual Interface Card, which is designed for use with Cisco UCS C-Series Rack-Mount Servers. Other adapter vendors are following, providing adapters that support this capability. Cisco Nexus 5500 platform and Cisco Nexus 2000 Series Fabric Extenders support the connection to the Cisco VIC with Adapter FEX and VM-FEX technology, across various adapter platforms such as those from Broadcom and Intel. Therefore, the offering can now be expanded outside of a Cisco UCS environment, to third-party server vendors that will support IEEE 802.1Qbh-capable adapters.

Adapter FEX technology provides exceptional scalability and flexibility for both virtualized and nonvirtualized environments, enabling on-demand, cost-effective solutions for data center server connectivity.

Features and Benefits

The Cisco FEX on the adapter that is enabled with Adapter FEX or VM-FEX technology addresses these main challenges:

- Organizations are deploying virtualized workloads to meet strong needs to save costs and reduce physical equipment. However, virtualization technologies and the increased number of CPU cores require servers to be equipped with many network connections. For example, a typical VMware server has six 1 Gigabit Ethernet ports and two 4-Gb/s Fibre Channel ports. These ports correspond to the eight cables that are needed to connect to the eight upstream network ports. This increase has a tremendous impact on CapEx and OpEx because the many adapters, cables, and switch ports directly affect power, cooling, and cable-management costs.
- Network administrators struggle to link the virtualized network infrastructure and virtualized server platforms to the physical network infrastructure without degrading performance and with a consistent feature set and operational model.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

- With a consolidated infrastructure, providing guaranteed bandwidth, latency, and isolation of traffic across multiple cores and VMs is challenging.
- Data center designs call for efficient cabling and reduced power and cooling because of stringent budgetary constraints.
- In virtualized environments, network administrators experience lack of visibility into the traffic that is exchanged among VMs that belong to the same host. Administrators also face challenges in establishing and enforcing policies and maintaining configurations and policies consistently across mobility events. They often also see a dramatic increase in the number of management points; disparate provisioning, management, and operational models; and inconsistency between the physical and virtual access layers.

Adapter FEX provides these benefits:

- Flexible and efficient deployment of static vNICs for nonvirtualized environments
- Scalability for traditional virtualized environments, using virtual embedded bridges or virtual Ethernet switches (vSwitches)
- Highly optimized VM connectivity for virtualized environments

Adapter FEX technology can be used in virtualized server environments, to make the network infrastructure VM-aware. A vNIC that is created in this context is referred to as a Cisco VM-FEX vNIC. This vNIC integrates with the server-virtualization management tool, allowing the user to bind a VM to a dynamic vNIC. This binding makes it possible to use an external hardware switch (such as a Cisco Nexus 5500 switch) to switch the VM traffic, creating a single point of management and policy enforcement on the switch. That management capability enables migration support for VMs, floating vNIC Cisco VN-Link interfaces, and floating vEth interfaces. Port profile consistency is enabled from one host and DCB port on the management node to the other host and DCB port on the same switch (or virtual port channel [vPC]-connected switch). Port profile consistency provides tools that have the same level of visibility, security, and troubleshooting for VMs as customers are accustomed to using for physical devices.

The portfolio of Cisco VM networking products provides various options that meet these needs:

- Hypervisor switching with Cisco Nexus 1000V Series Switches
- Hardware switching (VN-Link in hardware pass-through switching) with VM-FEX and Cisco Nexus 5500 platform
- Hardware switching (VN-Link in hardware pass-through switching) with VM-FEX and Cisco UCS

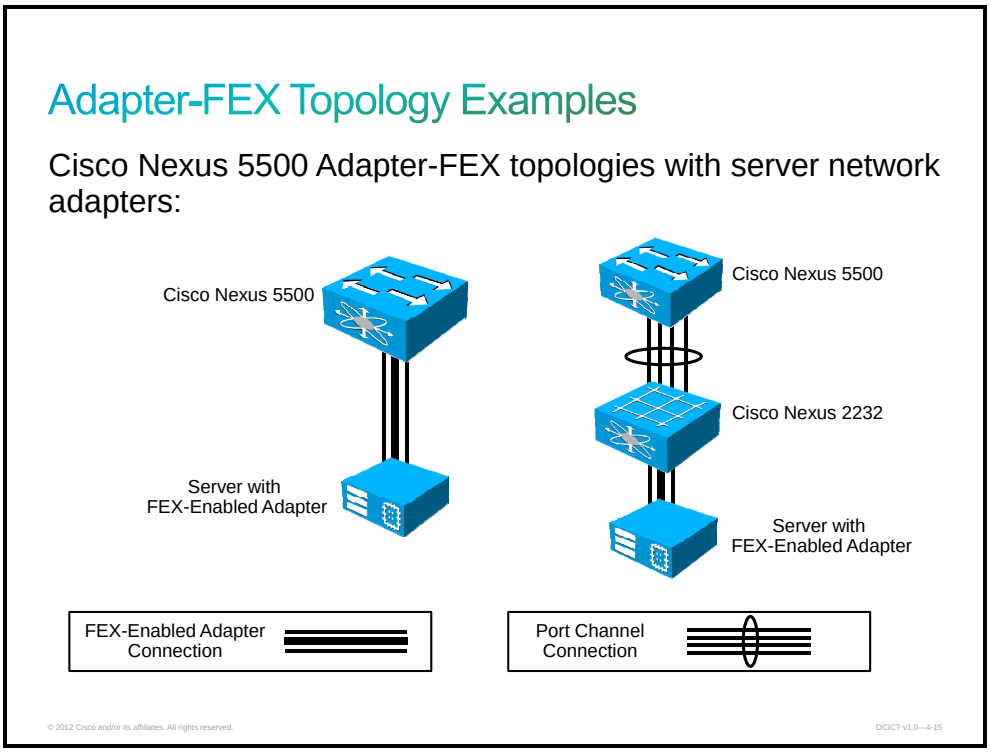
Across all these solutions, VM networking enables these capabilities:

- Policy-based VM networking
- Transparent network and security policy mobility with VM migration
- Nondisruptive operational model, with the network administrator using a consistent set of tools to manage both virtual and physical networking resources

https://t.me/learningnets

Adapter FEX technology provides these main business benefits to IT departments:

- **On-demand design and deployment of data center applications to enable cloud deployments, with reuse of existing equipment:** The wire-once model enables subsequent deployment of unified fabric and virtualization technologies.
- **Exceptional scalability:** 4000 logical host-facing ports are managed through one point of configuration.
- **Infrastructure efficiency through consolidation:** The network is simplified by reducing the number of adapters, cables, and network ports. The consolidation also reduces the number of network devices and the management overhead, thus lowering CapEx and OpEx.
- **Ease and consistency of management:** The server administrator can independently configure the adapter to a certain number of logical NICs. This configuration allows for repurposing of the physical NIC in real time as application needs evolve. These changes can be made with little impact on the network and storage teams. At the same time, the network team can preconfigure the advanced network configuration, lowering the overall management overhead. Both teams continue to use traditional management tools: BIOS- or operating system-based management tools on the server side and the CLI on the network side.



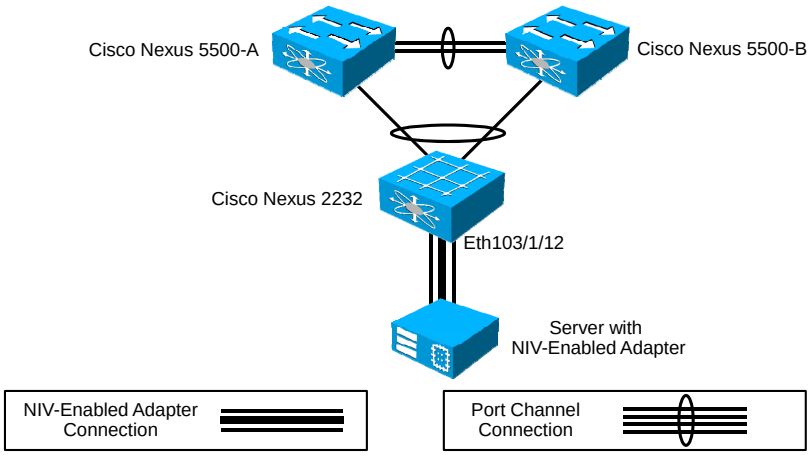
Adapter FEX requires a server network adapter that is connected to a parent switch that supports Adapter FEX functionality. Adapter FEX support is available with Cisco Nexus 5500 switches and Cisco Nexus 2200 fabric extenders and with the Cisco UCS systems that connect to a Cisco Nexus 5500 parent switch.

This implementation is designed to work with server network adapters, such as the Cisco UCS P81E Virtual Interface Card for the Cisco UCS C-Series Rack-Mount Server or Cisco blade servers or third-party VIC adapters that support the virtual network tag (VNTag) and Cisco FEX technology, such as the Broadcom BCM57712 Convergence Network Interface Card or an Intel VIC.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Adapter-FEX Topology Examples (Cont.)

Cisco Nexus 5500 Adapter-FEX topologies with server network adapters:

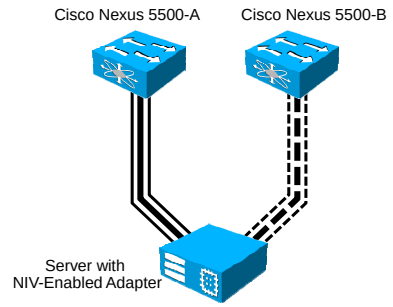


© 2012 Cisco and/or its affiliates. All rights reserved.

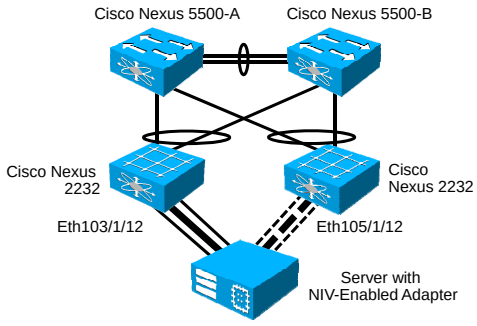
This figure depicts the third topology example for Cisco Nexus 5500 Adapter FEX implementations with server network adapters.

Adapter-FEX Topology Examples (Cont.)

Cisco Nexus 5500 Adapter-FEX active-standby topology:



Cisco Nexus 5500 Adapter-FEX dual-homed active-standby topology:





© 2012 Cisco and/or its affiliates. All rights reserved.

This figure shows topologies that are supported for server network adapters that support active/standby teaming of uplinks. The active/standby topologies that the figure shows have one active uplink and one standby uplink. With some server network adapters, you can select the active and standby uplinks per vNIC. In this case, each uplink is an active uplink for a specific vNIC and becomes a standby for the remaining uplinks.

https://t.me/learningnets

Selecting the active and standby uplinks per vNIC is a best practice.

When Adapter FEX is dual-homed, the vPC capability lets you synchronize the vEth MAC and Internet Group Management Protocol (IGMP) groups from active to standby. The vPC also sends gratuitous MACs and multicast registrations for vEth interfaces that fail over.

Note	The Cisco UCS P81E Virtual Interface Card supports active/standby uplink.
Note	The Cisco UCS P81E Virtual Interface Card allows each vNIC to choose the active uplink. The other uplink is configured as a standby uplink.
Note	For details, see the Cisco UCS P81E Virtual Interface Card documentation in the *Cisco UCS C-Series Servers Integrated Management Controller CLI Configuration Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/sw/cli/config/guide/131/b_Cisco_UCS_C-Series_CLI_Configuration_Guide_131.html.

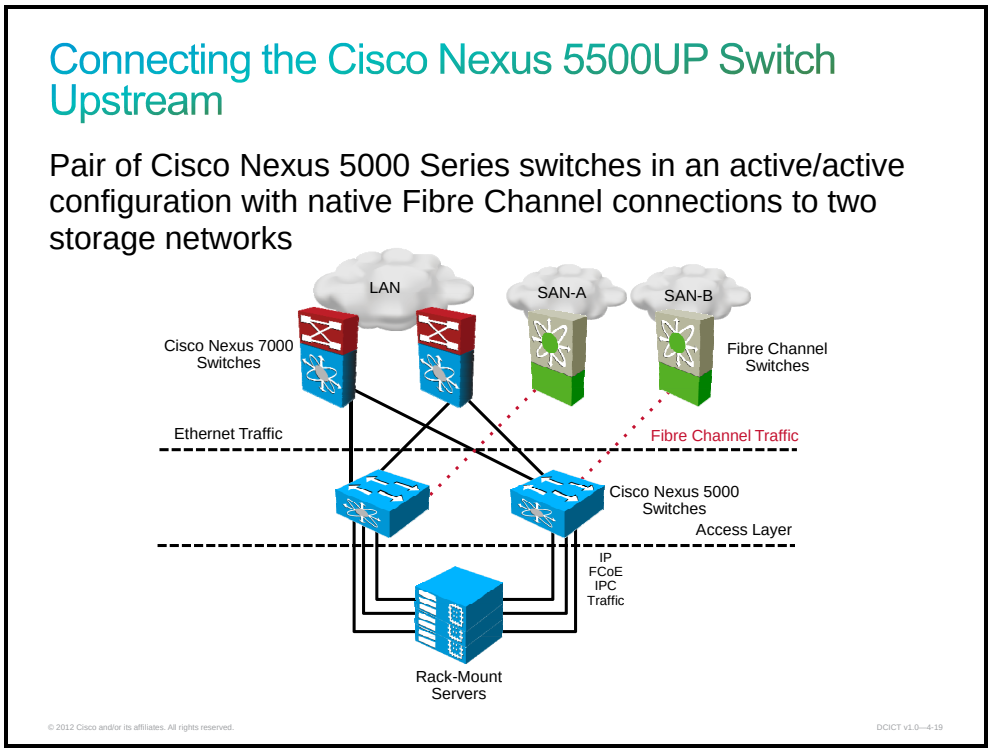
If NIC teaming is required, it is configured at the host operating-system level and requires two vNICs at the Cisco VIC level.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Connecting the Cisco Nexus 5500UP Unified Fabric Switch to Northbound LAN and SAN Fabrics

This topic describes northbound LAN connections to the Layer 3 LAN domain and the Fibre Channel SAN edge switch.



Unified Fabric with FCoE: I/O Consolidation

The Cisco Nexus 5500 platform consolidates multiple networks—LAN, SAN, and interprocess communication (IPC) cluster—on one unified fabric, saving the CapEx and OpEx that are associated with deployment of multiple parallel networks, switching infrastructure, and cabling. The Cisco Nexus 5500 platform is compatible with third-party I/O CNAs that present one or more separate Ethernet NICs and Fibre Channel HBAs to the server operating system. This approach allows existing drivers and Fibre Channel management software to work transparently with Fibre Channel over Ethernet (FCoE). Upstream, expansion modules support direct connections from the Cisco Nexus 5500 platform to existing native Fibre Channel SANs by using dedicated native Fibre Channel ports. The Cisco Nexus 5500 platform can also be connected to upstream FCoE-capable switches through industry-standard virtual expansion port (VE Port) interswitch links (ISLs).

The Cisco Nexus 5500 platform supports consolidated I/O by using FCoE on downlinks to servers. FCoE is Fibre Channel and so has familiar methods of configuration, operation, and management. On the uplinks to the network, customers can choose among these options:

- FCoE direct attachment to storage systems
- Native Fibre Channel direct attachment to storage systems
- DCBand FCoE connection to DCB- and FCoE-capable switches
- Native Fibre Channel connection to existing SANs
- Native Ethernet connection to an existing LAN

https://t.me/learningnets

Fibre Channel and FCoE can exist simultaneously in the Fibre Channel network, allowing customers to continue to use their existing administration and management tools. Cisco Fabric Manager SAN management software manages Fibre Channel and FCoE flows over both FCoE- and Fibre Channel-capable ports.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

Summary

- An SFP is a compact I/O transceiver used in data communications and telecommunications networks, providing connectivity between devices and conversions between optical and electrical signals.
- Transceivers support different distances within the data center. For short-reach communication Twinax copper cabling solutions would be used, and for longer-reach communication, optical cabling solutions.
- The Cisco UCS P81E Virtual Interface Card connects the server to a Cisco Nexus 5500 switch and provides multiple virtual adapters, which connect to virtual interfaces on the switch.
- The Cisco UCS P81E Virtual Interface Card supports I/O consolidation and FCoE. The upstream Cisco Nexus 5500 switch can split the traffic coming from the adapter by providing connectivity to both the LAN and SAN infrastructures.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 3

Describing Enhanced FCoE Scalability with Cisco Nexus 2232 10GE Fabric Extenders

Overview

The purpose of this lesson is to describe the role of Network Interface Virtualization (NIV) with the Cisco Nexus 2232 10GE fabric extender in Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) deployments.

Objectives

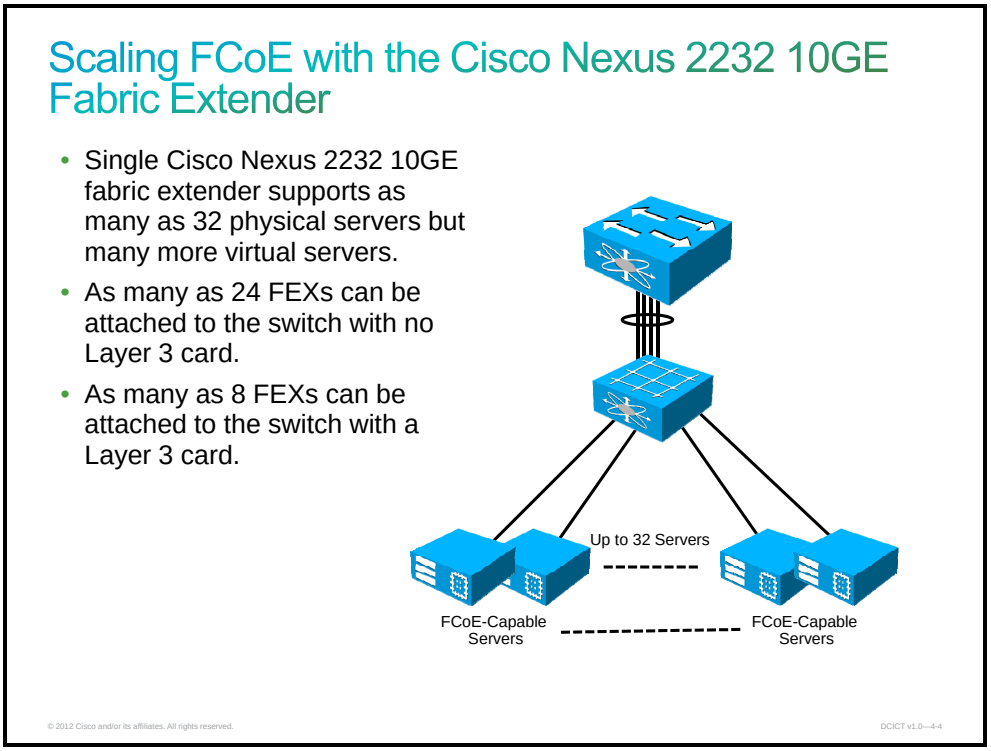
Upon completing this lesson, you will be able to describe the role of NIV with the Cisco Nexus 2232 10GE Fabric Extender in FCoE deployments. You will be able to meet these objectives:

- Describe how the Cisco Nexus 2232 10GE fabric extender increases the scalability of FCoE deployments
- Describe Cisco Nexus 2232 10GE fabric extender connectivity to the Cisco Nexus 5500 Switch
- Describe Adapter FEX on the Cisco Nexus 2232 10GE fabric extender
- Verify Adapter FEX on the Cisco Nexus 2232 10GE fabric extender

https://t.me/learningnets

Scaling the Data Center Virtualized Access Layer with the Cisco Nexus 2232 10GE Fabric Extenders

Cisco Nexus 2232 10GE fabric extenders have 32 10-Gb/s DCB interfaces, all of which are FCoE and DCB and FCoE aware. This topic describes how Cisco Nexus 2232 10GE fabric extenders extend the FCoE capacity of Cisco Nexus 5500 switches.



The Cisco Nexus 2232PP 10GE Fabric Extender supports FCoE and a set of network technologies that are known collectively as DCB. (FCoE is included in the DCB technologies.). DCB increases the reliability, efficiency, and scalability of Ethernet networks. These features allow the switches to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN, and cluster environments.

The Cisco Nexus 2232TM 10GE Fabric Extender supports scalable 1/10GBASE-T environments, ease of migration from 1GBASE-T to 10GBASE-T, and effective reuse of existing structured cabling. This fabric extender (FEX) comes with an uplink module that supports eight 10 Gigabit Ethernet fabric interfaces. This FEX supports DCB.

The combination of Cisco Nexus 5000 Series Switches and Cisco Nexus 2232PP 10GE Fabric Extender provides a unified network fabric that supports LAN and SAN consolidation. Another benefit of the Cisco Nexus 2000 Series architecture is the ability to collapse data center access and aggregation layers into one layer.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

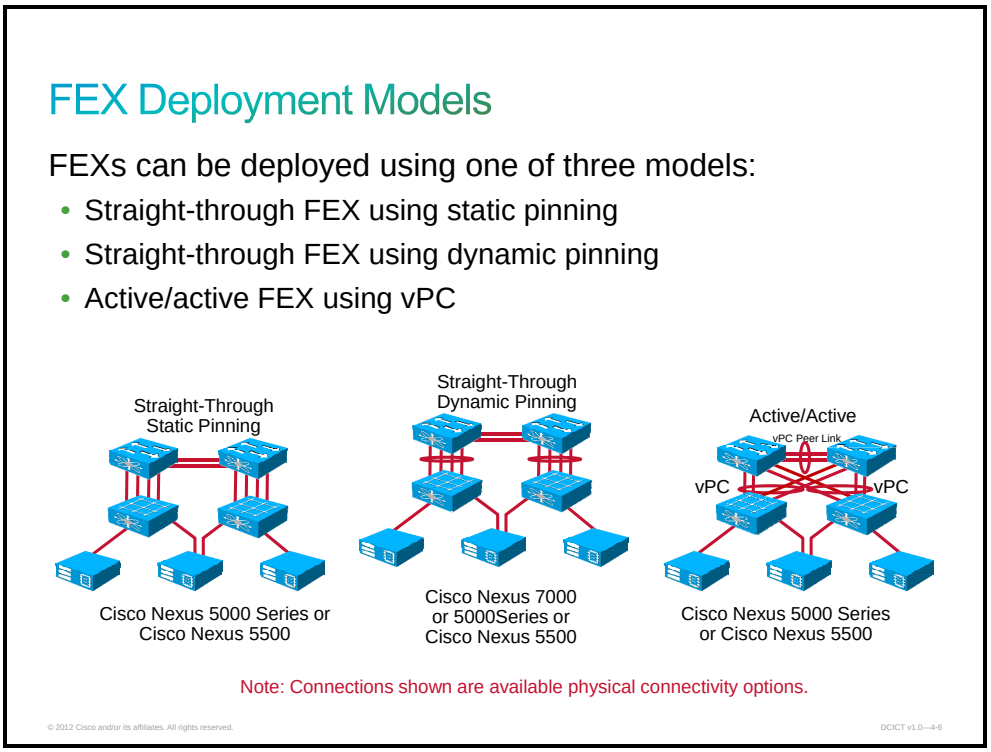
Cisco Nexus 2000 Series Fabric Extenders support an optimal cabling strategy that simplifies network operations and prepares for future technologies:

- **Short intra-rack runs of copper:** Twinax cables connect servers to top-of-rack (ToR) Cisco Nexus 2232PP 10GE fabric extenders. This model allows server vendors to preconfigure server racks and points of delivery (PoDs) so that they can be rolled into place and put into service upon arrival.
- **Longer inter-rack horizontal runs of fiber:** Cisco Nexus 2000 Series Fabric Extenders in each rack are connected to parent switches that are placed at the end or middle of the row: For long reach between the FEX and the parent switch, Cisco Fabric Extender Transceiver (FET), enhanced small form-factor pluggable (SFP+) short-reach (SR), and SFP+ long-reach (LR) optics over OM2 or OM3 cables can be used. Fiber, which supports upcoming Ethernet standards including 40 and 100 Gigabit Ethernet, protects investments into the future. If the distance to the Cisco Nexus 5000 Series switch or Cisco Nexus 7000 Series switch is less than 33 feet (10 m), then twinax (direct-attach CX1) cables can be used. Alternatively, copper cables can be used to reach end-of-row or middle-of-row 1GBASE-T or 10GBASE-T FEXs).

https://t.me/learningnets

Cisco Nexus 2232 10GE Fabric Extender-to-Cisco Nexus 5500 Switch Connectivity

This topic describes connectivity options from the Cisco Nexus 2232 10GE fabric extender to a Cisco Nexus 5500 switch.

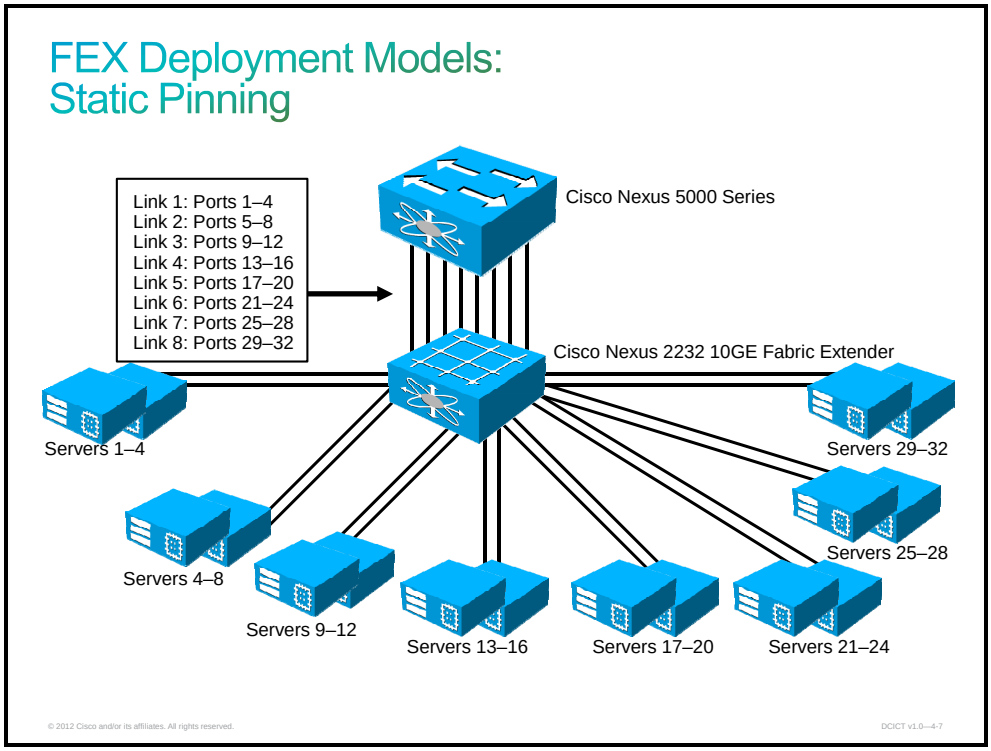


Three deployment models are used to deploy the Cisco Nexus 2232 10GE fabric extender with Cisco Nexus 5500 switches:

- **Straight-through using static pinning:** In the straight-through model, each FEX is connected to one Cisco Nexus switch. This single switch exclusively manages the ports on that FEX. Static pinning means that each downlink server port on the FEX is statically pinned to an uplink between the FEX and the switch. Traffic to and from a specific server port always uses the same uplink.
- **Straight-through using dynamic pinning:** This deployment model also uses the straight-through connection model between FEXs and switches. However, there is no static relation between the downlink server ports and the uplink ports. The port between the FEX and the switch are bundled into a port channel, and traffic is distributed across the uplinks, based on the PortChannel hashing mechanism.
- **Active/active FEX using virtual port channel (vPC):** In this deployment model, the FEX is dual-homed to two Cisco Nexus switches. vPC is used on the link between the FEX and the pair of switches. Traffic is forwarded between the FEX and the switches, based on vPC forwarding mechanisms.

Note	Cisco Nexus 7000 Series Switches currently support only straight-through deployment using dynamic pinning. Static pinning and active/active FEX are currently supported only on Cisco Nexus 5000 Series Switches.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



In static-pinning mode, the server ports on the Cisco Nexus 2000 Series Fabric Extenders are statically pinned to an uplink port. For example, when a Cisco Nexus 2000 Series Fabric Extender with 32 10 Gigabit Ethernet server ports is deployed in static-pinning mode (using 8 10 Gigabit Ethernet uplink ports) to a Cisco Nexus 5000 Series switch, 4 server ports are pinned to each uplink port:

- Ports 1 to 4 are pinned to the first uplink.
- Ports 5 to 8 to the second uplink.
- Ports 9 to 12 to the third uplink.
- Ports 13 to 16 to the fourth uplink.
- Ports 17 to 20 to the fifth uplink.
- Ports 21 to 24 to the sixth uplink.
- Ports 25 to 28 to the seventh uplink.
- Ports 29 to 32 to the eighth uplink.

This configuration results in a four-to-one oversubscription ratio. Each group of four 10 Gigabit Ethernet server ports share the bandwidth of one 10 Gigabit DCB uplink.

If an uplink between the Cisco Nexus 2000 Series fabric extender and the Cisco Nexus 5000 Series switch fails, then the FEX disables the physical server ports that are pinned to that uplink port. For example, if the fourth uplink fails, then the FEX disables server ports 13 to 16. Servers that are connected to these ports see the associated Ethernet link fail. If the servers are dual-homed or single-homed and use some form of network interface card (NIC) redundancy, then this mechanism is triggered. The server then fails over to the other NIC. A single-homed server simply loses connectivity if it is connected to a port that is pinned to the failed uplink port. The oversubscription ratio on the other ports remains unchanged because each of the other seven groups of four 10 Gigabit Ethernet ports still shares the same 10 Gigabit DCB uplink port as before.

https://t.me/learningnets

FEX Deployment Models: Dynamic Pinning

Dynamic pinning uses a port channel between the Cisco Nexus switch and the FEX.

- Traffic distribution across the uplinks is determined through port-channel hashing.
- When an uplink fails, traffic is rehashed onto the remaining links:
  - Server downlinks are not disabled when a single uplink fails.
  - Oversubscription ratio between FEX and switch changes on uplink failure.
  - Single-homed servers retain connectivity.
  - Dual-homed servers do not fail over to the other NIC.

© 2012 Cisco and/or its affiliates. All rights reserved.

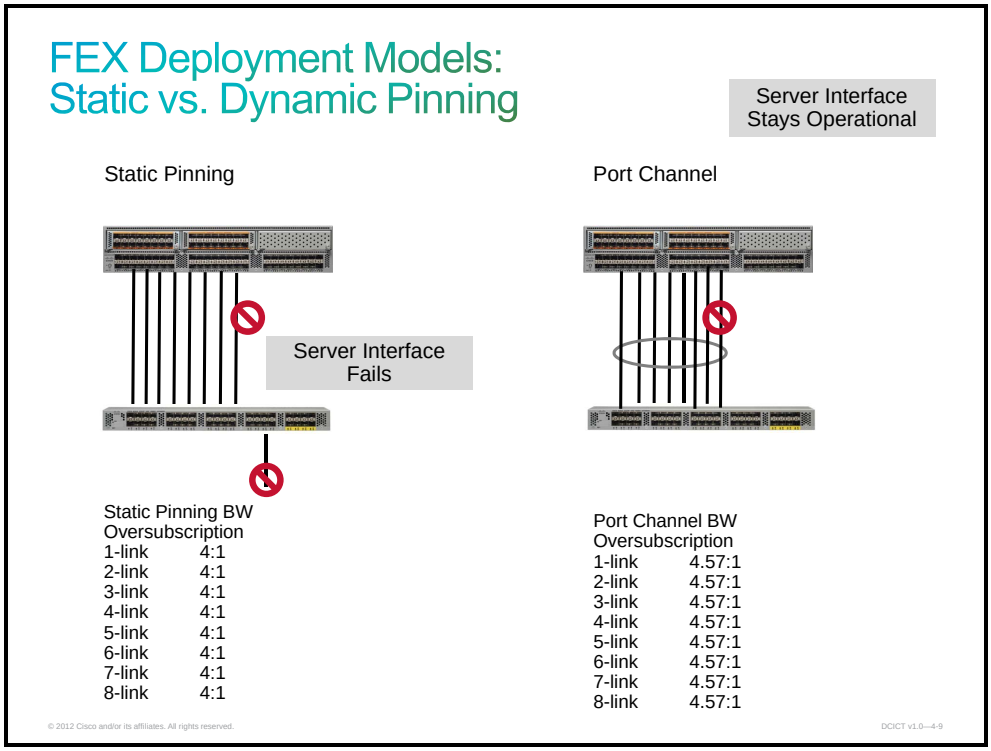
In dynamic-pinning mode, the server ports are not statically pinned to an uplink port. All server ports share the combined bandwidth of the uplink ports within the port channel. Hosts see the port channel as one logical connection. This mode is achieved by configuring a port channel between a Cisco Nexus 2000 Series fabric extender and a Cisco Nexus 5000 Series switch. Traffic is distributed over the port-channel members, based on the PortChannel load-balancing hash algorithm.

However, the behavior under failure conditions differs from the behavior when using static-pinning mode. If a physical uplink port between a Cisco Nexus 2000 Series fabric extender and a Cisco Nexus 5000 Series switch fails, the FEX does not disable any server ports. There is no longer a direct relation between the physical uplink port and the server port. The relationship is between the logical uplink port and the server port. Traffic for the 32 server ports is distributed over the 10 Gigabit Ethernet uplinks of the remaining seven uplink port-channel members. The servers that connect to the Cisco Nexus 2000 Series fabric extender do not notice the failure and keep forwarding traffic on the physical NIC that connects to the FEX. Single-homed servers do not lose connectivity when using dynamic pinning. Their traffic is simply redistributed over the remaining uplink port-channel members.

Dual-homed servers do not fail over to the redundant NIC. However, the oversubscription ratio for the remaining uplink ports changes. The oversubscription ratio for the remaining ports in the example is 32:7 = 4.5:1. If the uplink port channel members are already running close to maximum utilization, this oversubscription might cause traffic from all servers to be dropped, degrading performance for all servers. If congestion exists, any virtual host bus adapters (vHBAs) that send traffic across congested uplinks are stopped. For FCoE traffic, static pinning is recommended.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



The example in the figure compares the difference between static and dynamic pinning when an Ethernet uplink interface fails.

As explained, if a FEX physical uplink interface fails when static pinning is configured, then the associated FEX server interfaces fail as well. If a FEX physical uplink interface fails when dynamic pinning is configured, then the associated FEX server interfaces are repinned to the other FEX physical uplink interfaces.

https://t.me/learningnets

FEX Deployment Models: Active/Active FEX

- The active/active FEX allows a single FEX to be dual-homed to two Cisco Nexus 5000 Series switches.
- vPC is used to connect the FEX to the pair of Cisco Nexus switches.
- Ports on the FEX are associated with both switches and need to be configured consistently on both switches.
  - Effectively, ports on the FEX are vPCs and the same consistency checks are performed as for regular vPCs.
  - On the Cisco Nexus 5000 Series switches, the configurations can be synchronized automatically by using the configuration synchronization feature.
- Active/active FEX improves the availability of FEX-based solutions.
  - Server traffic is protected against FEX uplink and switch failures.

© 2012 Cisco and/or its affiliates. All rights reserved.

The active/active FEX deployment is more complex. A Cisco Nexus 2000 Series fabric extender can be controlled and configured by two Cisco Nexus 5000 Series Switches. The FEX must be configured consistently on both switches. A vPC must be set up to combine the FEX uplinks into one port channel.

Note	The two Cisco Nexus switches are configured independently. Therefore, configuration changes that are made to ports on a dual-homed Cisco Nexus 2000 Series fabric extender must be manually synchronized between the two Cisco Nexus switches. To simplify the configuration that this requirement causes, Cisco Nexus 5000 Series Switches support a configuration-synchronization feature, which allows the configuration for ports on the FEX to be synchronized between the switches.

Because a vPC-based port channel is used between the FEXs and the switches, dynamic pinning is automatically used. Traffic is balanced across the FEX uplinks, based on PortChannel load balancing. When one of the FEX uplinks is lost, traffic is balanced across the remaining uplinks for native Ethernet traffic. FCoE traffic is not distributed across both vPC link members.

In this deployment model, single-homed servers maintain connectivity if one Cisco Nexus switch fails for native Ethernet traffic. To provide the same redundancy for FCoE traffic, Small Computer Systems Interface (SCSI) multipathing software must be running on the host. This design increases the availability of single-homed servers to a level that is comparable to that of a single-homed server that connects to a chassis-based switch with dual redundant supervisors.

The choice between straight-through or active/active deployment models is not exclusive. If the FEXs are deployed in a typical ToR design, you can mix the two options. In some racks, the straight-through option can be used, whereas FEXs in others can be configured to use the active/active option.

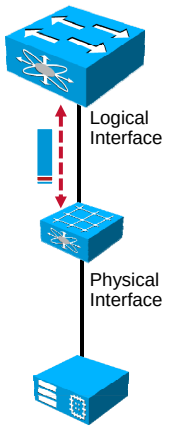
https://t.me/learningnets

Adapter FEX on the Cisco Nexus 2232 10GE Fabric Extender

The technology for connecting servers through a remote module that connects to a parent switch is known as the Cisco Adapter Fabric Extender (Adapter FEX) technology. This technology is supported on the Cisco Nexus 2232 10GE fabric extender when connected to an upstream parent Cisco Nexus switch. This topic describes Adapter FEX on the Cisco Nexus 2232 10GE fabric extender.

Cisco Nexus 5000 FEX Connectivity

- **Control and management:** Cisco Nexus 2000 Series fabric extender operates as a remote line card (local CPU with protocol offload).
- **Data plane:** Forwarding is performed on the parent switch ASICs.
- VNTag is an NIV technology that extends the parent switch port (logical interface) to the Cisco Nexus 2000 Series fabric extender (host interface).
  - VNTag is added to the packet between the fabric extender and parent switch.
  - VNTag is or can be extended down to the VIC adapter.
  - Link local state only: Internal switch header
- Port extension allows the fabric extender to act as a data path extension of the parent switch.



© 2012 Cisco and/or its affiliates. All rights reserved.

A Cisco Nexus 2000 Series fabric extender operates as a remote line card to a Cisco Nexus 5000 Series switch. The parent switch performs all control and management functions, as well as forwarding.

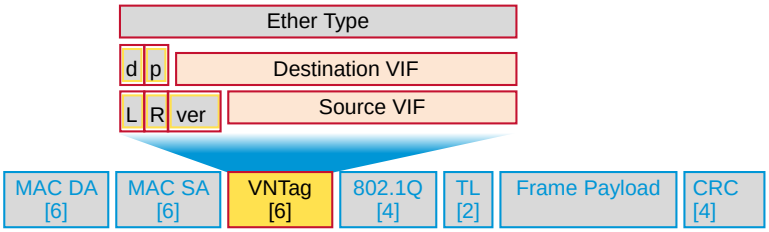
Physical host interfaces on a Cisco Nexus 2000 Series fabric extender are represented with logical interfaces on the parent Cisco Nexus 5000 Series switch.

https://t.me/learningnets

Cisco Nexus 5000 FEX Connectivity: VNTag Overview

VNTag technology has been submitted to the IEEE 802.1Qbh standards body:

- d: Direction bit (= 0 host-to-network forwarding; = 1 network-to-host)
- p: Pointer bit (set if this is a multicast frame and requires egress replication)
- L: Looped filter (set if sending back to a source Cisco Nexus 2000 Series fabric extender)
- VIF: Virtual interface



© 2012 Cisco and/or its affiliates. All rights reserved.

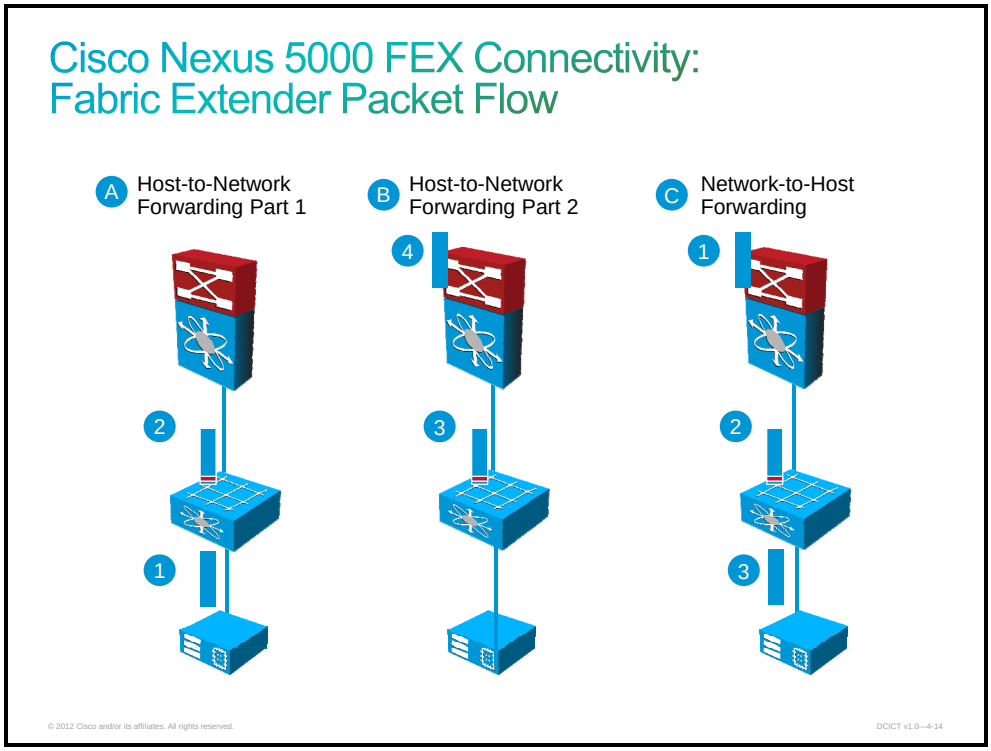
Virtual network tag (VNTag) technology has been submitted to the IEEE 802.1Qbh standards body.

Because a FEX does not support local switching, forwarding occurs on the parent switch. Forwarding from and to a FEX is based on VNTag. VNTag is an NIV technology. VNTag allows the FEX to act as a Cisco Nexus 5000 Series switch data path for all policy and forwarding. VNTag is added to the packet between the FEX and Cisco Nexus 5000 Series switch. VNTag is stripped before the packet is sent to hosts.

The figure describes VNTag fields in an Ethernet frame.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



The figure describes packet processing on a Cisco Nexus 2000 Series fabric extender.

When the host sends a packet to the network (diagrams A and B), these processes occur:

1. The frame arrives untagged from the host.
2. The Cisco Nexus 2000 Series fabric extender adds VNTag, and the packet is forwarded over a fabric link, using a specific VNTag. The Cisco Nexus 2000 Series fabric extender adds a unique VNTag for each Cisco Nexus 2000 Series host interface. The VNTag uses these field values:
   - The direction bit is set to 0, indicating host-to network forwarding.
   - The source virtual interface is set based on the ingress host interface.
   - The p (pointer), l (looped), and destination virtual interface are undefined (0).
3. The packet is received over the fabric link, using a specific VNTag. The Cisco Nexus 5500 switch extracts the VNTag, which identifies the logical interface that corresponds to the physical host interface on the actual Cisco Nexus 2000 Series fabric extender. The Cisco Nexus 5500 switch applies an ingress policy that is based on the physical Cisco Nexus 5500 switch port and logical interface:
   - Access control and forwarding are based on frame fields and logical interface policy.
   - Physical link-level properties are based on the Cisco Nexus 5500 switch port.
4. The Cisco Nexus 5500 switch strips the VNTag and sends the packet to the network.

https://t.me/learningnets

When the network sends a packet to the host (diagram C), these processes occur:

1. The frame is received on the physical or logical interface. The Cisco Nexus 5500 switch performs standard lookup and policy processing, when the egress port is determined to be a logical interface (Cisco Nexus 2000 Series) port. The Cisco Nexus 5500 switch inserts a VNTag with these fields:
   - The direction is set to 1 (network to host).
   - The destination virtual interface is set to be the Cisco Nexus 2000 Series switch port VNTag.
   - The source virtual interface is set if the packet was sourced from a Cisco Nexus 2000 Series switch port.
   - The l bit (looped) filter is set if sending back to a source Cisco Nexus 2000 Series fabric extender.
   - The p bit is set for a multicast frame and requires egress replication.
2. The packet is forwarded over the fabric link, using a specific VNTag.
3. The Cisco Nexus 2000 Series fabric extender strips the VNTag and the frame is forwarded to the host interface.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Verifying Adapter FEX on the Cisco Nexus 2232 10GE Fabric Extender

This topic describes how to verify Adapter FEX on a Cisco Nexus 2232 10GE fabric extender.

Verifying Adapter FEX

The Ethernet interfaces that connect to the FEX are now fabric interfaces.

```
N5K-1# show interface brief

--------------------------------------------------------------------------------
Ethernet      VLAN   Type Mode   Status  Reason                   Speed   Port
Interface                                                                Ch #
--------------------------------------------------------------------------------
Eth1/1        1      eth  access down    SFP validation failed       10G(D) --
Eth1/2        1      eth  access down    SFP not inserted            10G(D) --
Eth1/3        1      eth  access up      none                        10G(D) --
Eth1/4        1      eth  access up      none                        10G(D) --
Eth1/5        1      eth  access down    SFP not inserted            10G(D) --
Eth1/6        1      eth  access down    SFP not inserted            10G(D) --
Eth1/7        1      eth  access down    SFP not inserted            10G(D) --
Eth1/8        1      eth  access down    SFP not inserted            10G(D) --
Eth1/9        1      eth  fabric up      none                        10G(D) --
Eth1/10       1      eth  fabric up      none                        10G(D) --
Eth1/11       1      eth  access up      none                        10G(D) --
Eth1/12       1      eth  access up      none                        10G(D) --
```

© 2012 Cisco and/or its affiliates. All rights reserved.

Use the **show interface brief** command to verify that the interfaces are configured as FEX fabric interfaces. The mode of these interfaces is listed as fabric in the command output.

https://t.me/learningnets

Verifying Adapter FEX (Cont.)

Use the **show fex** command to determine whether the FEX has been recognized as a remote line card and to determine the hardware information of the FEX.

```
N5K-1# show fex 111

FEX: 111 Description: FEX 111, rack 1, top   state: Online
  FEX version: 5.0(2)N2(1) [Switch version: 5.0(2)N2(1)]
  Extender Model: N2K-C2232PP-10GE,  Extender Serial:
JAF1419ECAC
  Part No: 73-12748-05
 pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Eth1/9
  Fabric interface state:
    Eth1/9 - Interface Up. State: Active
    Eth1/10 - Interface Up. State: Active
```

© 2012 Cisco and/or its affiliates. All rights reserved.

To get a brief overview of the status of the FEX, use the **show fex** *fex-number* command. This command displays the state of the FEX, the pinning mode, max-links, and the state of the physical and logical uplink interfaces.

Verifying Adapter FEX Host Interfaces

After the FEX has been associated with the Cisco Nexus switch, the ports on the FEX can be seen as ports on the switch by using the FEX number as a virtual slot number.

```
N5K-1# show fex 111 detail

FEX: 111 Description: FEX 111, rack 1, top   state: Online
-- output omitted --

 pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Eth1/9
  Fabric interface state:
    Eth1/9 - Interface Up. State: Active
    Eth1/10 - Interface Up. State: Active
Fex Port        State  Fabric Port
      Eth111/1/1  Down       Eth1/9
      Eth111/1/2  Down       Eth1/9
      Eth111/1/3  Down       Eth1/9
      Eth111/1/4  Down       Eth1/9
      Eth111/1/5  Down       Eth1/9
      Eth111/1/6  Down       Eth1/9
      Eth111/1/7  Down       Eth1/9
-- output omitted –
```

© 2012 Cisco and/or its affiliates. All rights reserved.

Similar to the **show fex** command, the **show fex detail** command provides the description of the FEX, the pinning mode, the number of max-links, and the state of the physical and logical uplink. In addition, the command shows the FEX interfaces and the fabric interfaces to which they are pinned.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying Adapter FEX Host Interfaces (Cont.)

The FEX host interfaces appear in the running configuration and as interfaces of the Cisco Nexus 5000 Series Switch.

```
N5K-1# show running-config | begin "interface Ethernet111"

interface Ethernet111/1/1

interface Ethernet111/1/2

interface Ethernet111/1/3
-- output omitted --
```

```
N5K-1# show interface ethernet 111/1/1

Ethernet111/1/1 is down (Link not connected)
  Hardware: 100/1000 Ethernet, address: 5475.d0e1.1482 (bia
5475.d0e1.1482)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  Port mode is access
  auto-duplex, auto-speed
-- output omitted --
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The newly connected FEX interfaces appear in the running configuration of the parent Cisco Nexus 5500 switch. The interfaces are listed by using the FEX number/module/slot number designation. These interfaces are the host (server) interfaces. The fabric (uplink) interfaces of the FEX do not appear.

The **show interface** command displays the Ethernet parameters of the FEX host interfaces, which connect to the Cisco Nexus 5500 switch through the fabric interfaces and physically reside on the Cisco Nexus 2000 Series fabric extender.

https://t.me/learningnets

Summary

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

Summary

- Cisco Nexus 2232 10GE fabric extenders provides additional 10 Gb/s connectivity for servers and supports FCoE and DCB features.
- Connecting a Cisco Nexus 2232 10GE fabric extender to the Cisco Nexus 5500 switch takes place through the fabric uplinks.
- Cisco Nexus 2232 10GE fabric extenders can be connected through straight-through static pinning, straight-through dynamic pinning, or active/active connections. Careful consideration should be taken as to which connectivity option to use. During failover, this option can affect oversubscription ratios for server connectivity.
- The Cisco Nexus 2232 10GE fabric extender and the Cisco Nexus 5500 switch use VNTag NIV technology to identify remotely connected devices through the Cisco Nexus 2232 10GE fabric extender.
- Verify that the Cisco Nexus 2232 10GE fabric extender has been seen and that ports are remotely available for use, by using commands such as **show fex** and **show fex detail**.

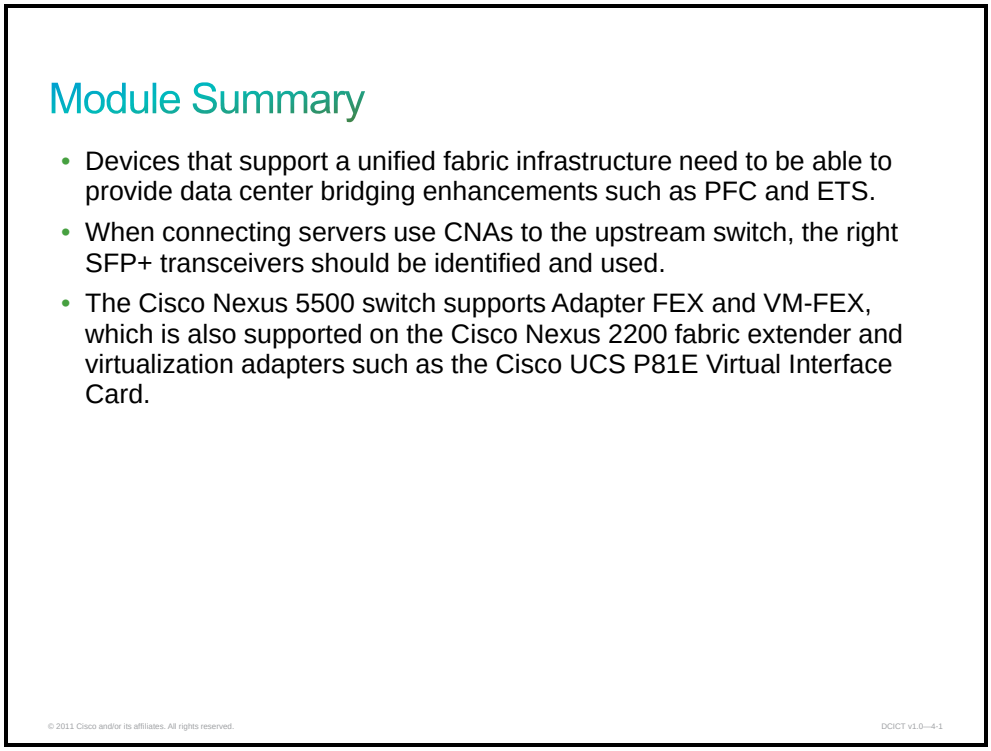
© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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



Supporting a unified fabric, including I/O consolidation, provides many benefits:

- Reduced cabling
- Reduced power requirements
- Fewer devices to manage

Devices in a unified fabric need to be able to support Data Center Bridging (DCB) standards:

- **Priority Flow Control (PFC):** Provides different levels of service to separate streams of traffic over the same connection
- **Enhanced Transmission Selection (ETS):** Guarantees levels of quality of service (QoS) and bandwidth requirements of separate streams of traffic over the same network architecture
- **Unified fabric congestion management (optional):** Provides end-to-end congestion management
- **DCB Exchange (DCBX):** Discovers and negotiates configuration requirements to edge devices

Cisco Nexus 5000 Series Switches support enhanced small form-factor pluggable (SFP+) transceivers for copper networks and optical fibers. For in-rack or adjacent-rack cabling, the SFP+ direct-attach 10 Gigabit Ethernet copper solution, which integrates transceivers with twinax cables is preferred. These cables require less power. For longer cable runs, the multimode, short-reach (SR) optical SFP+ transceivers should be used.

https://t.me/learningnets

Cisco Nexus 5500 switches also support servers that use virtualization cards such as the Cisco UCS P81E Virtual Interface Card. This card supports the Cisco Adapter FEX technology, in which virtual NICs (vNICs) in the server are connected logically to virtual interfaces on a Cisco Nexus 5500 switch.

When connecting servers that are consolidating I/O, the upstream switch receives Fibre Channel and Ethernet traffic across the 10-Gb/s Ethernet connection. The Fibre Channel over Ethernet (FCoE) traffic is forwarded to another switch or device that supports FCoE. The FCoE traffic might be stripped back to a Fibre Channel frame and sent to a Fibre Channel switch or device. Therefore, a Cisco Nexus 5500 switch might need connectivity to both the LAN and SAN infrastructures.

To scale the number of FCoE-supported interfaces on a Cisco Nexus 5500 switch, one or more Cisco Nexus 2232 10GE fabric extenders can be connected as remote I/O modules. Each Cisco Nexus 2232 10GE fabric extender has an additional 32 10-Gb/s Ethernet interfaces and is DCB aware. The Cisco Nexus 2232 10GE fabric extender can be connected in either a straight-through dynamic, straight-through static, or active/active topology.

Because the Cisco Nexus 2232 10GE fabric extender is a remote I/O module with remote interfaces to the parent switch, it falls under the Adapter FEX technology. To ensure that the FEX knows which virtual server the traffic was sourced from and which is the receiver, the Cisco Nexus 5500 switch and Cisco Nexus 2232 10GE fabric extender use a virtual network tag (VNTag).

References

For additional information, refer to these resources:

- Cisco Systems, Inc. *Twinax Cables Certification Matrix for Cisco Nexus 2000 and 5000 Series Switches Data Sheet.* http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps10110/data_sheet_c78-568589_ps9670_Products_Data_Sheet.html.
- Cisco Systems, Inc. *Cisco Nexus 5548, 5548UP, 5596UP, and 5596T Switches Data Sheet.* http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-618603.html.
- Cisco Systems, Inc. *Data Center Bridging.* http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns783/at_a_glance_c45-460907.pdf. 2010.
- Cisco Systems, Inc. *Priority Flow Control: Build Reliable Layer 2 Infrastructure.* http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-542809_ns783_Networking_Solutions_White_Paper.html.
- Cisco Systems, Inc. *Unified Fabric: Cisco's Innovation for Data Center Networks.* http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns783/white_paper_c11-462422.html.
- Cisco Systems, Inc. *Network Adapter Virtualization Design (Adapter-FEX) with Cisco Nexus 5548 Switches and Cisco Nexus 2232 Fabric Extenders.* http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/guide_c07-690080_ns1118_Networking_Solutions_White_Paper.html.
- Cisco Systems, Inc. *Cisco Adapter Fabric Extender.* http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns945/ns1134/ns1118/data_sheet_c78_657397.pdf.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

- Cisco Systems, Inc. *Cisco Nexus 5000 Series NX-OS Fibre Channel over Ethernet Configuration Guide, Release 5.1(3)N1(1).* http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/fcoe/513_n1_1/b_Cisco_n5k_fcoe_config_gd_re_513_n1_1.html.
- Cisco Systems, Inc. *Use of Mode Conditioning Patch Cables in Gigabit Ethernet and 10 Gigabit Ethernet Laser-Based Transmissions.* http://www.cisco.com/en/US/prod/collateral/modules/ps5455/product_bulletin_c25-530836.html.
- Cisco Systems, Inc. *Cisco UCS C-Series Servers Integrated Management Controller CLI Configuration Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/sw/cli/config/guide/131/b_Cisco_UCS_C-Series_CLI_Configuration_Guide_131.html.

https://t.me/learningnets

© 2012 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 enhancement provides lossless delivery for selected CoSs? (Source: Describing DCB)

A) TRILL
B) QCN
C) PFC
D) ETS

Q2) What is the default CoS value for FCoE traffic over a 10-Gb/s Ethernet connection? (Source: Describing DCB)

A) 2
A) 3
B) 5
C) 7

Q3) Which enhancement provides differentiation of traffic by creating priority groups? (Source: Describing DCB)

A) ETS
B) QCN
C) TRILL
D) PFC

Q4) Which protocol is used to discover the DCB-capable devices in the network? (Source: Describing DCB)

A) SNMP
B) TRILL
C) LACP
D) DCBX

Q5) Which version of DCBX is supported on T11-compliant second-generation CNAs? (Source: Describing DCB)

A) CIN-DCBX
B) CEE-DCBX

Q6) Which deployment model for Cisco Nexus 2232 10GE fabric extender connections to the Cisco Nexus 5548 Switch causes the corresponding server ports to fail during a FEX uplink port failure? (Source: Describing Enhanced FCoE Scalability with Cisco Nexus 2232 10GE Fabric Extenders)

A) straight-through using static pinning
B) straight-through using dynamic pinning
C) active/active FEX using vPC
D) active/standby FEX using vPC

https://t.me/learningnets

Q7) On the Cisco Nexus 5548 Switch, which interfaces are created to represent the physical host interfaces on a Cisco Nexus 2232 10GE fabric extender? (Source: Describing Enhanced FCoE Scalability with Cisco Nexus 2232 10GE Fabric Extenders)

A) physical host interfaces
B) physical Ethernet interfaces
C) logical interfaces
D) virtual interfaces

Q8) Which **show** command can you use to verify the fabric interfaces to which the FEX interfaces are pinned? (Source: Describing Enhanced FCoE Scalability with Cisco Nexus 2232 10GE Fabric Extenders)

A) **switch# show interfaces detail**
B) **switch# show fex**
C) **switch# show interfaces brief**
D) **switch# show fex detail**

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Module Self-Check Answer Key

Q1) C

Q2) B

Q3) A

Q4) D

Q5) B

Q6) A

Q7) C

Q8) D

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Module 5

Cisco UCS

Overview

This module describes the Cisco Unified Computing System (UCS) product family.

Module Objectives

Upon completing this module, you will be able to articulate and verify correct operation of the components, connectivity, and management of Cisco UCS devices. This ability includes being able to meet these objectives:

- Describe the components of the Cisco UCS B-Series product family
- Describe the components of the Cisco UCS C-Series product family
- Describe how Cisco UCS B-Series Blade Servers connect to the chassis and achieve external LAN and SAN connectivity
- Describe the required information to perform initial setup of a Cisco UCS B-Series cluster and to validate cluster operation
- Describe how Cisco UCS Manager discovers new hardware and manages and monitors all Cisco UCS components
- Describe the Cisco UCS Manager features that allow rapid provisioning and consistent application of policy

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 1

Describing the Cisco UCS B-Series Product Family

Overview

The purpose of this lesson is to describe and verify the components of the Cisco Unified Computing System (UCS) B-Series product family.

Objectives

Upon completing this lesson, you will be able to describe the components of the Cisco UCS B-Series product family. You will be able to meet these objectives:

- Describe fabric interconnect options
- Describe the Cisco UCS 5108 Blade Server Chassis
- Describe the Cisco UCS B200 M3 Blade Server
- Describe the Cisco UCS B230 M2 Blade Server
- Describe the Cisco UCS B250 M2 Extended Memory Blade Server
- Describe the Cisco UCS B440 M2 High-Performance Blade Server
- Describe the features of mezzanine cards for Cisco UCS B-Series Blade Servers
- Describe the best practices for populating DRAM in Cisco UCS B-Series Blade Servers

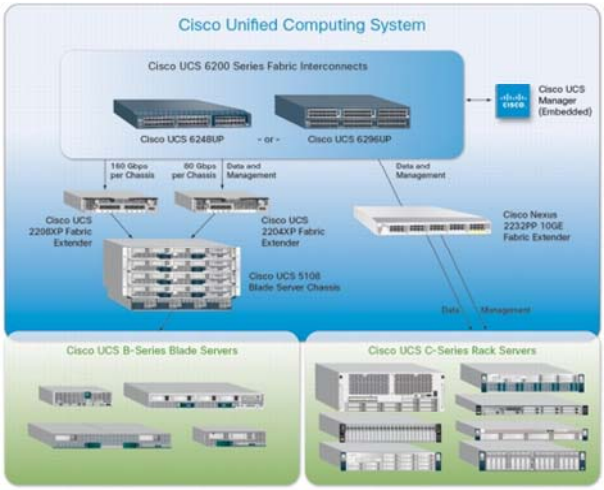
https://t.me/learningnets

Cisco UCS 6100 and 6200 Series Fabric Interconnects

This topic describes the Cisco UCS 6100 and 6200 Series Fabric Interconnects and available options.

Overview of Cisco UCS Family

The Cisco UCS family includes Cisco UCS B-Series and C-Series.



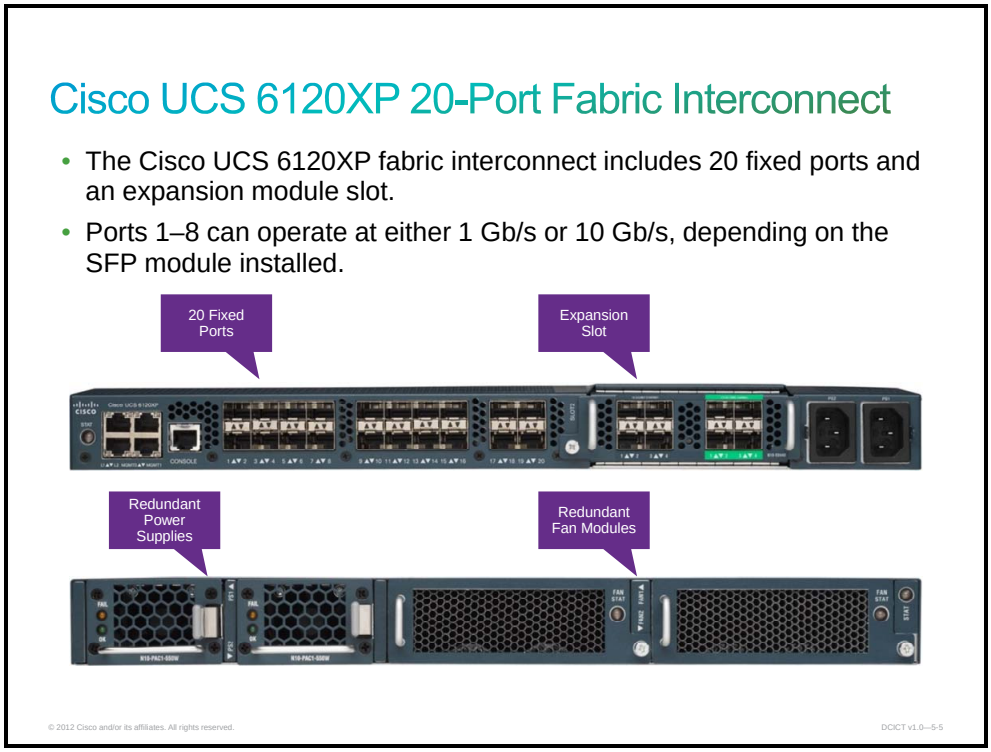
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS family of products includes blade servers, rack servers, and the associated components that are necessary for connectivity. Cisco UCS connectivity is covered in more detail in the lesson "Connecting Cisco UCS B-Series Blade Servers."

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 6120XP 20-Port Fabric Interconnect



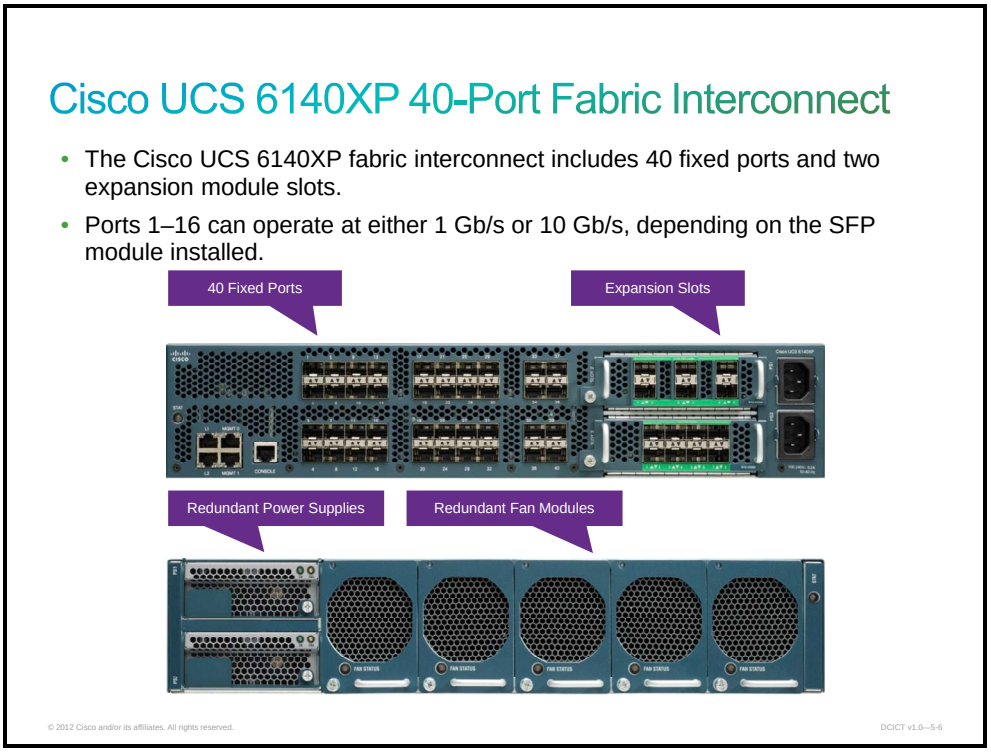
The Cisco UCS 6120XP 20-Port Fabric Interconnect is typically deployed in clustered pairs and provide unified access to both networks (LAN) and storage (SAN).

The Cisco UCS 6120XP fabric interconnect is a 1-rack unit (RU) interconnect that provides these features:

- 20 fixed 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) enhanced small form-factor pluggable (SFP+) ports
- 520 Gb/s of throughput
- A single expansion module bay that can support as many as eight Fibre Channel ports or six 10 Gigabit Ethernet ports using the SFP+ interfaces
- Support for as many as 160 servers or 20 chassis as one seamless system

https://t.me/learningnets

Cisco UCS 6140XP 40-Port Fabric Interconnect



The Cisco UCS 6140XP 40-Port Fabric Interconnect is typically deployed in clustered pairs and provides unified access to both networks and storage.

The Cisco UCS 6140XP fabric interconnect is a 2-RU interconnect that provides these features:

- 40 fixed 10 Gigabit Ethernet and FCoE SFP+ ports
- 1.04 Tb/s of throughput
- Two expansion module bays, each supporting as many as eight Fibre Channel ports or six 10 Gigabit Ethernet ports using the SFP+ interfaces
- Support for as many as 320 servers or 40 chassis as one seamless system

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Comparison of Cisco UCS 6120XP and Cisco UCS 6140XP Fabric Interconnects

Comparison of Cisco UCS 6120XP and Cisco UCS 6140XP Fabric Interconnects

The table summarizes the features of the Cisco UCS 6120XP and Cisco UCS 6140XP fabric interconnects.

Item	Cisco UCS 6120XP	Cisco UCS 6140XP
Description	20-port fabric interconnect	40-port fabric interconnect
Form factor	1 RU	2 RU
Number of fixed 10-GB interfaces	20	40
Throughput	520 Gb/s	1.04 Tb/s
Expansion slots	1	2
Fan modules	1+1	4+1

© 2012 Cisco and/or its affiliates. All rights reserved.

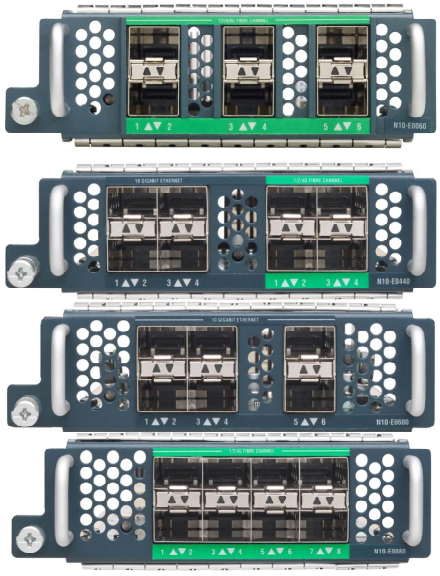
The figure summarizes the features of the Cisco UCS 6120XP and 6140XP fabric interconnects.

Cisco UCS 6100 Series Expansion Modules

Cisco UCS 6100 Series Expansion Modules

GEMs are available for the Cisco UCS 6120XP and 6140XP, in four models:

- 6-port 1-, 2-, 4-, or 8-Gb/s Fibre Channel
- 4-port 10 Gigabit Ethernet or 4-port Fibre Channel
- 6-port 10 Gigabit Ethernet
- 8-port 1-, 2-, or 4-Gb/s Fibre Channel



© 2012 Cisco and/or its affiliates. All rights reserved.

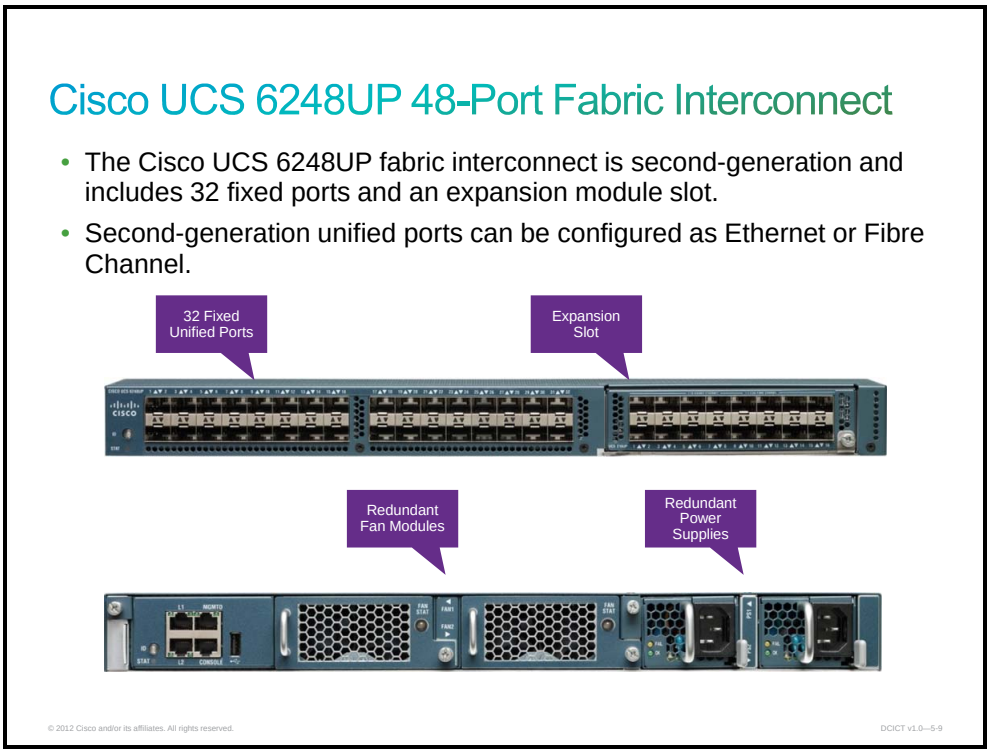
https://t.me/learningnets

The expansion slots of Cisco UCS 6100 Series Fabric Interconnects support four Generic Expansion Modules (GEMs). GEMs provide native Fibre Channel connectivity for FCoE and Fibre Channel direct-attach storage (DAS). The GEMs also offer additional 10 Gigabit Ethernet for uplinks and uplink port channels.

GEMs are available in these offerings:

- 6-port 1/2/4/8-Gb Fibre Channel
  - 1/2/4-Gb/s operation with standard Fibre Channel small form-factor pluggable (SFP) transceivers
  - 2/4/8-Gb/s operation with SFP+ transceivers
- 8-port 1/2/4-Gb Fibre Channel
- 4-port 10 Gigabit Ethernet/4-port Fibre Channel
- 6-port 10 Gigabit Ethernet

Cisco UCS 6248UP 48-Port Fabric Interconnect



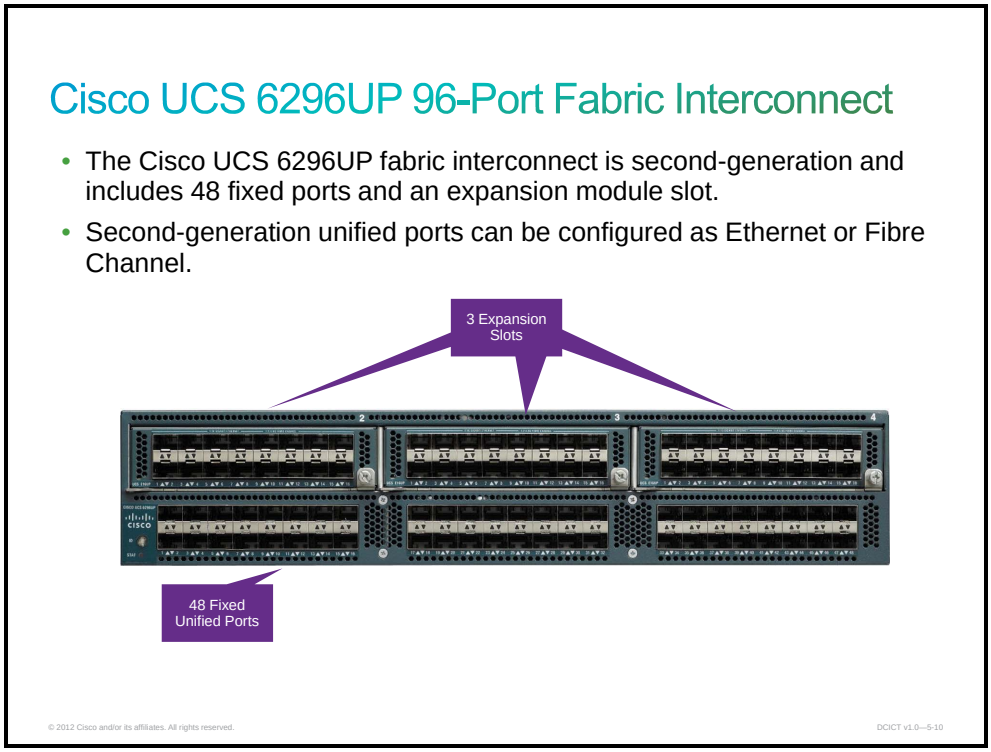
The Cisco UCS 6248UP 48-Port Fabric Interconnect offers next-generation performance to Cisco UCS B-Series Blade Servers. These blade servers can leverage 80- to 160-Gb/s bandwidth to service the most demanding workloads. The table lists the features of this fabric interconnect.

Cisco UCS 6248UP Fabric Interconnect Features

Feature	Description
Power supply	AC or DC power supplies available
Unified ports	Can configure all ports as either 10 Gigabit Ethernet or Fibre Channel
Throughput	960 Gb/s

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS 6296UP 96-Port Fabric Interconnect



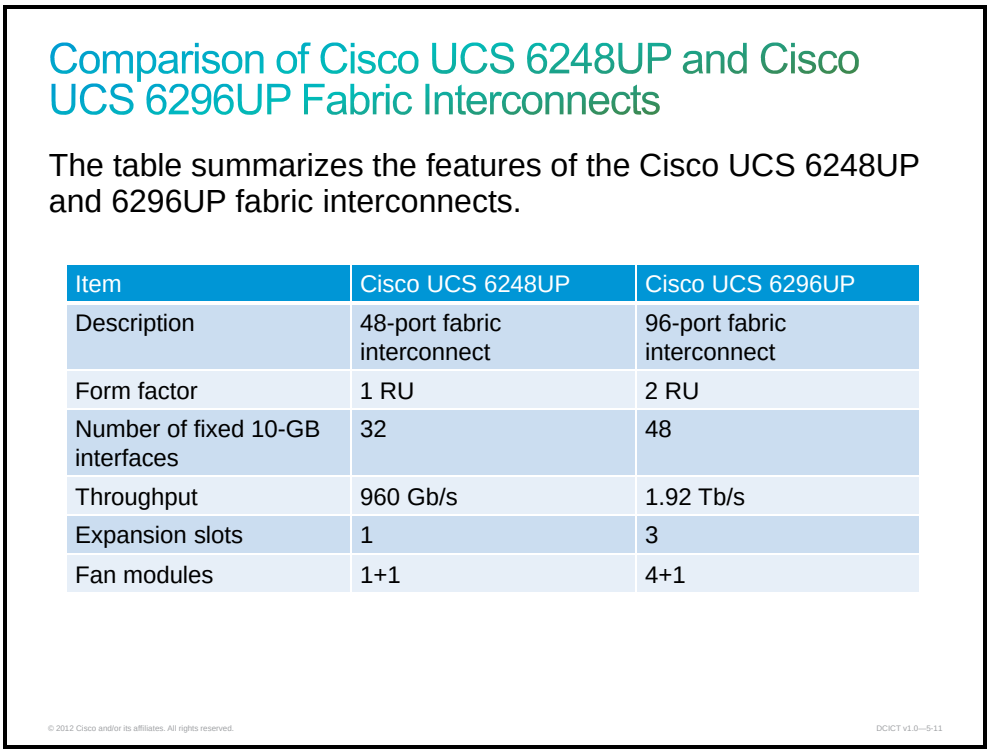
The Cisco UCS 6296UP 96-Port Fabric Interconnect also offers next-generation performance to Cisco UCS B-Series Blade Servers. These blade servers can leverage 80- to 160-Gb/s bandwidth. The table lists the features of this fabric interconnect.

Cisco UCS 6296UP Fabric Interconnect Features

Feature	Description
Power supply	AC or DC power supplies available
Unified ports	Can configure all ports as either 10 Gigabit Ethernet or Fibre Channel
Throughput	1920 Gb/s

https://t.me/learningnets

Comparison of Cisco UCS 6248UP and Cisco UCS 6296UP Fabric Interconnects



This figure summarizes the features of the Cisco UCS 6248UP and Cisco UCS 6296UP fabric interconnects.

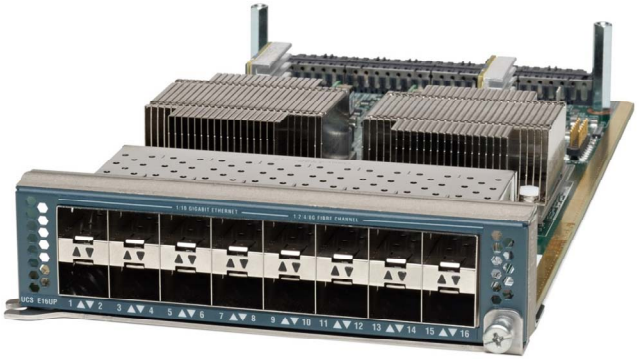
© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 6200 Series Expansion Module

Cisco UCS 6200 Series Expansion Module

- The 16-port unified port expansion module for the Cisco UCS 6200 Series Fabric Interconnects allows Ethernet only, Fibre Channel only, or mixed ports.
- The port personality is determined by the configuration and requires the appropriate SFP module.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS 6248UP fabric interconnect includes 32 fixed unified ports and an expansion slot. The unified port 16-port expansion module increases the port density to 48 unified ports. Port personality is determined by configuration in Cisco UCS Manager and choice of SFP or SFP+ transceiver.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 5108 Blade Server Chassis

This topic describes the components of the Cisco UCS 5108 Blade Server Chassis.

Cisco UCS 5108 Blade Server Chassis—Front

- Blade servers and power supplies are installed and removed from the front of the chassis.
- Servers and power supplies are hot-swappable.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS 5108 Blade Server Chassis is 6-RU high, mounts in any industry-standard 19-inch 4-post rack or cabinet, and uses standard front-to-back cooling. One chassis can accommodate as many as eight half-slot or four full-slot Cisco UCS B-Series Blade Servers or other combinations.

Blade servers and power supplies are inserted from the front of the chassis.

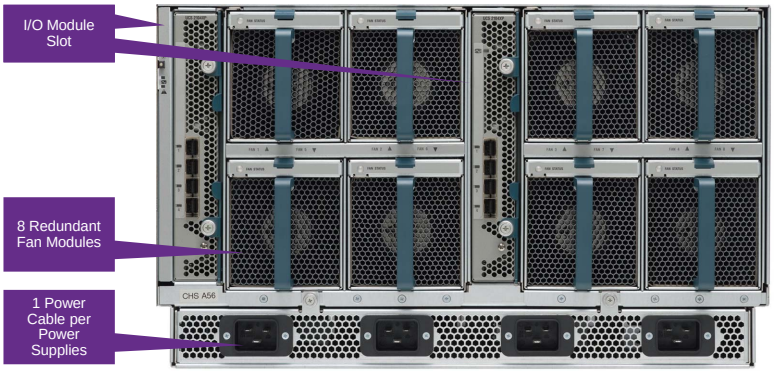
 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 5108 Blade Server Chassis: Rear

Cisco UCS 5108 Blade Server Chassis—Rear

- Power inputs, I/O modules, and fan modules are installed and removed from the rear of the chassis.
- All rear-accessible components are hot-swappable.



© 2012 Cisco and/or its affiliates. All rights reserved.

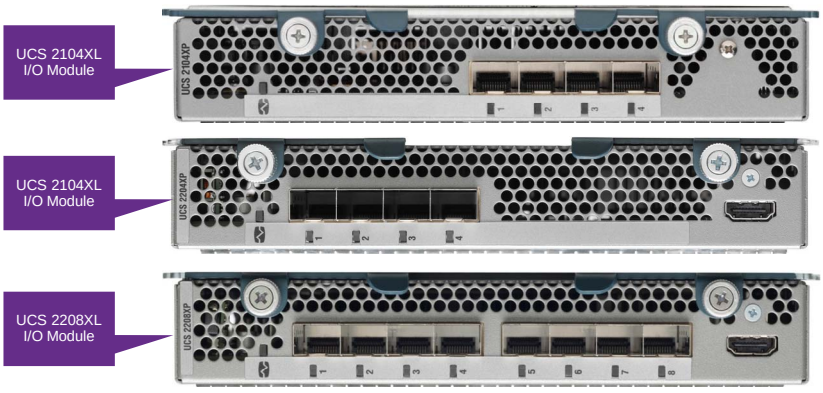
The rear of the chassis provides access to the I/O modules, hot-swappable fan modules, and power input.

The Cisco UCS 5108 Blade Server Chassis can be ordered with AC or DC power input.

Cisco UCS I/O Modules

Cisco UCS I/O Modules

- The first-generation Cisco UCS 2104XP supplies up to 40 Gb/s to each fabric.
- The second-generation Cisco UCS 2204XP supplies up to 40 Gb/s to each fabric.
- The second-generation Cisco UCS 2208XP supplies up to 80 Gb/s to each fabric.



© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

The purpose of an I/O module is to multiplex Ethernet and FCoE connections from servers to the fabric interconnects.

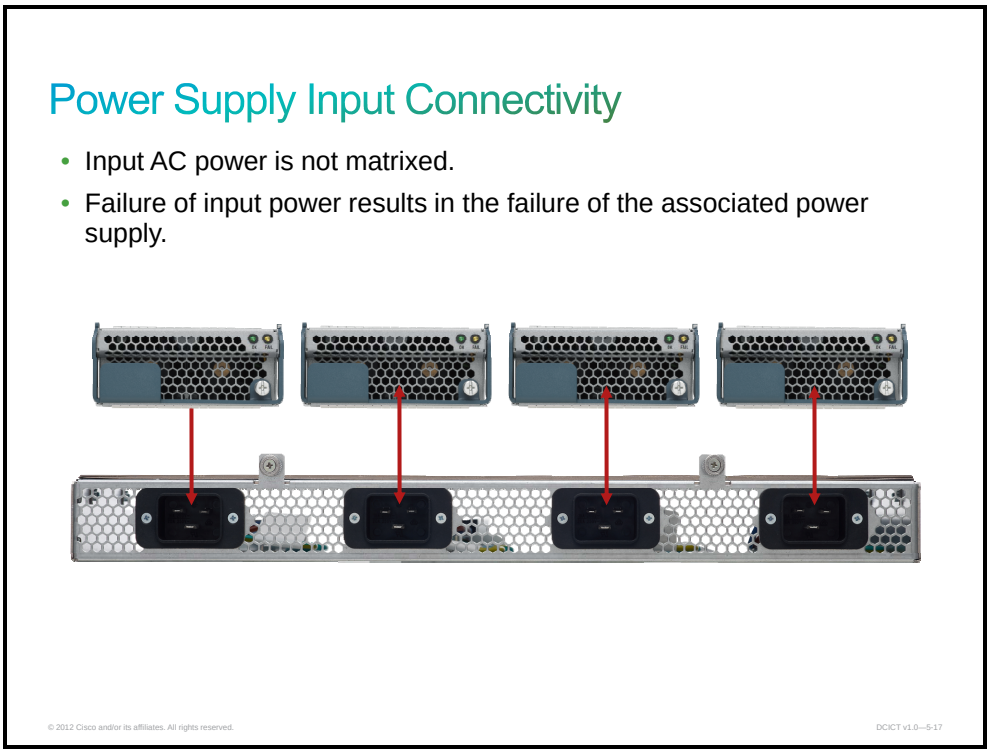
The Cisco UCS 2104XP I/O module is a first-generation module that supplies as much as 40 GB/s bandwidth per fabric to blade servers. The Cisco UCS 5108 Blade Server Chassis has I/O module slots for fabric A and fabric B.

The Cisco UCS 2204XP I/O module is a second-generation module that provides as much as 40 GB/s bandwidth to each fabric. This module is deployed with the second-generation Cisco UCS 6248UP or Cisco UCS 6296UP fabric interconnect and Cisco UCS Virtual Interface Card 1280 or Virtual Interface Card 1240 mezzanine cards. The Cisco UCS 2204XP I/O module can support fabric port channels.

Note Mezzanine refers to the fact that the adapter cards for Cisco UCS B-Series Blade Servers lie above and connect to the motherboard. This arrangement is referred to as a mezzanine, like a mezzanine area that sits between two floors of a building. Photos of this form-factor appear later in this lesson.

The Cisco UCS 2208XP I/O module is a second-generation module that provides as much as 80 GB/s to each fabric. This module is deployed with the second-generation Cisco UCS 6248UP or Cisco UCS 6296UP fabric interconnect and Cisco UCS Virtual Interface Card 1280 or Virtual Interface Card 2140 mezzanine cards. The Cisco UCS 2208XP I/O module can support fabric port channels.

Power Supply Input Connectivity

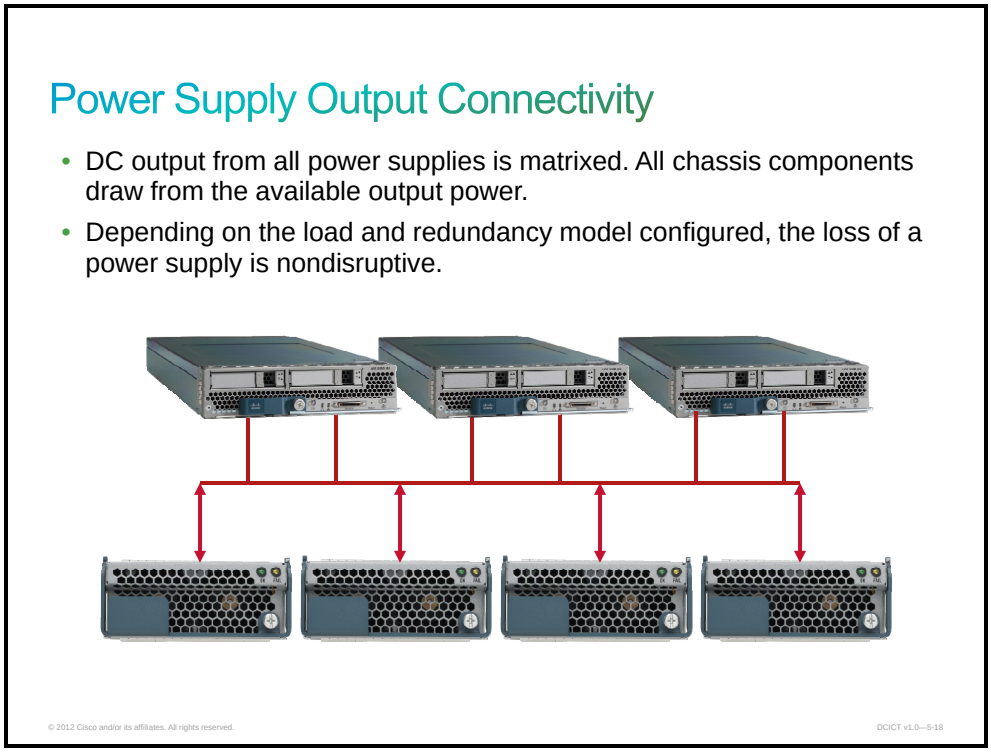


Each 220-V AC input connector is electrically associated to the power supply in that slot. If the power cord s removed or loses power, then the connected power supply fails.

The input connectivity characteristics of AC input also apply to the Cisco UCS 5108 Blade Server DC chassis.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Power Supply Output Connectivity



Unlike input power, the sum of all available DC power that the power supplies produce is matrixed, or pooled, and made available to all components. Even when a power supply or power supply input is removed, chassis components continue to operate so long as the minimum power requirements are met.

https://t.me/learningnets

Cisco UCS B200 M3 Blade Server

This topic describes the Cisco UCS B200 M3 Blade Server.

Cisco UCS B200 M3

The Cisco UCS B200 M3 Blade Server balances performance and density for virtualization and other mainstream data-center workloads.

Specification	Description
DRAM	Up to 384 GB
Hard drives	Up to 2
CPU	1 or 2 Intel Xeon E5-2600 series
RAID	RAID 0 or 1
Adapters (mezzanine slots)	1



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS B200 M3 Blade Server is a half-width, 2-socket blade server with substantial throughput and scalability. The table lists the features of this blade server.

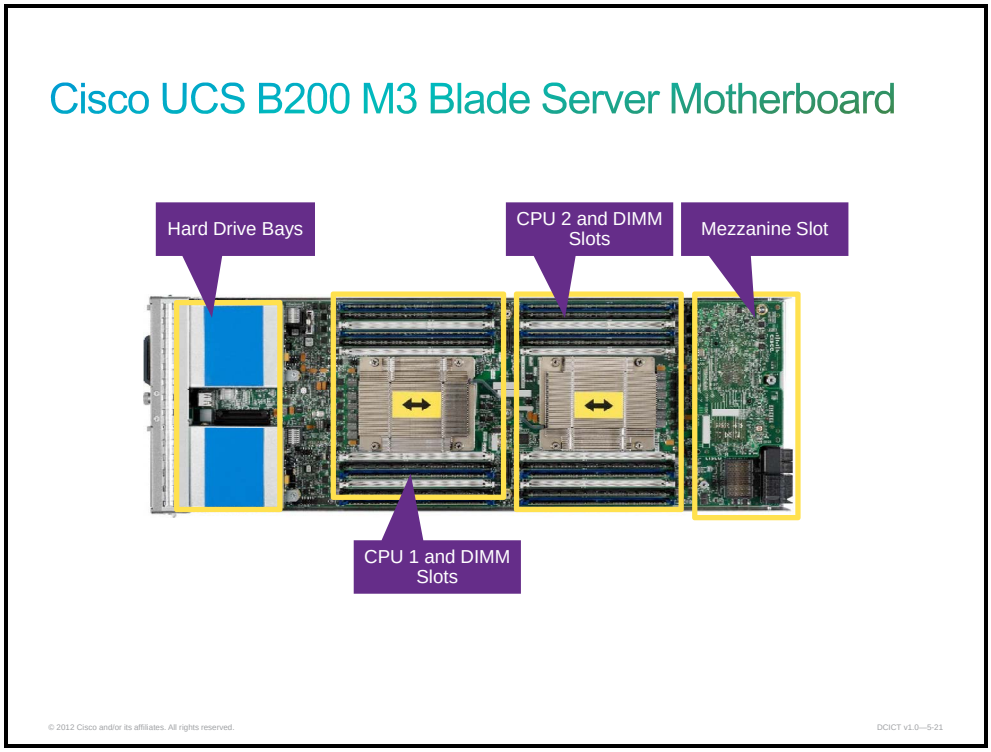
Cisco UCS B200 M3 Blade Server Core Features

Feature	Description
Processor sockets	2
Processors supported	Intel Xeon processor E5-2600 Series
Memory capacity	24 DIMMs; up to 384 GB
Memory size and speed	4-, 8-, and 16-GB double data rate 3 (DDR3); 1066 MHz and 1600 MHz
Internal disk drive	Two 2.5-in. small form-factor (SFF) Serial Attached SCSI (SAS) or 15-mm Serial Advanced Technology Attachment (SATA) solid state drive (SSD)
Maximum internal storage	Up to 2.0 TB
Integrated RAID	0,1
I/O throughput	Up to 80 Gb/s (with Cisco UCS Virtual Interface Card 1280)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS B200 M3 Blade Server Motherboard



On the motherboard, each CPU owns 12 DIMM slots. The maximum memory capacity requires two installed CPUs.

https://t.me/learningnets

Cisco UCS B230 M2 Blade Server

This topic describes the Cisco UCS B230 M2 Blade Server.

Cisco UCS B230 M2 Blade Server

The Cisco UCS B230 M2 Blade Server allows higher density virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 512 GB
Hard drives	Up to two
CPU	1 or 2 Intel E7-2800 series
RAID	RAID 0 or 1
Adapters (mezzanine slots)	1



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS B230 M2 Blade Server extends the capabilities of Cisco UCS and delivers new levels of performance and density in a half-slot form factor. The table lists the features of this blade server.

Cisco UCS B230 M2 Blade Server Core Features

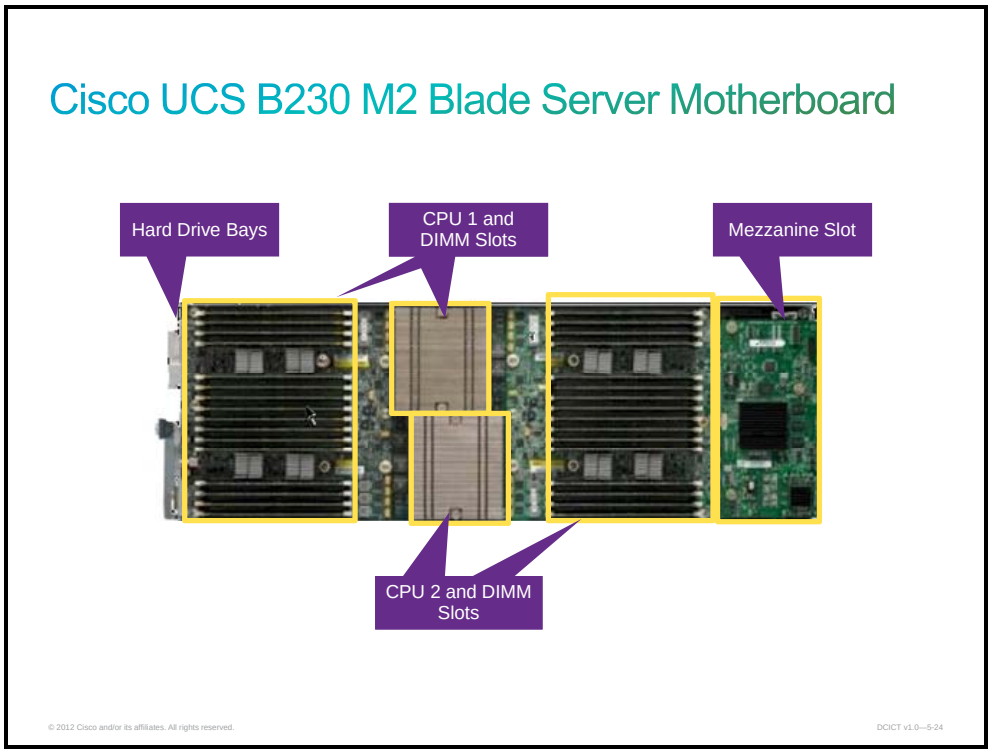
Feature	Description
Processor sockets	2
Processors supported	Intel Xeon processor E7-2800 product family
Memory capacity	32 DIMMs; up to 512 GB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz
Internal disk drive	Two 2.5-in. 7-mm SSD
Maximum internal storage	Up to 200 GB (SSDs)
Integrated RAID	0,1
I/O throughput	Up to 80 Gb/s (with Cisco UCS Virtual Interface Card 1280)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS B230 M2 Blade Server Motherboard

This topic describes the Cisco UCS B230 M2 Blade Server.



On the motherboard, each CPU owns 16 DIMM slots. The maximum memory capacity requires two installed CPUs.

https://t.me/learningnets

Cisco UCS B250 M2 Extended Memory Blade Server

This topic describes the Cisco UCS B250 M2 Extended Memory Blade Server.

Cisco UCS B250 M2 Extended Memory Blade Server

The Cisco UCS B250 M2 blade server provides expanded memory and up to 40-Gb/s bandwidth to serve I/O-intensive workloads.

Specification	Description
DRAM	Up to 384 GB
Hard drives	Up to two
CPU	1 or 2 Intel Xeon 5600 series
RAID	RAID 0 or 1
Adapters (mezzanine slots)	2



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS B250 M2 Extended Memory Blade Server helps to maximize performance and capacity for the most demanding virtualization and large dataset applications. This blade server features more than double the industry-standard memory of traditional 2-socket x86 servers. The table lists the features of this blade server.

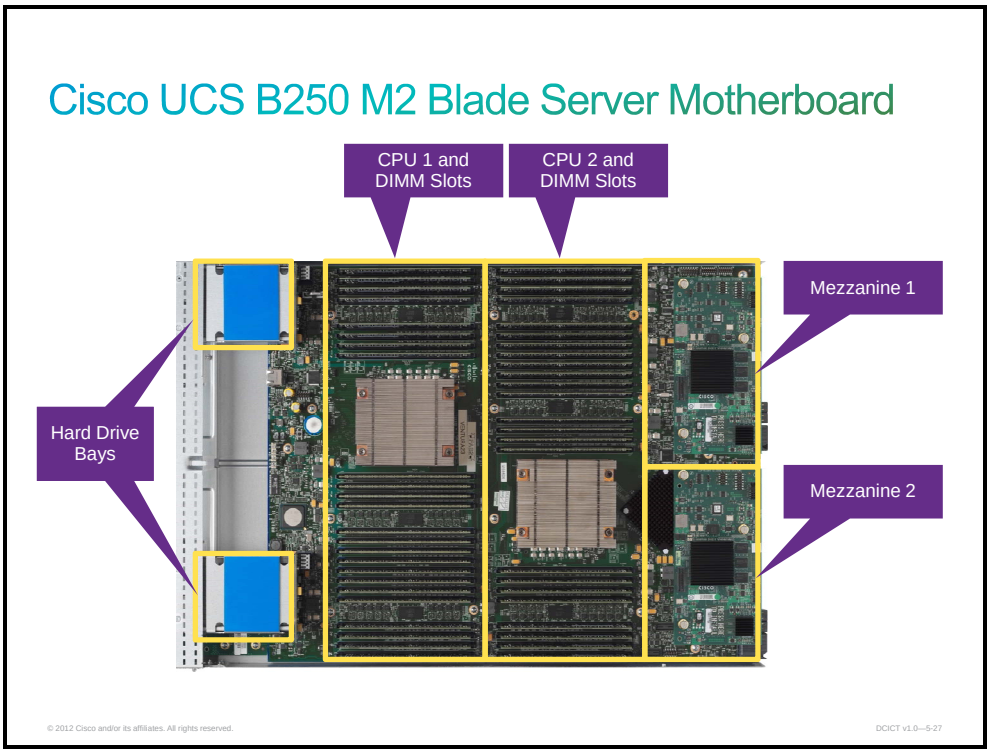
Cisco UCS B250 M2 Extended Memory Blade Server Core Features

Feature	Description
Processor sockets	2
Processors supported	Intel Xeon processor 5600 Series
Memory capacity	48 DIMMs; up to 384 GB
Memory size and speed	4- and 8-GB DDR3; 1066 MHz and 1333 MHz
Internal disk drive	Two 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 1.2 TB
Integrated RAID	0,1
I/O throughput	Up to 160 Gb/s (with 2 Cisco UCS Virtual Interface Card 1280)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS B250 M2 Extended Memory Blade Server Motherboard



On the motherboard, each CPU owns 24 DIMM slots. The maximum memory capacity requires two installed CPUs. Two mezzanine slots allow a maximum server bandwidth of 40 Gb/s with first-generation adapters and up to 160 Gb/s with Cisco UCS Virtual Interface Card 1280 adapters.

https://t.me/learningnets

Cisco UCS B440 M2 High-Performance Blade Server

This topic describes the Cisco UCS B440 M2 High-Performance Blade Server.

Cisco UCS B440 M2 High-Performance Blade Server

The Cisco UCS B440 M2 blade server is designed to power the most demanding data center applications.

Specification	Description
DRAM	Up to 512 GB
Hard drives	Four
CPU	2 or 4 Intel E7-4800 series
RAID	*RAID 0, 1, 5, 6, 10, 50 & 60
Adapters (mezzanine slots)	2



***** The actual RAID groups available depend on a hardware license in the form of a RAID key.

© 2012 Cisco and/or its affiliates. All rights reserved.

With the Intel Xeon processor E7-4800 product family, the 4-socket Cisco UCS B440 M2 High-Performance Blade Server extends Cisco UCS by offering new levels of performance and scalability. The table lists the features of this blade server.

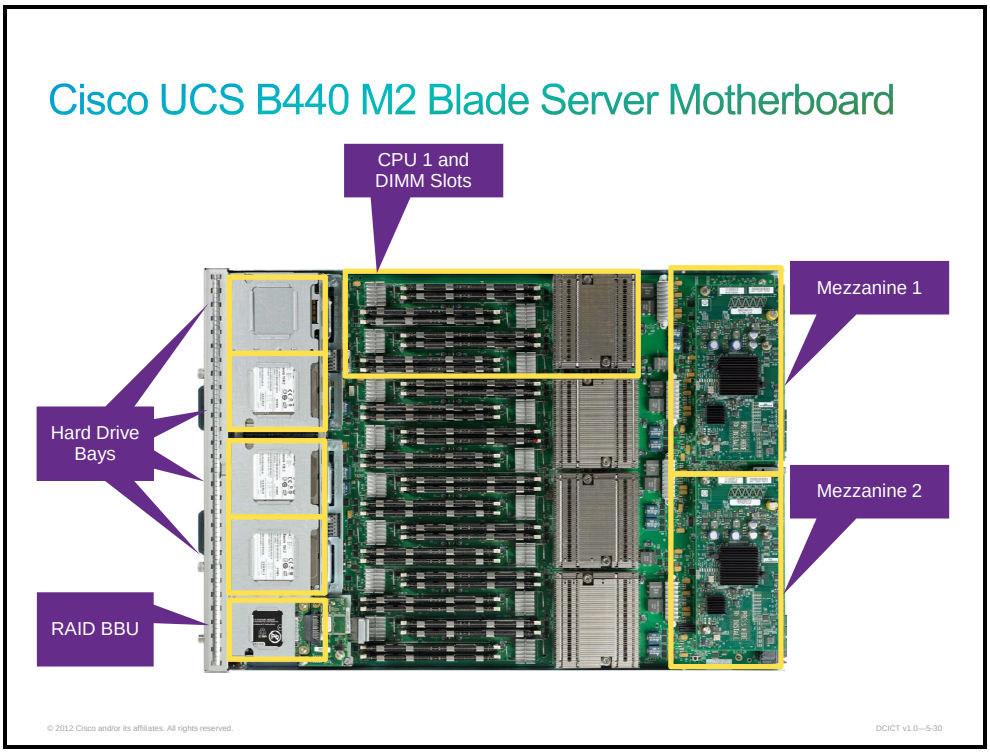
Cisco UCS B440 M2 High-Performance Blade Server Core Features

Feature	Description
Processor sockets	4
Processors supported	Intel Xeon processor E7-4800 product family
Memory capacity	32 DIMMs; up to 512 GB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz
Internal disk drive	Four 2.5-in. SFF SAS or SATA
Maximum internal storage	Up to 2.4 TB
Integrated RAID	0,1,5,6
I/O throughput	Up to 160 Gb/s (with 2 x Cisco UCS Virtual Interface Card 1280)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

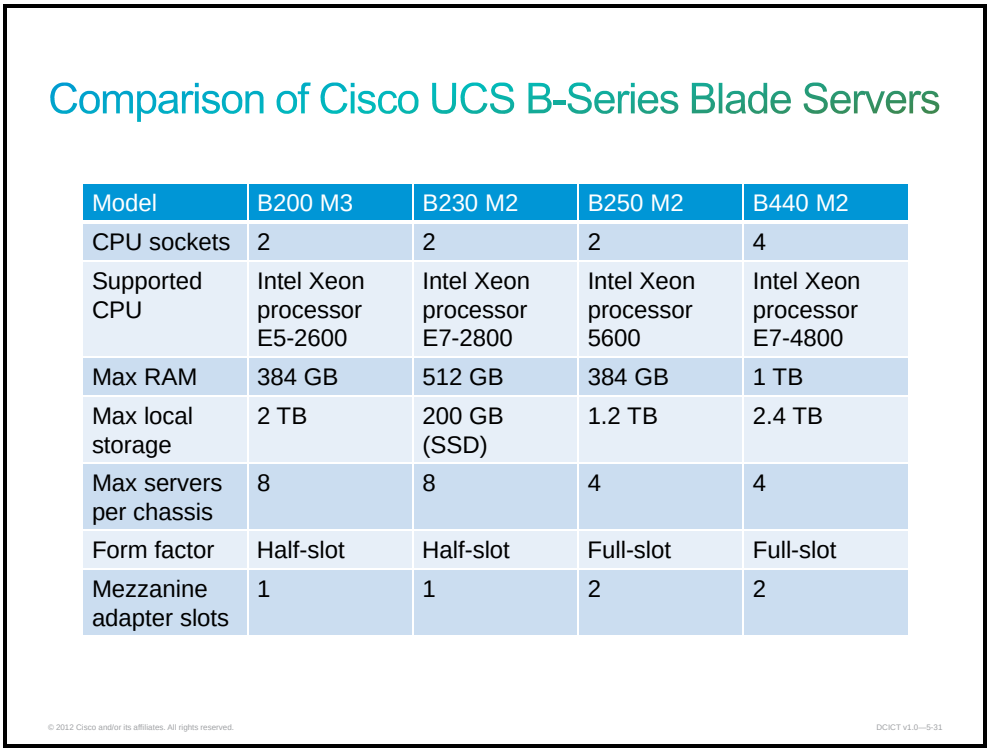
Cisco UCS B440 M2 High-Performance Blade Server Motherboard



On the motherboard, each CPU owns eight DIMM slots. The maximum memory capacity requires four installed CPUs. Two mezzanine slots allow a maximum server bandwidth of 40 Gb/s with first-generation adapters and up to 160 Gb/s with Cisco UCS Virtual Interface Card 1280 adapters.

https://t.me/learningnets

Comparison of Cisco UCS B-Series Blade Servers



This figure summarizes the comparison of Cisco UCS B-Series Blade Servers.

© 2012 Cisco Systems, Inc.

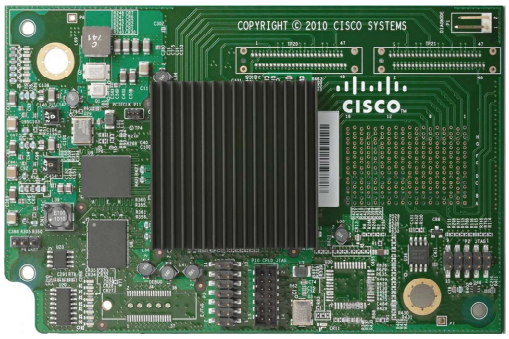
https://t.me/learningnets

Mezzanine Card Options for Cisco UCS B-Series Blade Servers

This topic describes the features of Cisco UCS B-Series Blade Server mezzanine cards.

Cisco UCS Virtual Interface Card 1280

- The Cisco UCS Virtual Interface Card 1280 is a required component to take advantage of the additional bandwidth available with the Cisco UCS 6200 fabric interconnects and Cisco UCS 2208XP I/O module.
- The Cisco UCS Virtual Interface Card 1280 provides 80-Gb/s mezzanine slot to Cisco UCS B-Series blade servers and offers hardware-based failover for nanosecond range failover in the event of fabric failure.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS Virtual Interface Card 1280 mezzanine adapter provides acceleration for the various new operational modes that server virtualization introduces. The adapter is supported by all current Cisco UCS B-Series Blade Servers. The Cisco UCS Virtual Interface Card 1280 operates in conjunction with other second-generation hardware and requires the Cisco UCS 6200 Series Fabric Interconnects and the Cisco UCS 2208XP I/O module.

In virtualized environments, this highly configurable adapter provides these benefits and features.

- Near-bare-metal I/O performance
- Ease of management
- Exceptional network visibility

The table lists the features of this adapter.

Cisco UCS Virtual Interface Card 1280 Features

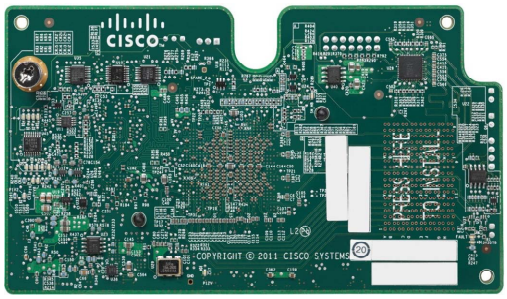
Feature	Description
Total interfaces	256
Interface type	Dynamic
Ethernet interfaces	0–256
Fibre Channel interfaces	0–256
Cisco Virtual Machine-Fabric Extender (VM-FEX)	VM-FEX or Cisco Adapter FEX
Failover handling	Hardware, no driver needed

https://t.me/learningnets

Cisco UCS Virtual Interface Card 1240

Cisco UCS Virtual Interface Card 1240

- The Cisco UCS Virtual Interface Card 1240 is a required component to take advantage of the additional bandwidth available with the Cisco UCS 6200 Series Fabric Interconnects and Cisco UCS 2204XP or 2208XP I/O module.
- The Cisco UCS Virtual Interface Card 1240 provides up to 40 Gb/s per fabric to Cisco UCS B-Series M3 blade servers and offers hardware-based failover for nanosecond range failover in the event of fabric failure.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS Virtual Interface Card 1240 mezzanine adapter provides acceleration for the various new operational modes that server virtualization introduces. This adapter is a new option that is available to Cisco UCS B-Series M3 blade servers. The Cisco UCS Virtual Interface Card 1240 operates in conjunction with other second-generation hardware and requires the Cisco UCS 6200 Series Fabric Interconnects and the Cisco UCS 2208XP I/O module.

In virtualized environments, this highly configurable adapter provides these benefits and features.

- Near-bare-metal I/O performance
- Ease of management
- Exceptional network visibility

The table lists the features of this adapter.

Cisco UCS Virtual Interface Card 1240 Features

Feature	Description
Total interfaces	256
Interface type	Dynamic
Ethernet interfaces	0–256
Fibre Channel interfaces	0–256
VM-FEX	VM-FEX or Adapter FEX
Failover handling	Hardware, no driver needed

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS M81KR Virtual Interface Card

Cisco UCS M81KR Virtual Interface Card

- The Cisco UCS M81KR CNA allows the creation of up to 128 10 Gigabit Ethernet or Fibre Channel interfaces on each fabric.
- Virtual Fiber Channel and 10 Gigabit Ethernet interfaces appear as physical PCIe devices to the operating system or hypervisor.
- This adapter offers hardware-based failover for nanosecond range failover in the event of fabric failure.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS M81KR mezzanine adapter provides acceleration for the various new operational modes that are introduced by server virtualization. It is supported by all current Cisco UCS B-Series Blade Servers. The Cisco UCS M81KR is fully supported with generation 1 and 2 hardware.

In virtualized environments, this highly configurable adapter provides the following benefits and features:

- Near-bare-metal I/O performance
- Ease of management
- Exceptional network visibility

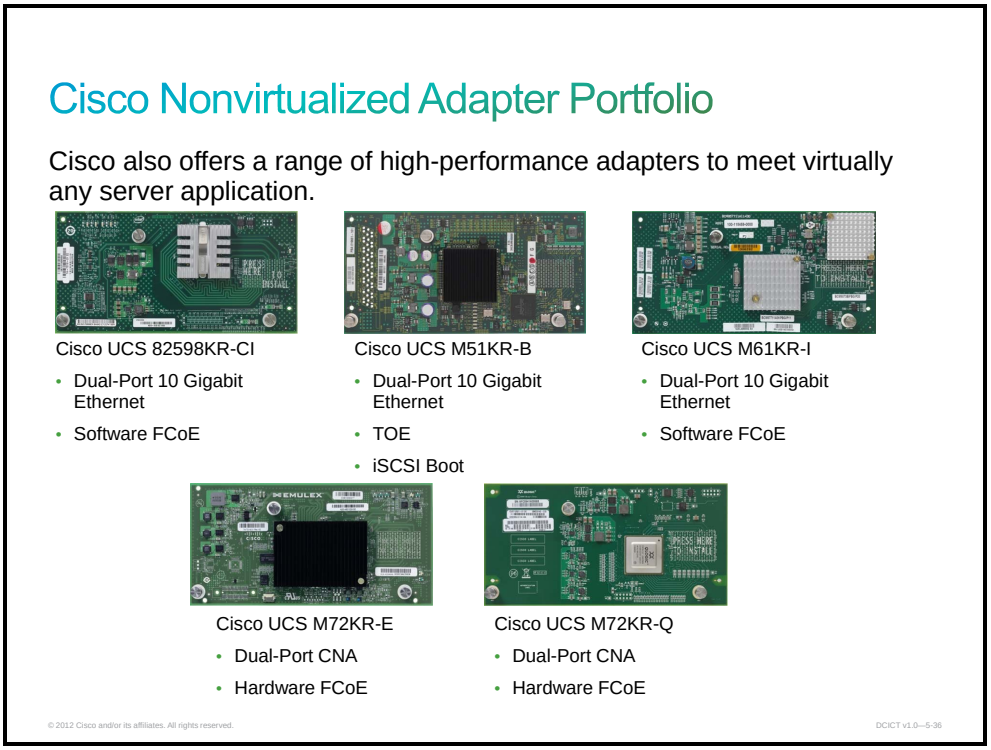
The table lists the features of this blade server.

Cisco UCS Virtual Interface Card M81KR Features

Feature	Description
Total interfaces	128
Interface type	Dynamic
Ethernet interfaces	0–128
Fibre Channel interfaces	0–128
VM-FEX	VM-FEX or Adapter FEX
Failover handling	Hardware, no driver needed

https://t.me/learningnets

Cisco UCS Nonvirtualized Adapter Portfolio



The table lists Cisco UCS nonvirtualized adapters and their features.

Cisco UCS Nonvirtualized Adapter Features

	Cisco UCS 82598KR-CI	Cisco UCS M51KR-B	Cisco UCS M61KR-I	Cisco UCS M72KR-E	Cisco UCS M72KR-Q
Ethernet interfaces	2	2	2	2	2
Fibre Channel interfaces	0	0	0	2	2
Failover handling	Software network interface card (NIC) bonding driver	Software NIC bonding driver	Software NIC bonding driver	Software NIC bonding driver	Software NIC bonding driver

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Memory Population Guidelines for Cisco UCS B-Series Blade Servers

This topic describes the best practices for populating DRAM in Cisco UCS B-Series Blade Servers.

General Memory Population Guidelines

- Take appropriate ESD precautions.
- Fill all slots in a memory channel.
- Do not mix DIMMs of various speeds (lowest speed is always used).
- Refer to the model-specific Cisco UCS B-Series blade server Installation and Service Note for detailed specifications.



© 2012 Cisco and/or its affiliates. All rights reserved.

All memory components in Cisco UCS B-Series Blade Servers are subject to damage or destruction by ESD. Take precautions against ESD by installing memory at a grounded workstation. Always wear an ESD wrist strap to dissipate any charge on your body.

The Intel Nehalem memory architecture comprises three channels that form a logical DIMM. Memory interleaving takes place most efficiently when all three channels in a logical DIMM are filled.

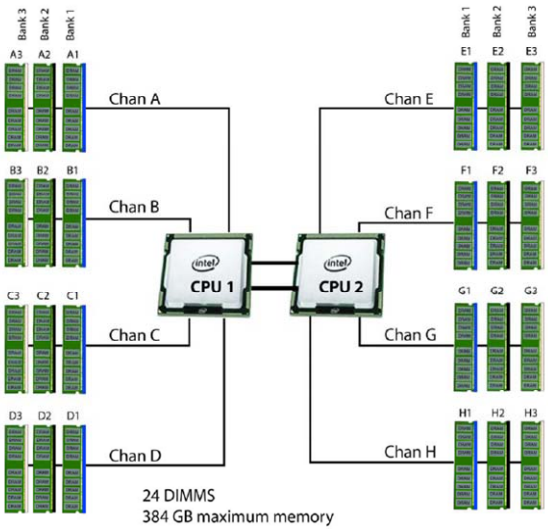
If you mix DIMMs of different speeds in a Cisco UCS B-Series server, the memory bus operates at the lowest common denominator. If there is a mix of 1066-MHz and 1333-MHz modules in a blade server, then the memory bus operates at 1066 MHz.

https://t.me/learningnets

Cisco UCS B200 M3 Blade Server Memory Population

Cisco UCS B200 M3 Blade Server Memory Architecture

- 24 DIMM slots
- Up to 3 DIMMs per channel (DPC)
- Minimum 1 DIMM per CPU (2 DIMMs per server)
- 1600-MHz speed supported up to 2 DPC
- Above 2 DPC (more than 16 DIMMs) speed drops to 1066-MHz



© 2012 Cisco and/or its affiliates. All rights reserved.

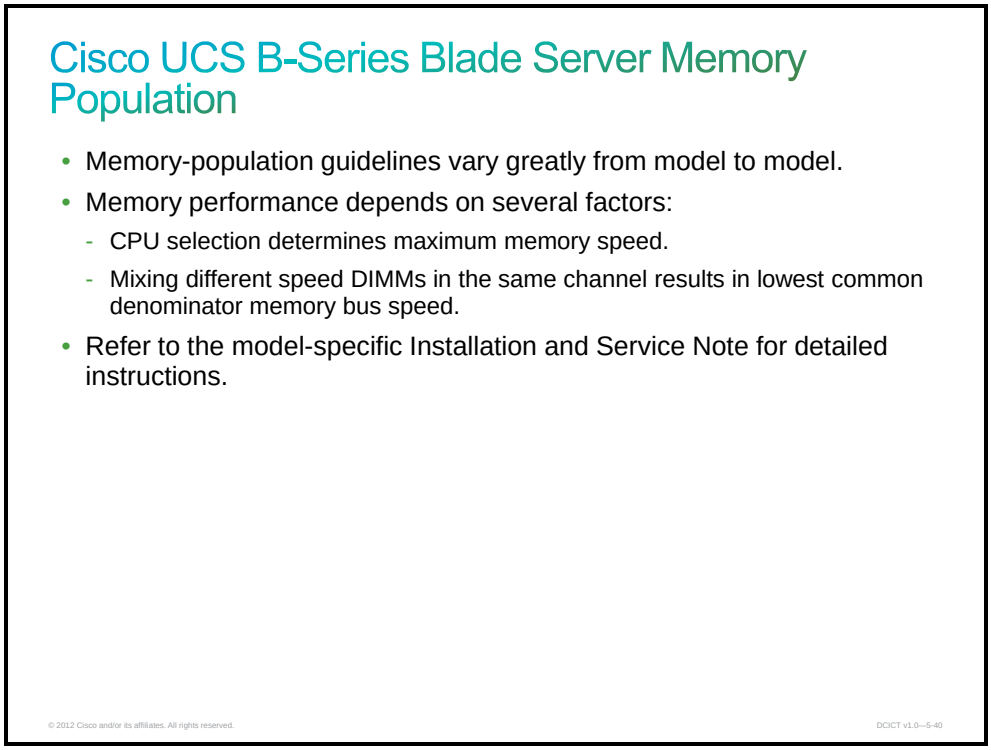
The new Intel E5-2600 chipset allows four memory channels per socket with three DIMMs per channel. Like in any CPU architecture, optimal configurations maximize memory density and speed. At the writing of this course, the optimal configuration for memory density and speed maximization is 192-GB RAM at 1600 MHz.

For detailed specifications, compatibility, and performance data, refer to the *Cisco UCS B200 Blade Server Installation and Service Note* at http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/blade.html.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS B-Series Blade Server Memory Population



The memory-population guidelines for Cisco UCS B-Series Blade Servers vary greatly from one model to another. Failure to populate memory properly can significantly reduce performance or cause a blade server to fail power-on self-test (POST).

Refer to these resources for memory-population guidelines for detailed specifications, compatibility, and performance data:

- *Cisco UCS B200 Blade Server Installation and Service Note* at http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/blade.html
- *Cisco UCS B230 Blade Server Installation and Service Note* athttp://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/B230.html
- *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note* at http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/fullblade.html
- *Cisco UCS B440 High Performance Blade Server Installation and Service Note* at http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/quadblade.html

You will also benefit from these technical documents:

- *Cisco UCS Manager B-Series Troubleshooting Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/ts/guide/UCSTroubleshooting.html
- *UCS C-Series Servers Troubleshooting Tips* at http://www.cisco.com/en/US/products/ps10493/products_configuration_example09186a0080b10d5e.shtml

https://t.me/learningnets

Summary

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

Summary

- Cisco UCS 6100 Series Fabric Interconnects offer up to 80 Gb/s per chassis, while the Cisco UCS 6200 Series Fabric Interconnects offer up to 160 Gb/s to each chassis.
- The Cisco UCS 5108 Blade Server Chassis supports eight half-slot blades, four full-slot blades, or a mix of half-slot and full slot blades.
- The Cisco UCS B200 M3 Blade Server offers dual-socket Intel Xeon E5-2800 series CPUs and up to 384-GB RAM.
- The Cisco UCS B230 M2 Blade Server offers dual-socket Intel E7-2800 series CPUs and up to 512-GB RAM.
- The Cisco UCS B250 M2 Extended Memory Blade Server offers dual-socket Intel Xeon 5600 series CPUs and up to 384-GB RAM.
- The Cisco UCS B440 M2 High-Performance Blade Server offers quad-socket Intel E7-4800 series CPUs and up to 512-GB RAM.
- Cisco UCS B-Series Blade Servers have the choice of nine mezzanine adapters to suit virtually any use case.
- The memory performance of Cisco UCS B-Series Blade Servers depends on following best practices for DIMM population.

© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Describing the Cisco UCS C-Series Product Family

Overview

The purpose of this lesson is to describe and verify the components of the Cisco Unified Computing System (UCS) C-Series product family.

Objectives

Upon completing this lesson, you will be able to describe the components of the Cisco UCS C-Series product family. You will be able to meet these objectives:

- Describe the Cisco UCS C-Series product family
- Describe the Cisco UCS C200 M2 High-Density Rack Server
- Describe the Cisco UCS C210 M2 General-Purpose Rack Server
- Describe the Cisco UCS C220 M3 Rack Server
- Describe the Cisco UCS C240 M3 Rack Server
- Describe the Cisco UCS C250 M2 Extended-Memory Rack Server
- Describe the Cisco UCS C260 M2 Rack Server
- Describe the Cisco UCS C460 M2 High-Performance Rack Server
- Describe the features of Cisco UCS C-Series network adapter cards
- Describe RAID adapter options for Cisco UCS C-Series Rack Servers
- Describe the best practices for populating DRAM in Cisco UCS C-Series Rack Servers

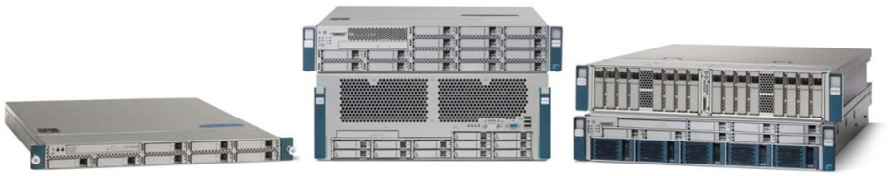
https://t.me/learningnets

Cisco UCS C-Series Product Family

This topic describes the Cisco UCS C-Series product family.

Cisco UCS C-Series Rack Server Family

- The Cisco UCS C-Series rack server family offers an entry point into unified computing.
- Cisco UCS C-Series Rack Servers offer the flexibility of standalone management and configuration or integration with Cisco UCS Manager and Cisco UCS B-Series Blade Servers.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C-Series product family begins with the cost-effective Cisco UCS C200 M2 High-Density Rack Server and Cisco UCS C210 M2 General-Purpose Rack Server and scales to the massively capable Cisco UCS C460 servers. Cisco UCS C-Series Rack Servers bring Cisco UCS virtualization benefits into a familiar rack-mount server form factor that can operate standalone or integrate with Cisco UCS Manager.

Note	The Cisco UCS C-Series deployment options are discussed in more detail in the lesson "Describing Cisco UCS Manager Operations."

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS C200 M2 High-Density Rack Server

This topic describes the Cisco UCS C200 M2 High-Density Rack Server.

Cisco UCS C200 M2 High-Density Rack Server

The Cisco UCS C200 M2 rack server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 192 GB
Hard drives	Up to 8
CPU	1 or 2 Intel Xeon 5600 series
Built-in RAID	RAID 0 or 1 (SATA Only)
Optional RAID	0, 1, 5, 6, 10, 50, and 60[1]
Adapters slots	2 half-length PCIe x8 slots: 1 full-height and 1 low-profile slot

[1]Not available for the 4-disk C200 M2 that has RAID 0, 1, 5, 6 but not RAID 10, 50, 60



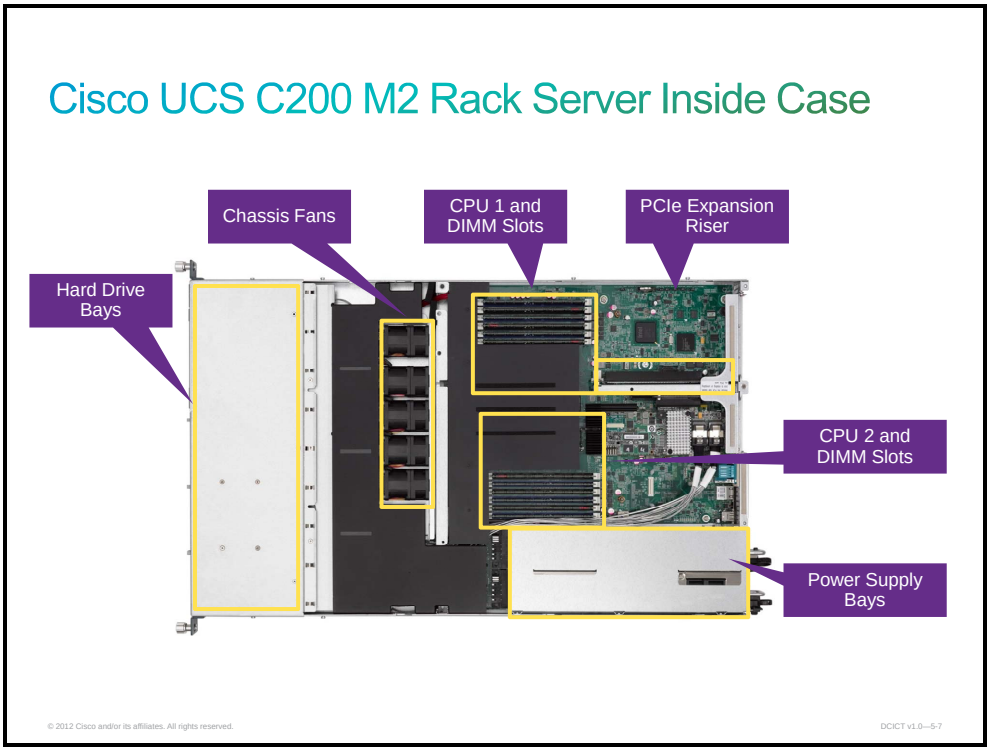
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C200 M2 High-Density Rack Server is a 1-rack unit (RU), 2-socket rack server with substantial throughput and scalability. The table shows the core features of this server.

Cisco UCS C200 M2 High-Density Rack Server Core Features

Parameter	Description
Memory capacity	12 DIMMs; up to 192 GB
Memory size and speed	4-, 8-, and 16-GB double data rate 3 (DDR3); 1066 MHz and 1333 MHz
Internal disk drive	Two 2.5-in. small form-factor (SFF) Serial Attached SCSI (SAS) or 15-mm Serial Advanced Technology Attachment (SATA) solid state drive (SSD)
Maximum internal storage	Up to 1.2 TB
Integrated RAID	0,1

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these field-replaceable units (FRUs) require removal of the server from the rack and the disconnection of power:

- CPU
- DIMM
- Fan modules
- PCI Express (PCIe) adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

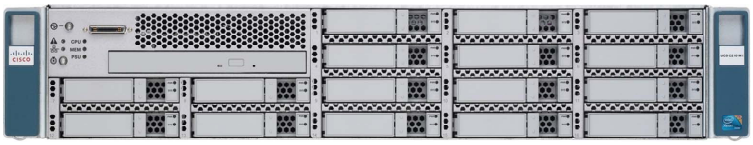
Cisco UCS C210 M2 General-Purpose Rack Server

This topic describes the Cisco UCS C210 M2 General-Purpose Rack Server.

Cisco UCS C210 M2 General-Purpose Rack Server

The Cisco UCS C210 M2 rack server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 192 GB
Hard drives	Up to 16
CPU	1 or 2 Intel Xeon 5600 series
Built-in RAID	0 and 1 (5 SATA drives only)
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	5 full-height PCIe x8 slots: 2 full-length and 3 half-length slots



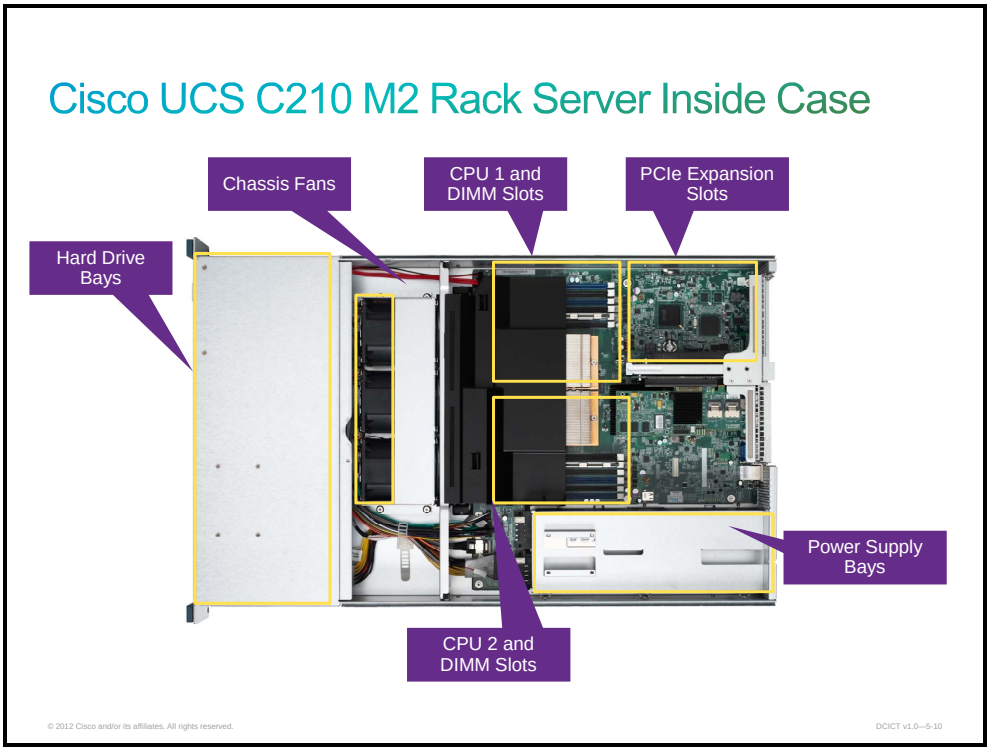
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C210 M2 General-Purpose Rack Server is a 2-RU, 2-socket rack server with substantial throughput and scalability. The table shows the core features of this server.

Cisco C210 M2 General-Purpose Rack Server Core Features

Feature	Description
Memory capacity	12 DIMMs; up to 192 GB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz and 1333 MHz
Internal disk drive	Two 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 16 TB
Integrated RAID	0,1

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS C220 M3 Rack Server

This topic describes the Cisco UCS C220 M3 Rack Server.

Cisco UCS C220 M3 Rack Server

The Cisco UCS C220 M3 Rack Server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 256 GB
Hard drives	Up to 8
CPU	2 Intel Xeon E5-2600 series
Built-in RAID	None
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	2 PCIe third-generation slots: 1 x8 half-height and half-length; 1 x16 full height and three-quarters length



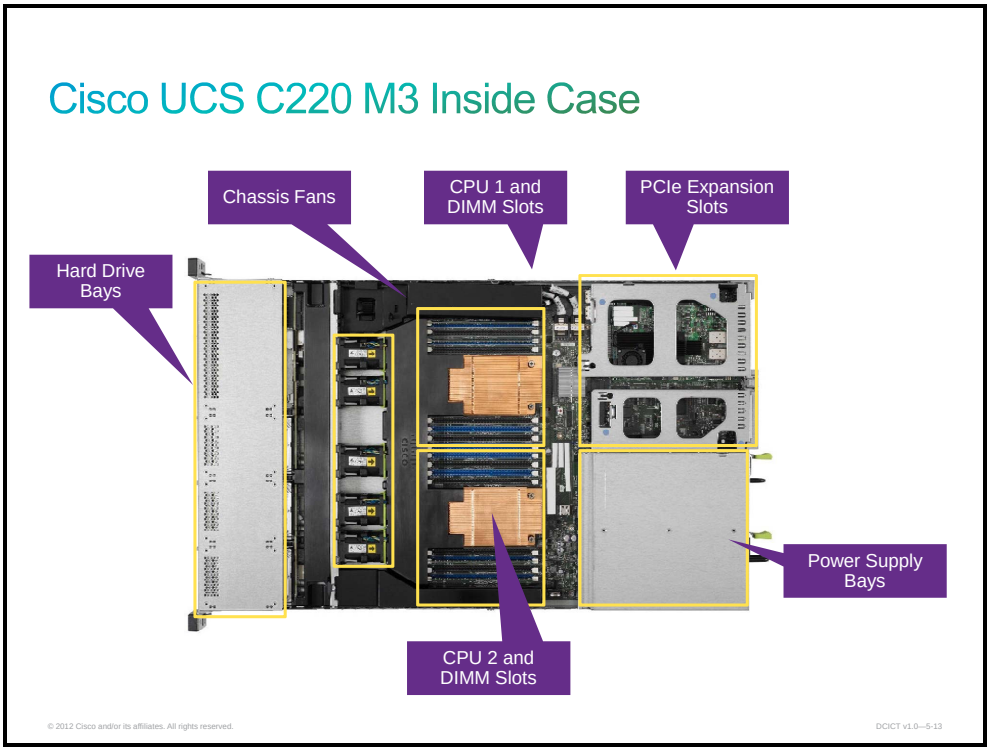
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C220 M3 Rack Server is a 2-RU, 2-socket rack server with substantial throughput and scalability. The table shows the core features of this server.

Cisco C220 M3 Rack Server Core Features

Feature	Description
Memory capacity	16 DIMMs; up to 256 GB
Memory size and speed	DDR3; 1600 MHz
Internal disk drive	Eight 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 8 TB
Integrated RAID	None

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS C240 M3 Rack Server

This topic describes the Cisco UCS C240 M3 Rack Server.

Cisco UCS C240 M3 Rack Server

The Cisco UCS C240 M3 Rack Server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 384 GB
Hard drives	Up to 24
CPU	1 or 2 Intel Xeon E5-2600 series
Built-in RAID	None
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	2 PCIe third-generation full-height x16 slots: 1 half-length and one three-quarters length; 2 PCIe third-generation x8 slots: both full height and one half-length; 1 PCIe third-generation half-height and half-length x8 slots



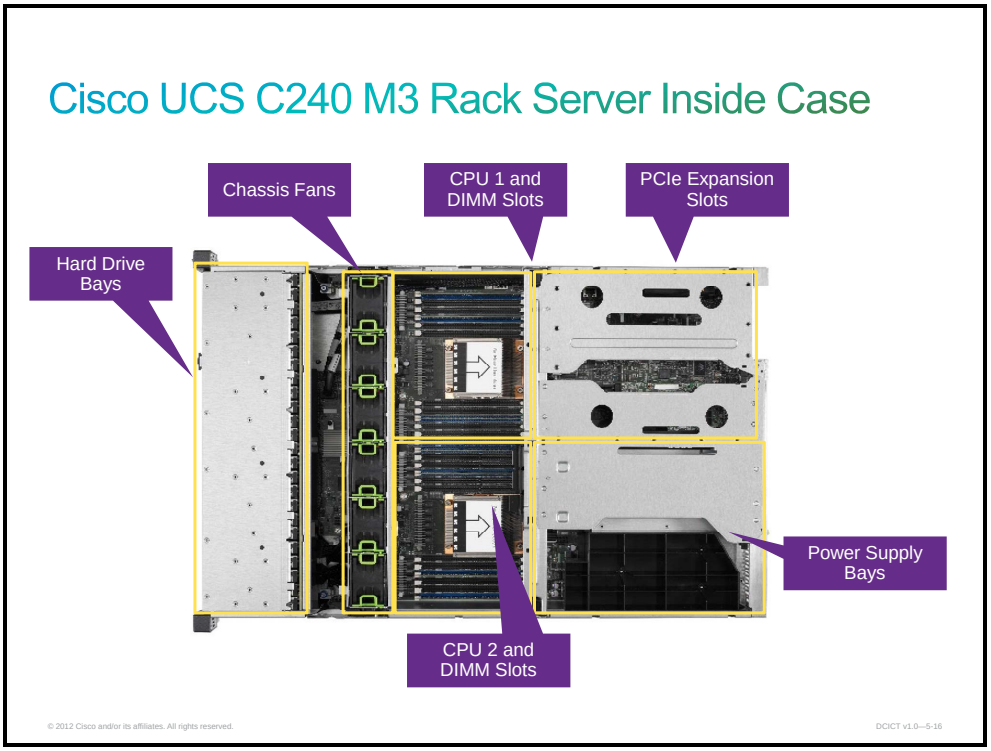
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C240 M3 Rack Server is a 2-RU, 2-socket rack server with substantial throughput and scalability. The table shows the core features of this server.

Cisco C240 M3 Rack Server Core Features

Feature	Description
Memory capacity	24 DIMMs; up to 384 GB
Memory size and speed	DDR3; 1600 MHz
Internal disk drive	Twenty-four 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 24TB
Integrated RAID	None

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS C250 M2 Extended-Memory Rack Server

This topic describes the Cisco UCS C250 M2 Extended-Memory Rack Server.

Cisco UCS C250 M2 Rack Server

The Cisco UCS C250 M2 rack server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 384 GB
Hard drives	Up to 8
CPU	1 or 2 Intel Xeon 5600 series
Built-in RAID	None
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	5 PCIe slots: 3 low-profile, half-length x8 slots; 2 full-height, half-length x16 slots



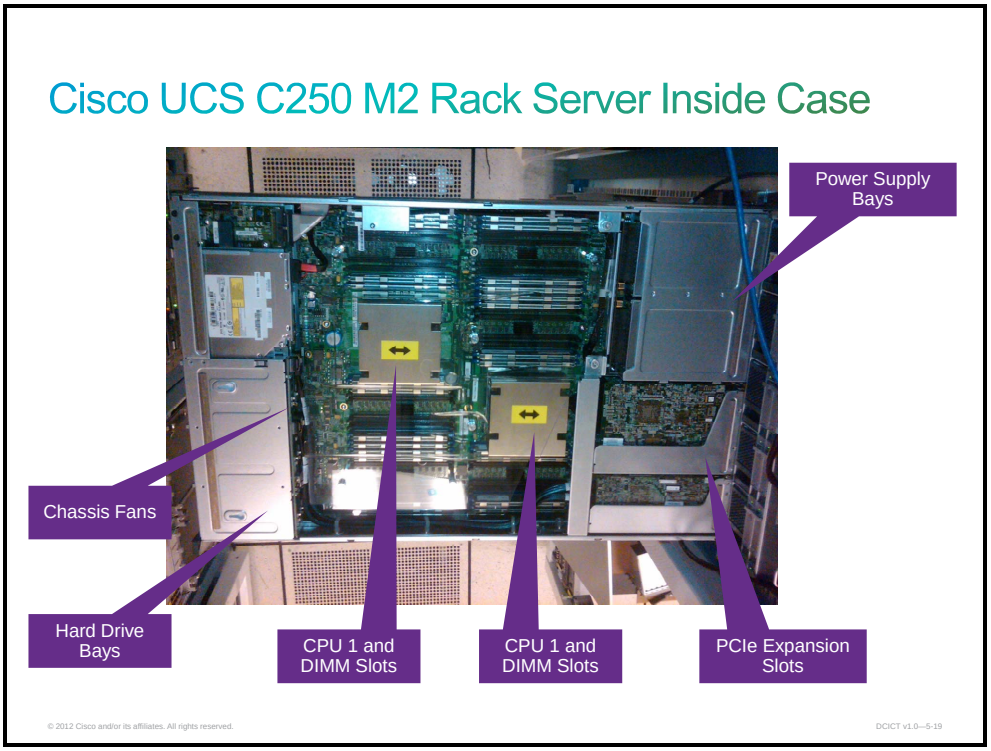
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C250 M2 Extended-Memory Rack Server is a 2-RU, 2-socket rack server with expanded memory capabilities to support memory-constrained applications such as virtual desktop infrastructure (VDI). The table shows the core features of this server.

Cisco UCS C250 M2 Extended-Memory Rack Server Core Features

Parameter	Description
Memory capacity	48 DIMMs; up to 384 GB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz and 1333 MHz
Internal disk drive	Eight 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 8 TB
Integrated RAID	None

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS C260 M2 Rack Server

This topic describes the Cisco UCS C260 M2 Rack Server.

Cisco UCS C260 M2 Rack Server

The Cisco UCS C260 M2 Rack Server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 1 TB
Hard drives	Up to 16
CPU	1 or 2 Intel Xeon E7-2800 series
Built-in RAID	None
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	3 low-profile, half-length x8 slots; 2 full-height, half-length x16 slots; 1 low-profile, half-length x4



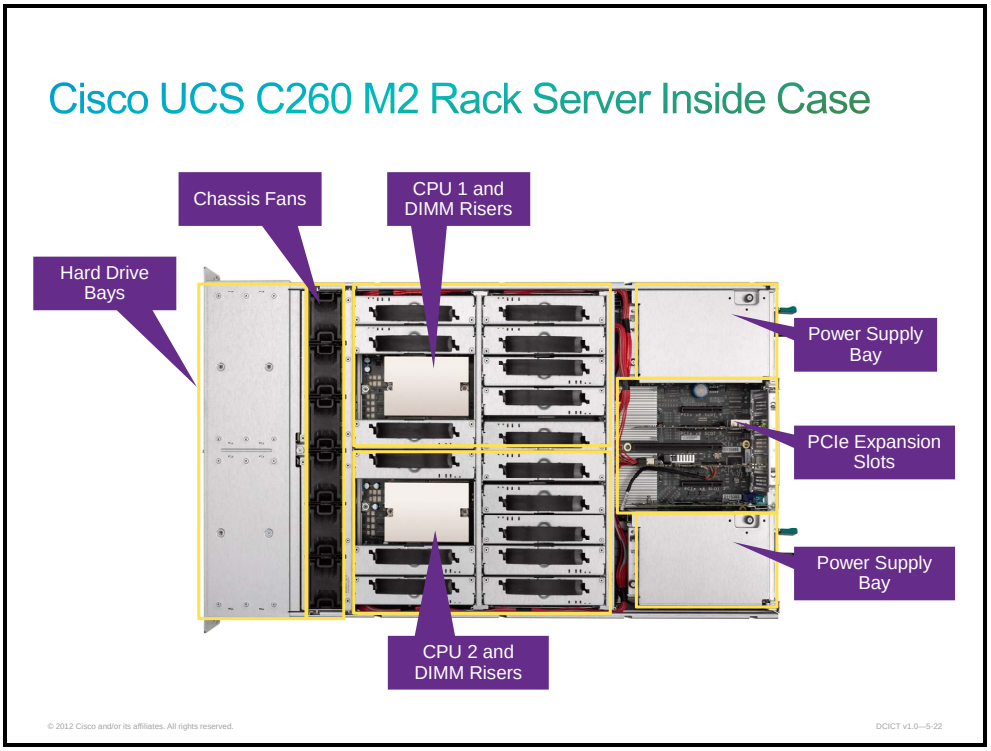
© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C260 M2 Rack Server is a 2-RU, 2-socket rack server with expanded memory capabilities to support memory-constrained applications such as VDI. The table shows the core features of this server.

Cisco UCS C260 M2 Rack Server Core Features

Feature	Description
Memory capacity	32 or 64 DIMMs (depending on choice of DIMM riser); up to 1TB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz and 1333 MHz
Internal disk drive	Sixteen 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 16 TB
Integrated RAID	None

https://t.me/learningnets



Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- DIMM risers
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

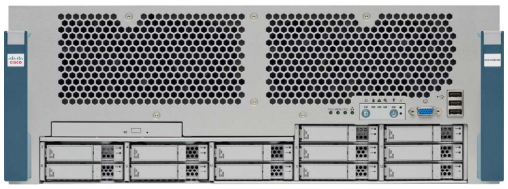
Cisco UCS C460 M2 High-Performance Rack Server

This topic describes the Cisco UCS C460 M2 High-Performance Rack Server.

Cisco UCS C460 M2 High-Performance Rack Server

The Cisco UCS C460 M2 rack server balances performance and density for virtualization and other mainstream data center workloads.

Specification	Description
DRAM	Up to 1 TB
Hard drives	Up to 12
CPU	2 or 4 Intel Xeon E7-4800 series
Built-in RAID	None
Optional RAID	0, 1, 5, 6, 10, 50, and 60
Adapters slots	10 full-height PCIe slots: 4 half-length slots, 6 three-quarter length slots



© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C460 M2 High-Performance Rack Server is a 4-RU, 4-socket rack server with expanded memory capabilities to support memory and CPU-constrained applications such as databases, high-performance computing (HPC), and VDI. The table shows the core features of this server.

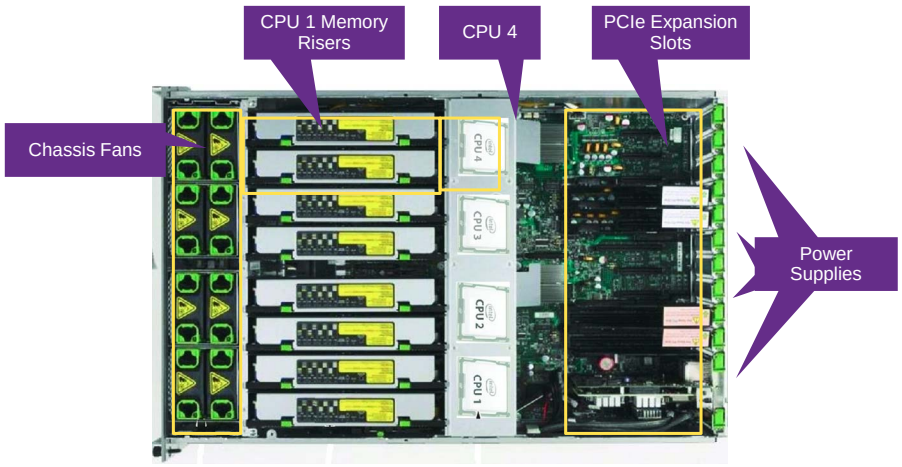
Cisco UCS C260 M2 Core Features

Feature	Description
Memory capacity	64 DIMMs; up to 1TB
Memory size and speed	4-, 8-, and 16-GB DDR3; 1066 MHz and 1333 MHz
Internal disk drive	Twelve 2.5-in. SFF SAS or 15-mm SATA SSD
Maximum internal storage	Up to 12 TB

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS C460 M2 Inside Case



© 2012 Cisco and/or its affiliates. All rights reserved.

Hard disks and power supplies can be hot-swapped without removing the server from the rack. The installation and removal of these FRUs require removal of the server from the rack and disconnection of power:

- CPU
- DIMM
- Fan modules
- DIMM risers
- PCIe adapter riser
- PCIe adapters

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco USC C-Series Comparison Summary

Comparison of Cisco UCS C-Series Rack Servers

Model	C200 M2	C210 M2	C220 M3	C240 M3
CPU sockets	2	2	2	2
Supported CPU	Intel Xeon processor 5600	Intel Xeon processor 5600	Intel Xeon processor E5-2600	Intel Xeon processor E5-2600
Max RAM	192 GB	192 GB	256 GB	384 GB
Max local storage	8 TB	16 TB	8 TB	24 TB
Form factor	1 RU	2 RU	2 RU	2 RU

© 2012 Cisco and/or its affiliates. All rights reserved.

Comparison of Cisco UCS C-Series Rack Servers (Cont.)

Model	C250 M2	C260 M2	C460 M2
CPU sockets	2	2	4
Supported CPU	Intel Xeon processor 5600	Intel Xeon processor E7-2800	Intel Xeon processor E7-4800
Max RAM	384 GB	1 TB	1 TB
Max local storage	8 TB	16 TB	12 TB
Form factor	2 RU	2 RU	4 RU

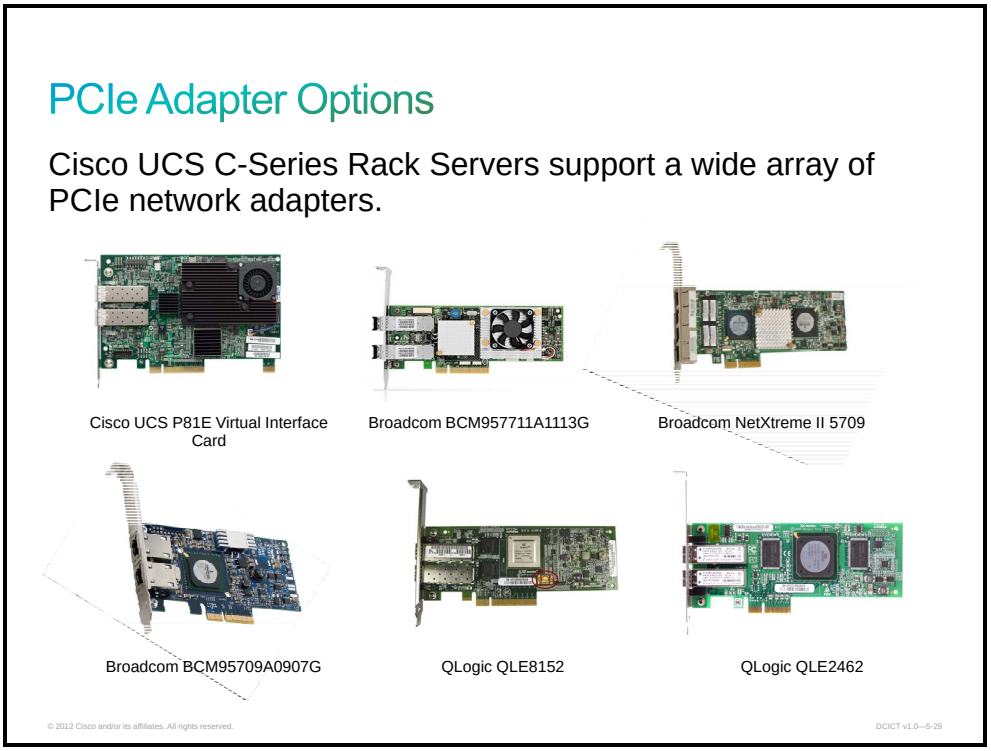
© 2012 Cisco and/or its affiliates. All rights reserved.

These figures provide a quick reference for comparing the specifications of the Cisco UCS C-Series rack server product portfolio.

https://t.me/learningnets

PCIe Adapter Options for Cisco UCS C-Series Rack Servers

This topic describes the features of Cisco UCS C-Series adapter cards.



Cisco UCS C-Series Rack Servers enjoy a wide range of PCIe adapters for Ethernet, Fibre Channel, and Cisco Network Assistant for Fibre Channel over Ethernet (FCoE) communication.

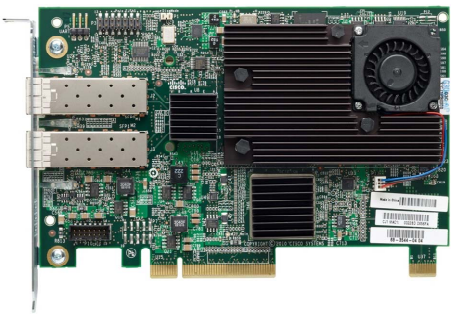
Refer to the data sheets at http://www.cisco.com/en/US/products/ps10493/products_data_sheets_list.html to access more detailed product specifications for supported adapters.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS P81E Virtual Interface Card

- The Cisco UCS P81E Virtual Interface Card is a virtualization-optimized FCoE PCIe 2.0 x8 10-Gb/s adapter.
- The VIC is a dual-port 10 Gigabit Ethernet PCIe adapter that can support up to 128 PCIe standards-compliant virtual interfaces.
- In addition, the Cisco UCS P81E card can support NIV and Cisco VN-Link technology.



© 2012 Cisco and/or its affiliates. All rights reserved.

Unique to Cisco UCS, the Cisco UCS P81E Virtual Interface Card was designed for virtualized environments and for data centers that want to reduce total cost of ownership (TCO) by reducing network interface cards (NICs), host bus adapters (HBAs), cabling, and switches.

The Cisco UCS P81E Virtual Interface Card can present as many as 128 virtual interfaces to a given operating system or hypervisor. Initially, the Cisco UCS P81E Virtual Interface Card supports as many as 2 Fibre Channel and 16 Ethernet devices. The number of virtual interfaces increases over time. With the Cisco UCS P81E Virtual Interface Card, deployment of applications that require or benefit from multiple Ethernet and Fibre Channel interfaces is no longer constrained by the available physical adapters.

To an operating system or a hypervisor that runs on a Cisco UCS C-Series rack mount server, the virtual interfaces appear as physical devices on the PCIe bus. In a virtualized environment, Cisco Virtual Machine-Fabric Extender (VM-FEX) technology allows virtual links to be centrally configured and managed, without the complexity that traditional approaches impose.

https://t.me/learningnets

Comparison of PCIe Adapters for Cisco UCS C-Series Rack Servers

Model	Interfaces	Interface Type	Ethernet Interfaces	Fibre Channel Interfaces	NIC Teaming
Cisco UCS P81E	128	Dynamic	128	128	*Hardware, no driver needed
Emulex OneConnect Universal CNA	4	Fixed	2	2	Software, NIC bonding driver
QLogic QLE8152 CNA		Fixed	2	2	Software, NIC bonding driver
Intel Gigabit ET2		Fixed	4	0	Software, NIC bonding driver
Broadcom NetXtreme II 5709		Fixed	4	0	Software, NIC bonding driver
Emulex LPe11002	2	Fixed	0	2	—
QLogic QLE2462	2	Fixed	0	2	—

*Hardware failover requires integration with Cisco UCS Manager.

© 2012 Cisco and/or its affiliates. All rights reserved.

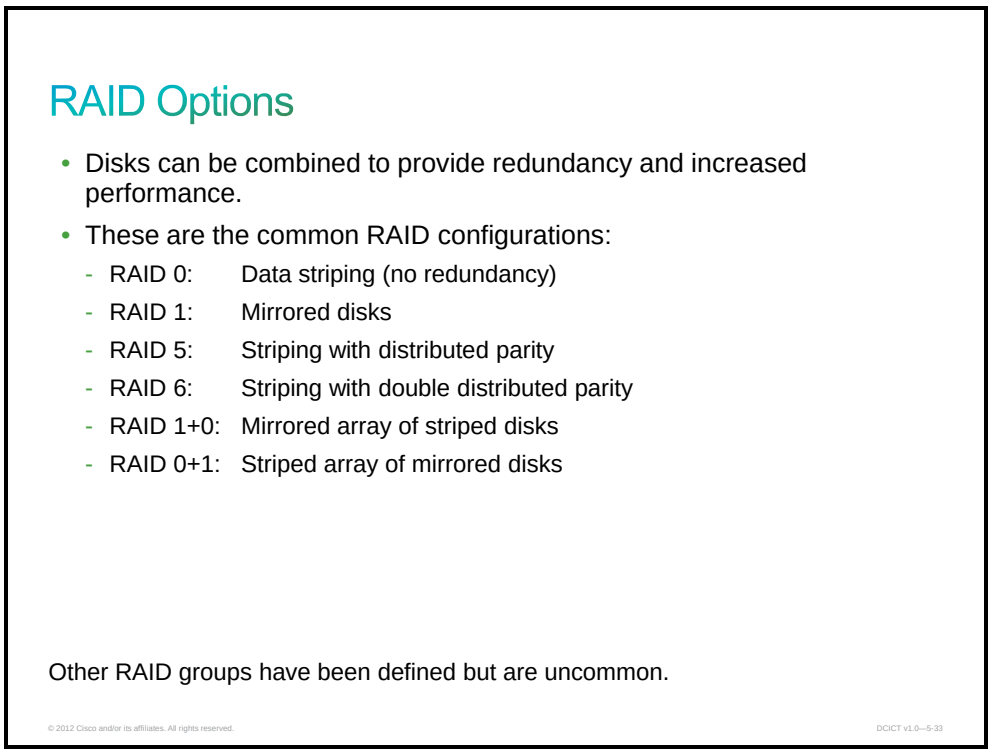
The figure summarizes the PCIe adapters that are available for Cisco UCS C-Series Rack Servers. Refer to the previously mentioned data sheets for detailed specifications for each adapter.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

RAID Adapter Options for Cisco UCS C-Series Rack Servers

This topic describes the RAID adapter options for Cisco UCS C-Series Rack Servers.



To provide redundancy, disks can be combined into arrays in several ways, which are referred to as RAID levels.

This list reflects the most commonly implemented RAID levels:

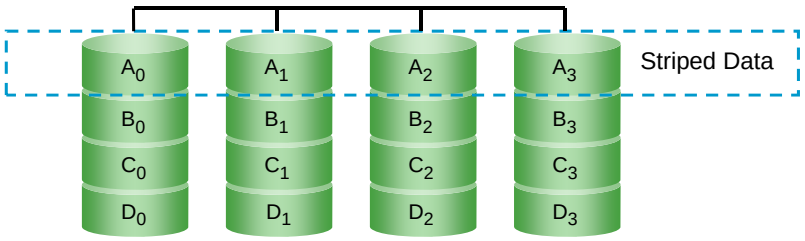
- **RAID 0:** RAID 0 provides data striping and good performance but no redundancy.
- **RAID 1:** RAID 1 provides mirrored disks and offers redundancy, but at twice the cost.
- **RAID 5:** RAID 5 provides data striping with distributed parity and good read performance, but writes are slower because the parity must be written for each stripe. Data on one failed disk might be recovered from the parity disk.
- **RAID 6:** RAID 6 provides data striping with double distributed parity and good read performance, but writes are slower because the parity must be written twice, using different encoding schemes. RAID 6 is used when the disk drives are more unreliable, so data can be recovered from two failed disks in a stripe.
- **RAID 1+0:** RAID 1+0 provides mirrored array of striped disks and combines the performance of RAID 0 with the redundancy of RAID 1. RAID 1+0 mirrors two RAID 0 stripes. If one disk fails anywhere, then the array becomes RAID 0 with no redundancy.
- **RAID 0+1:** RAID 0+1 provides striped array of mirrored disks and combines the performance of RAID 0 with the redundancy of RAID 1. RAID 0+1 mirrors each disk, then stripes all the mirrors. RAID 0+1 can survive multiple disk failures if they are in different mirrors.

https://t.me/learningnets

RAID Level 0: Data Striping

RAID 0 is really “AID,” as there is no redundancy.

- Data is distributed across the disks.
- Each disk contains only a portion of the data.
- RAID 0 does not provide any redundancy. If a disk fails, all data is lost.



© 2012 Cisco and/or its affiliates. All rights reserved.

RAID 0, or disk striping, is the process of breaking data into blocks of predetermined size and writing those blocks to separate disk drives.

The figure displays four physical drives that have been combined into an array. Each drive is sufficient in size to accommodate four stripes. If the individual blocks were 512 KB in size, then a single stripe file of 4 x 512 KB could be written across a stripe in one I/O cycle. By accessing the drives in parallel, the effective transfer rate equals the sum of the transfer rates of all the drives in the RAID set.

A major risk is associated with simple data striping: If a disk drive fails, then all the data that is stored on the array is lost. Giving up redundancy increases performance because there is no parity calculation. This lack of redundancy might be important for those seeking the best possible performance.

RAID 0 includes these advantages:

- Greatly improved I/O performance
- Extremely simple design
- Ease of implementation

RAID 0 includes these disadvantages:

- No fault tolerance
- Failure of a single drive results in the loss of the data on the entire array

Data striping (RAID 0) can be used to benefit solutions that have high-bandwidth requirements:

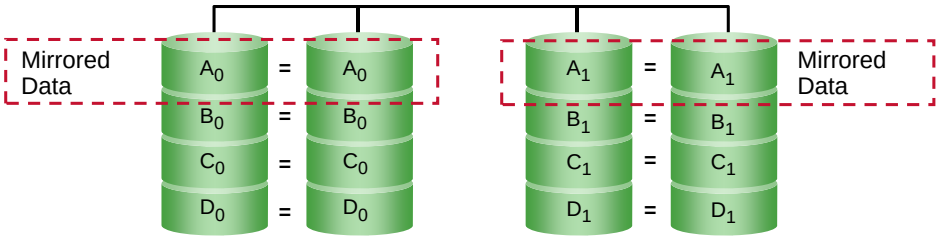
- Video production and editing
- Image editing
- Prepress applications

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

RAID Level 1: Disk Mirroring

- Data is duplicated on a second disk.
- If one disk of the mirrored set fails, data can be read from the other disk.
- RAID 1 has 100% redundancy and requires twice as much storage.



© 2012 Cisco and/or its affiliates. All rights reserved.

RAID 1, or disk mirroring, maintains identical copies of the same data on two separate disk drives. When a write occurs, the data is written to both drives. Duplication of data on the second drive can be accomplished by using either a hardware RAID controller or software.

The characteristics of data I/O do not change much with disk mirroring. When data is being written, it must be written to both drives. When a read occurs, data can be read from either drive. The ability to read the data from either drive provides redundancy in disk mirroring. If one of the drives fails, the data is still available from the other drive. RAID 1 can also employ duplexing, in which the controller card is duplicated as well as the drive, thus providing tolerance against failures of either a drive or a controller.

RAID 1 includes these advantages:

- One write or two reads are possible per mirrored pair.
- 100-percent redundancy ensures that no reconstruction of data is necessary following a disk failure. Data can simply be copied to a replacement drive.
- RAID 1 is simple.

A disadvantage of RAID 1 is that it has the highest disk overhead of all RAID levels because it requires double the number of disks.

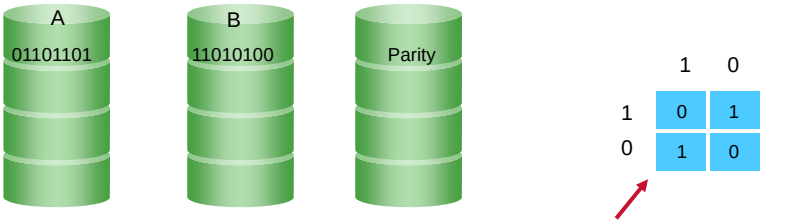
Many applications can benefit from RAID 1. This level of RAID is popular for applications that require fault tolerance but do not need the highest possible read performance. These types of applications can benefit from RAID 1:

- Accounting or other financial data
- Small database application data
- Enterprise servers

https://t.me/learningnets

Redundancy via Parity

- Instead of mirroring data, redundancy can be provided by parity blocks, which reduce the overhead of RAID 1.
- Parity is computed by an XOR operation.
- If any disk within a stripe fails, then the data from that disk can be rebuilt, using the data contained within the parity blocks and the remaining disk drives in the stripe.



If A XOR B = Parity, then A XOR Parity = B

© 2012 Cisco and/or its affiliates. All rights reserved.

Mirroring provides protection against data loss that is caused by a device failure. However, mirroring has significant overhead because two copies of the data are recorded.

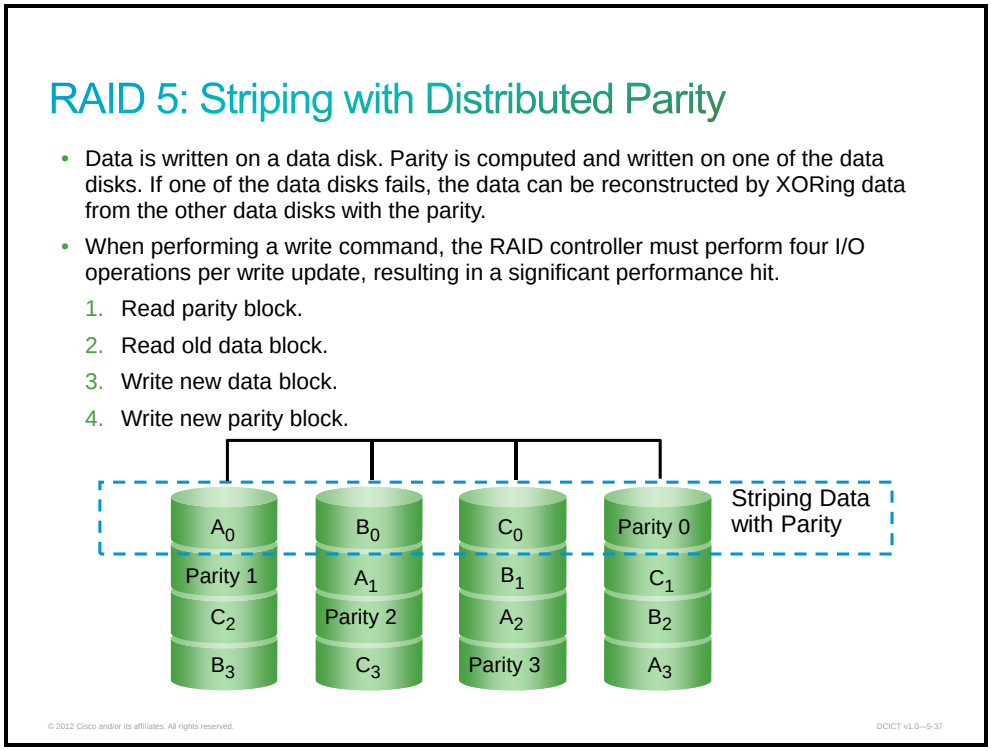
Parity is a technique that is used to provide redundancy with less overhead. Parity is computed by an exclusive OR (XOR) operation. XOR is applied to data bytes to create a parity byte, then any one data byte can be recreated by applying XOR on the remaining bytes with the parity byte.

This property is used by RAID 3, RAID 4, RAID 5, and RAID 6:

- If A XOR B XOR C = Parity, then A XOR B XOR Parity = C

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



RAID 5 uses a combination of block-level striping and distributed parity information. RAID 5 writes the data to one or more data disks in the RAID set and writes the parity information to the other disk drive in the set.

For example, a file that is represented by blocks A0, B0, and C0 is written across the first three drives of the stripe. The parity information for the file is written to the fourth drive. Data protection is provided by using the parity block.

Although every data write also causes a write to the parity disk, congestion is not a problem because the parity information is distributed among the disks.

The parity helps to reconstruct the data in the event of a disk failure:

1. The data for the first stripe is written to the disk.
2. The parity value is calculated as A0 XOR B0 XOR C0 = Parity.
3. If the second disk fails, then the data on the disk that is associated with the first stripe can be reconstructed by using the calculation of A0 XOR C0 XOR Parity0 = B0.
4. This approach results in a loss in read performance, but the data remains available to applications.

As disk prices have been reduced and capacities increased, RAID 5 has become less popular.

RAID 5 has these advantages:

- Very high read transaction rate
- Low ratio of parity disks to data disks
- Good aggregate transfer rate

RAID 5 has these disadvantages:

- Medium write performance
- Difficult and inefficient rebuilding of data if a disk fails
- Individual block transfer rates that are the same as a single disk

https://t.me/learningnets

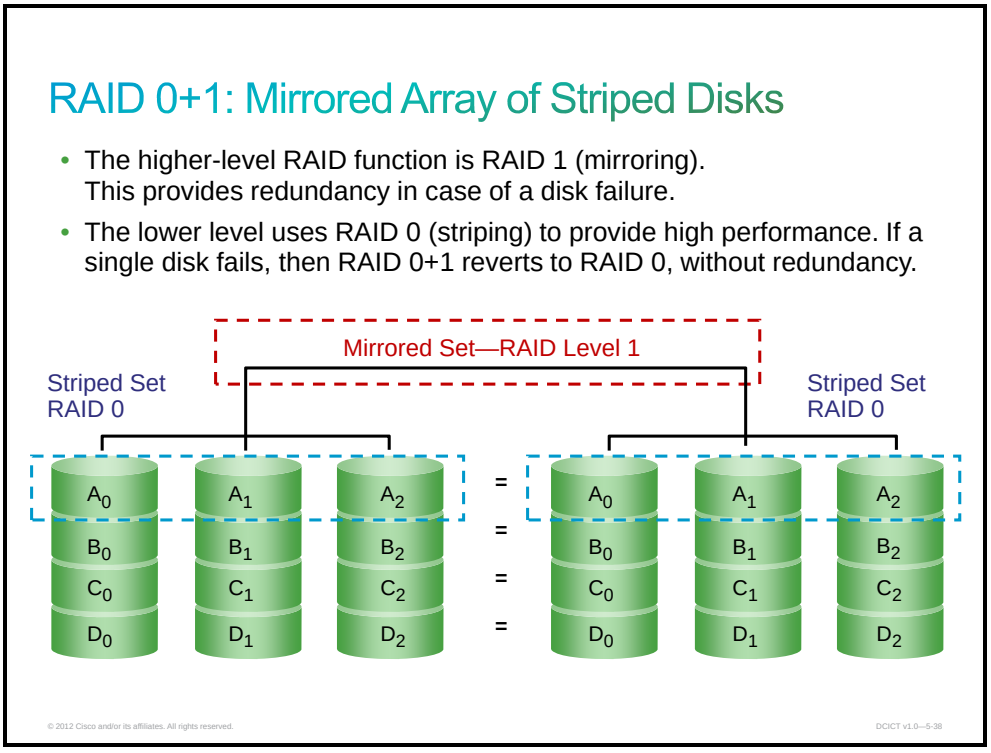
Several applications can benefit from RAID 5. These applications tend to be transaction processing-intensive or relational-database applications. These types of applications can benefit from RAID 5:

- Web servers
- Email servers
- News servers
- File and application servers
- Database servers

I/O multiplication is associated with RAID 5, RAID 4, and RAID 6. These levels of RAID incur additional overhead on write updates because of the need to update parity data. When performing a **write** command, this series of I/O operations must take place on the RAID controller:

1. Read the old data from the data disk
2. Write the new data to the data disk
3. Apply XOR to the old data with the new data, to compute an intermediate parity
4. Read the old parity from the parity disk
5. Apply XOR to the old parity with the intermediate parity that was just computed
6. Write the new parity to the parity disk

This series results in four I/O operations per write update.



RAID 0+1 is a combination of RAID 0 and RAID 1. The configuration contains a stripe that is mirrored to a second stripe. This configuration allows you to use striping to provide high I/O performance and to use mirroring for fault tolerance.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

RAID 0+1 includes these advantages:

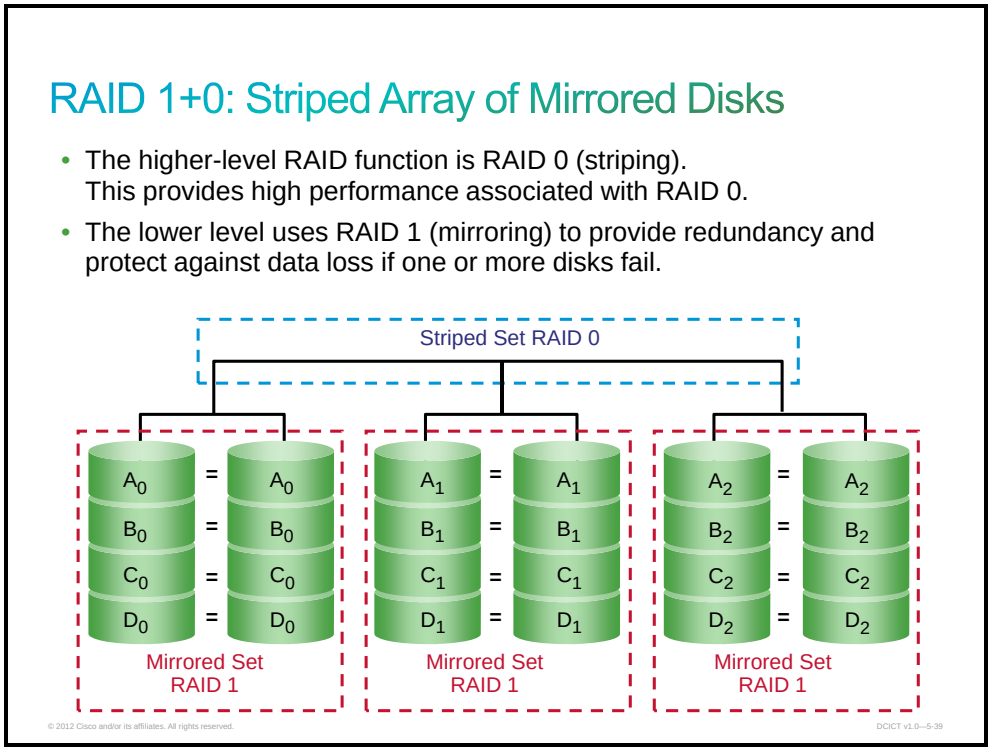
- Fault tolerance
- High I/O rates

RAID 0+1 includes these disadvantages:

- Need for two sets of drives, making the cost of implementation high
- Limited scalability

RAID 0+1 can benefit applications that require both high performance and reliability. These types of applications can benefit from RAID 0+1:

- Imaging applications
- General file servers



RAID 1+0 is a two-level RAID configuration. The highest-level of RAID is RAID 0, known as striping. Instead of striping the data across individual disks, the lower level uses RAID 1 (mirroring) to provide redundancy and to protect against data loss if a disk fails. In this configuration, the data is simultaneously mirrored and striped. This configuration is used to combine high reliability via mirroring with high I/O performance via striping.

RAID 1+0 includes these advantages:

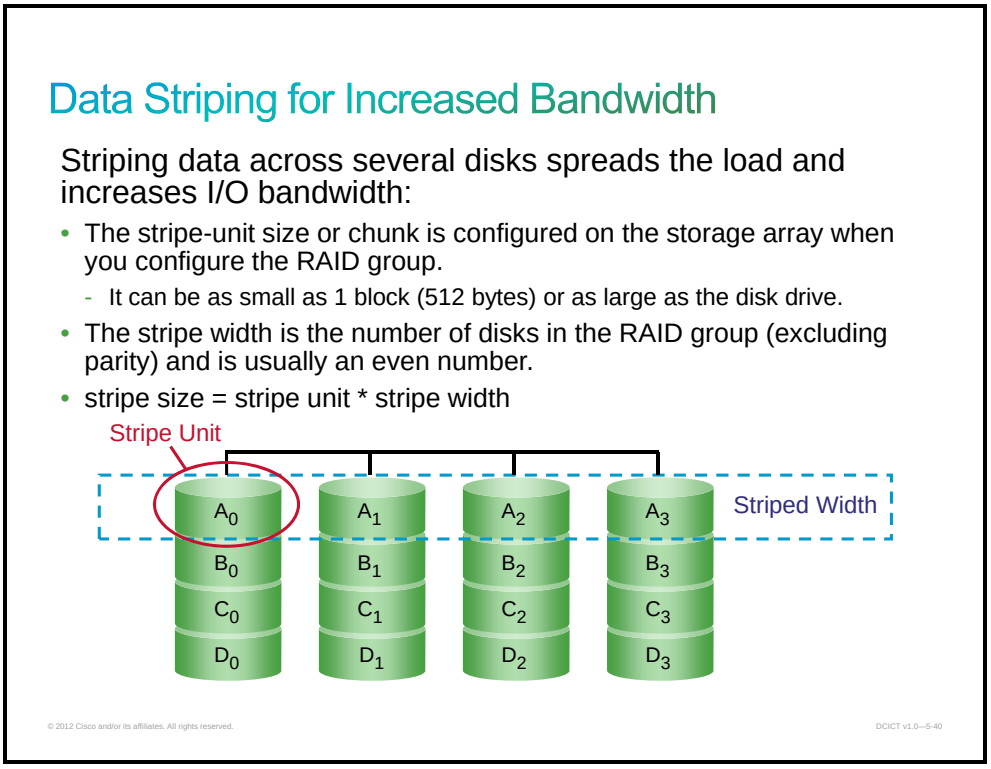
- Extremely high fault tolerance—enough to survive failure of multiple disks
- High I/O rates

RAID 1+0 includes these disadvantages:

- Great expense
- High overhead
- Limited scalability

https://t.me/learningnets

RAID 1+0 can benefit applications that require both high performance and reliability. The types of applications that can benefit from RAID 1+0 include database servers requiring high performance and fault tolerance.



Data Striping for Increased Bandwidth

Striping data across several disks spreads the load and increases the I/O bandwidth.

The stripe-unit size or chunk is configured on the storage array when you configure the RAID group. This chunk can be as small as one block (512 bytes) or as large as the disk drive.

The stripe width is the number of disks in the RAID group (excluding parity) and is usually an even number:

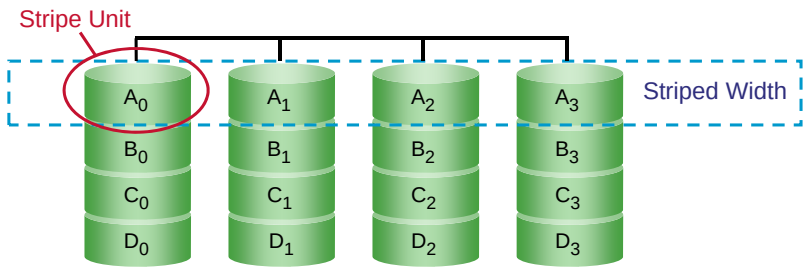
- stripe size = stripe unit * stripe width

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Data Striping and Performance

- Misalignment between application blocks and disk stripe boundaries will slow down your system.
- Best practice is to align blocks with hardware stripe boundaries:
  - The large cylinder represents four striped disks named A0, B0, C0, and D0.
  - A0 represents the configured stripe unit of the array.
  - Create logical disks by using a number of blocks that are evenly divisible by the specified stripe size.



© 2012 Cisco and/or its affiliates. All rights reserved.

When you configure your disks, you can easily overlook a configuration detail that can cause your physical I/O subsystem to execute more single-block I/O calls than your database requests.

In the figure, the large cylinders represent four striped disks. Each disk contains *n* chunks, each of which is the size that is defined by the stripe unit of the array.

The data is striped across the cylinders in stripe units.

The best practice is to align blocks with hardware stripe boundaries.

https://t.me/learningnets

Cisco UCS C-Series Model-Specific RAID Options

- Only the Cisco UCS C200 and C210 models, include built-in RAID controllers.
- Refer to the Cisco UCS C-Series server Installation and Service Guides for specific use cases.

Model	RAID Adapter Options
C200 M2	Intel ICH10R , LSI 1064E, LSI 6G MegaRAID 9260-4i
C210 M2	Intel ICH10R, LSI 1064E, LSI MegaRAID 9280-4i4e, LSI MegaRAID 9261-8i
C250 M2	LSI 3081, LSI SAS3081E-R, LSI MegaRAID 9261-8i
C260 M2	LSI MegaRAID SAS 9261-8i
C460 M2	LSI MegaRAID SAS 9240-8i or LSI MegaRAID SAS 9260-8i

© 2012 Cisco and/or its affiliates. All rights reserved.

Cisco UCS C-Series Rack Servers enjoy a wide range of RAID adapters to suit virtually any local RAID solution. Refer to the previously mentioned data sheets to access more detailed product specifications for supported RAID adapters.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Memory Population Guidelines for Cisco UCS C-Series Rack Servers

This topic describes the best practices for populating DRAM in the Cisco UCS C-Series Rack Servers.

General Memory Population Guidelines

- Take appropriate ESD precautions.
- Fill all slots in a memory channel.
- Do not mix DIMMs of various speeds (lowest speed is always used).
- Refer to the model-specific Cisco UCS C-Series server Installation and Service Guides for detailed specifications.



© 2012 Cisco and/or its affiliates. All rights reserved.

All memory components in Cisco UCS C-Series servers are subject to damage or destruction by ESD. Take precautions against ESD by installing memory at a grounded workstation. Always wear an ESD wrist strap to dissipate any charge on your body that can be transmitted to the components inside the server.

The Nehalem memory architecture comprises three channels that form a logical DIMM. Memory interleaving takes place most efficiently when all three channels in a logical DIMM are filled.

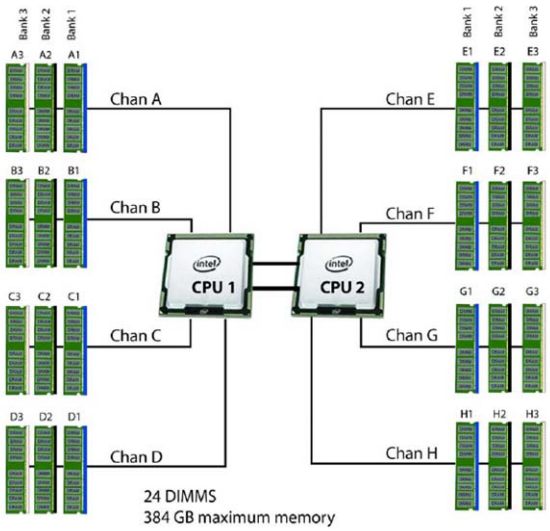
If you mix DIMMs of different speeds in a Cisco UCS C-Series server, then the memory bus operates at the lowest common denominator. If there is a mix of 1066-MHz and 1333-MHz modules in a server, then the memory bus operates at 1066 MHz.

https://t.me/learningnets

Cisco UCS C220 M3 Rack Server Memory Population

Cisco UCS C220 M3 Memory Architecture

- 24 DIMM slots
- Up to 3 DIMMs per channel (DPC)
- Minimum 1 DIMM per CPU (2 DIMMs per server)
- 1600-MHz speed supported up to 2 DPC
- Above 2 DPC (more than 16 DIMMs) speed drops to 1066-MHz



© 2012 Cisco and/or its affiliates. All rights reserved.

The new Intel E5-2600 chipset allows four memory channels per socket with 3 DIMMs per channel. Like in any CPU architecture, optimal configurations maximize memory density and speed. At the writing of this course, the optimal configuration for memory density and speed maximization is 192-GB RAM at 1600 MHz.

Refer to the *Cisco UCS C220 Server Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/C220.html for detailed specifications, compatibility, and performance data.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS C-Series Memory Population

- Memory population guidelines vary greatly from model to model.
- Memory performance depends on several factors:
  - CPU selection determines maximum memory speed.
  - Mixing different-speed DIMMs in the same channel results in lowest common denominator memory bus speed.
- Refer to the model-specific Installation and Service Guides for detailed instructions.

© 2012 Cisco and/or its affiliates. All rights reserved.

Cisco UCS C-Series rack server memory-population guidelines vary greatly from one model to another. Failure to populate memory properly can significantly reduce performance or cause a server to fail power-on self-test (POST).

Refer to these resources for memory-population guidelines for detailed specifications, compatibility, and performance data:

- *Cisco UCS C200 Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C200M1/install/c200M1.html
- *Cisco UCS C210 Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C210M1/install/C210M1.html
- *Cisco UCS C220 Server Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/C220.html
- *Cisco UCS C240 Server Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C240/install/C240.html
- *Cisco UCS C250 Server Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C250M1/install/c250M1.html
- *Cisco UCS C460 Server Installation and Service Guide* at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C460/install/C460.html

You will also benefit from this technical document:

- UCS C-Series Servers Troubleshooting Tips at http://www.cisco.com/en/US/products/ps10493/products_configuration_example09186a0080b10d5e.shtml

https://t.me/learningnets

Summary

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

Summary

- The Cisco UCS C-Series product family is the entry point into Cisco UCS and unified fabric in the data center.
- The Cisco UCS C200 M2 High-Density Rack Server is a dual-socket Intel Xeon 5600 server with up to 192-GB RAM and up to 8-TB local storage.
- The Cisco UCS C210 M2 General-Purpose Rack Server is a dual-socket Intel Xeon 5600 server with up to 192-GB RAM and up to 16-TB local storage.
- The Cisco UCS C220 M3 Rack Server is a dual-socket Intel Xeon 5600 server with up to 192-GB RAM and up to 16-TB local storage.
- The Cisco UCS C240 M3 Rack Server is a dual-socket Intel Xeon 5600 server with up to 192-GB RAM and up to 16-TB local storage.
- The Cisco UCS C250 M2 Extended Memory Rack Server is a dual-socket Intel Xeon 5600 server with up to 384-GB RAM and up to 8-TB local storage.
- The Cisco UCS C260 M2 Rack Server is a dual-socket Intel Xeon E7-2800 server with up to 1-TB RAM and up to 16-TB local storage.

© 2012 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Describe the Cisco UCS C440 M2 High-Performance Rack Server is a quad-socket Intel Xeon E7-4800 server with up to 1-TB RAM and up to 12-TB local storage.
- There are a wide variety of Cisco UCS C-Series network adapters to suit virtually any LAN and SAN use case.
- RAID adapter options for Cisco UCS C-Series Rack Servers provide enterprise configuration options and are model-specific.
- The memory performance of Cisco UCS C-Series Rack Servers depends on following best practices for DIMM population.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 3

Connecting Cisco UCS B-Series Blade Servers

Overview

The purpose of this lesson is to describe how Cisco Unified Computing System (UCS) B-Series Blade Servers connect to the chassis and achieve external LAN and SAN connectivity.

Objectives

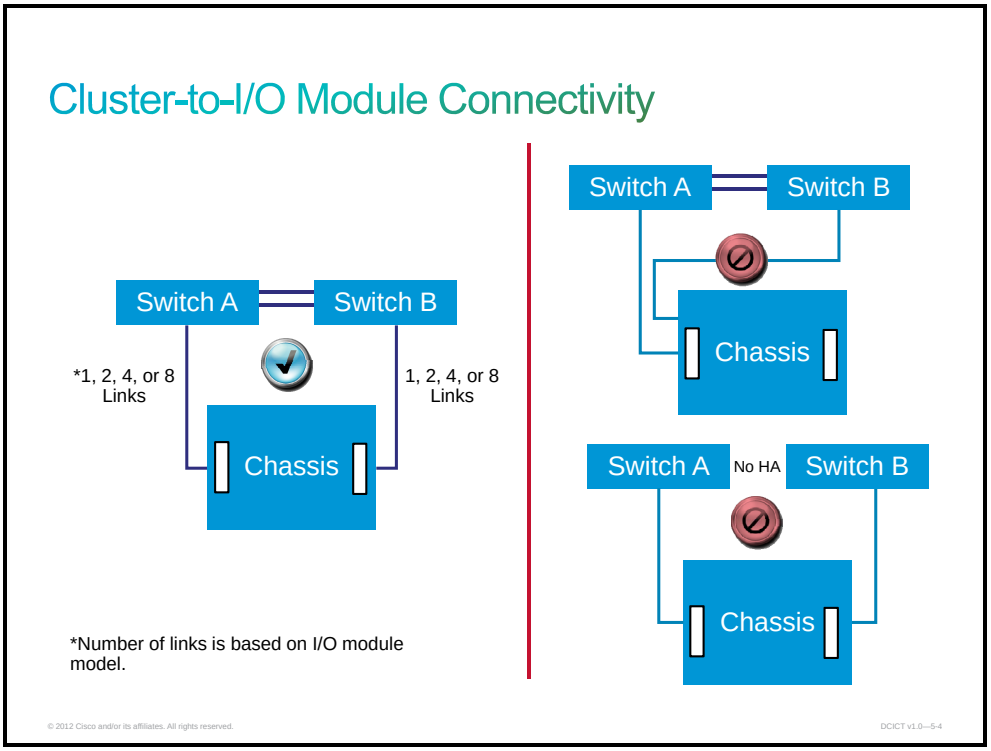
Upon completing this lesson, you will be able to describe how Cisco UCS B-Series Blade Servers connect to the chassis and achieve external LAN and SAN connectivity. You will be able to meet these objectives:

- Describe the options for connecting a blade server chassis to fabric interconnects
- Describe how physical connections from the fabric interconnect to the I/O module are routed to various internal components of the blade servers
- Describe the three main features of the Cisco Integrated Management Controller
- Describe the three basic states in which a 10 Gigabit Ethernet interface can be configured on the fabric interconnect
- Describe the discovery process and how to monitor using FSM output

https://t.me/learningnets

Chassis-to-Fabric Interconnect Physical Connectivity

This topic describes the options for connecting a blade server chassis to fabric interconnects.



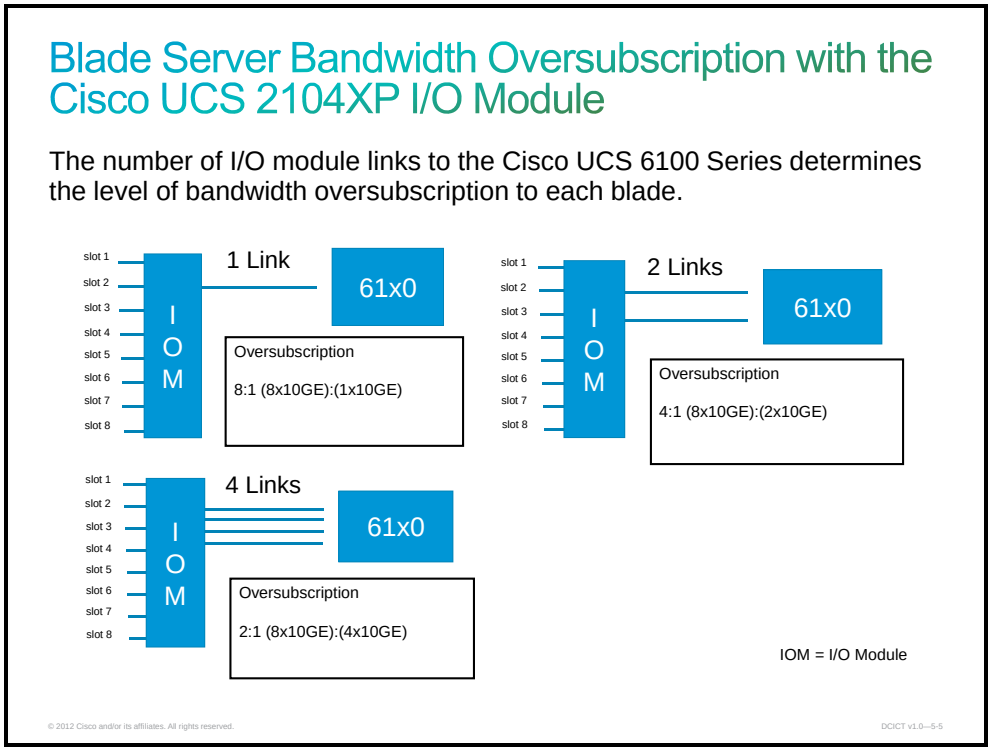
The diagram on the left in the figure illustrates the proper connectivity of a Cisco UCS 5108 Blade Server Chassis, from the I/O module to the fabric interconnects. I/O module A connects to fabric interconnect A with 1, 2, 4, or 8 links. I/O module B connects to fabric interconnect B with 1, 2, 4, or 8 links.

Note	The number of available links depends on which I/O module is deployed. The Cisco UCS 2104XP I/O module supports as many as four links, and the Cisco UCS 2208XP I/O module supports as many as eight links. Regardless of the I/O module, the number of links must be a power of 2.

The top-right diagram in the figure shows two links from I/O module A connecting to fabric interconnect A and two links connecting to fabric interconnect B. When the links attempt to register with fabric B, Cisco UCS Manager rejects the association because that I/O module is already registered to fabric A. The links of any I/O module can be connected to only one fabric interconnect.

The bottom-right diagram in the figure represents a situation in which the I/O modules are connected to the correct fabric interconnects, but the fabric interconnects are configured as standalone (nonclustered) devices with no cluster links. When the I/O module attempts to register with fabric interconnect B, it receives a conflicting cluster ID and will not associate with fabric interconnect B.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

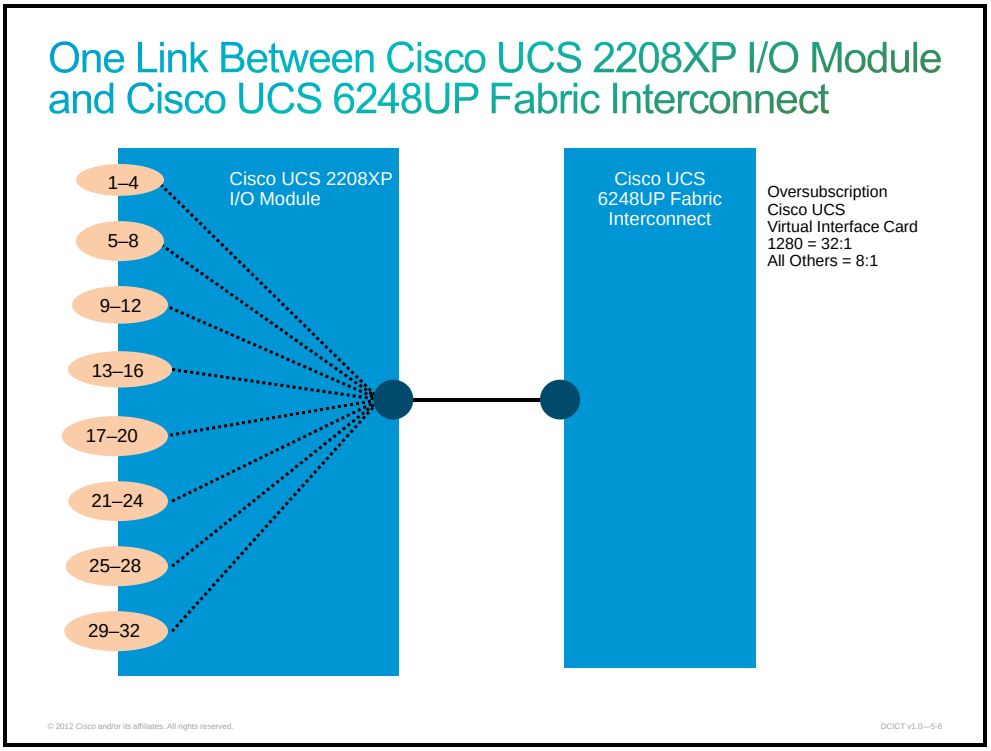


With the Cisco UCS 2104XP I/O module, bandwidth oversubscription is determined by the number of links from the I/O module to the fabric interconnects. The table lists the supported oversubscription rates:

Cisco UCS 210XP I/O Module Oversubscription Rates

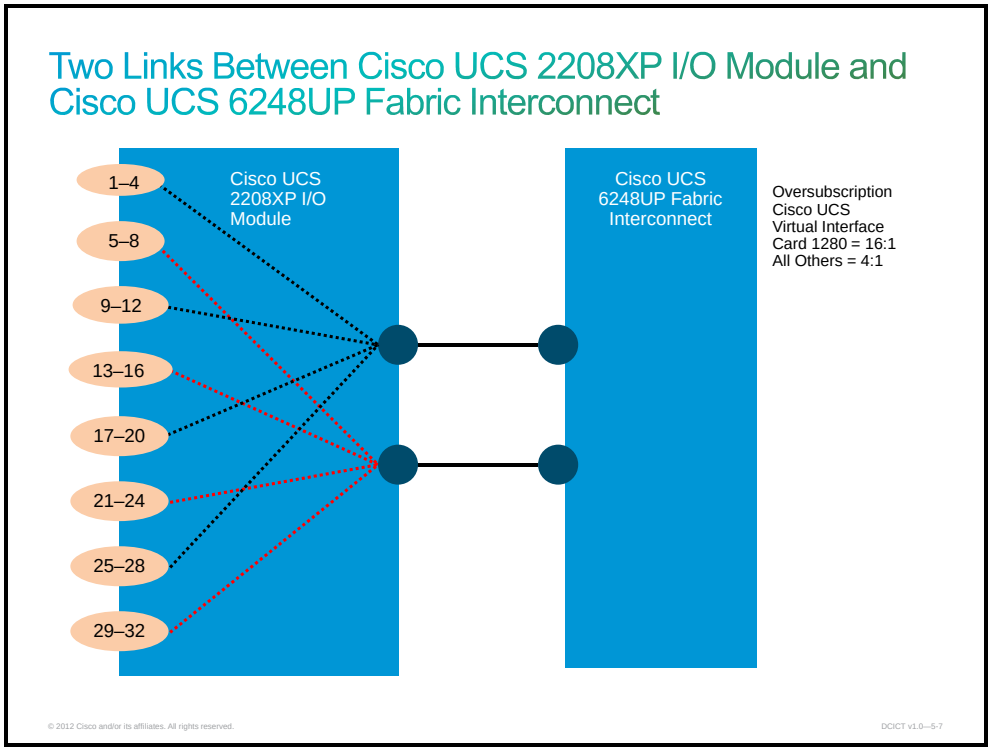
I/O Module Links	Server Oversubscription Rate
1	8:1
2	4:1
4	2:1

https://t.me/learningnets



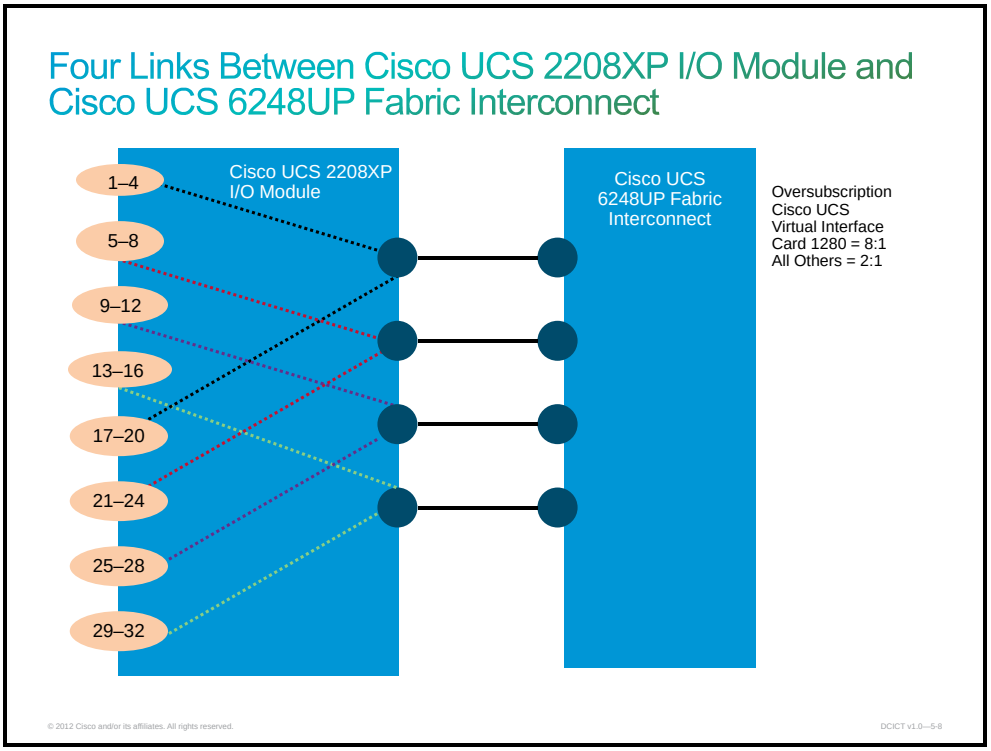
With the release of the second-generation Cisco UCS hardware, new topologies and oversubscription rates are possible. This figure illustrates the oversubscription rates using Cisco UCS 6248UP 48-Port Fabric Interconnects, Cisco UCS 2208XP I/O modules, and Cisco UCS Virtual Interface Card 1280 adapters on each of the eight half-slot Cisco UCS B-Series Blade Servers in the chassis. The Cisco UCS Virtual Interface Card 1280 has four 10 Gigabit Ethernet interfaces on each fabric. The numbers 1 through 4 represent the four 10-GB adapters on server 1. The other numbers represent the four 10 Gigabit Ethernet adapters on the remaining seven blade servers. In this example, all 32 10-GB connections are pinned to the one available uplink. This configuration results in an oversubscription rate of 32:1 when using Cisco UCS Virtual Interface Card 1280 adapters. If any other mezzanine card is used in the servers, then only one 10 Gigabit Ethernet interface per fabric is available on the adapter. The resulting oversubscription rate would be 8:1.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



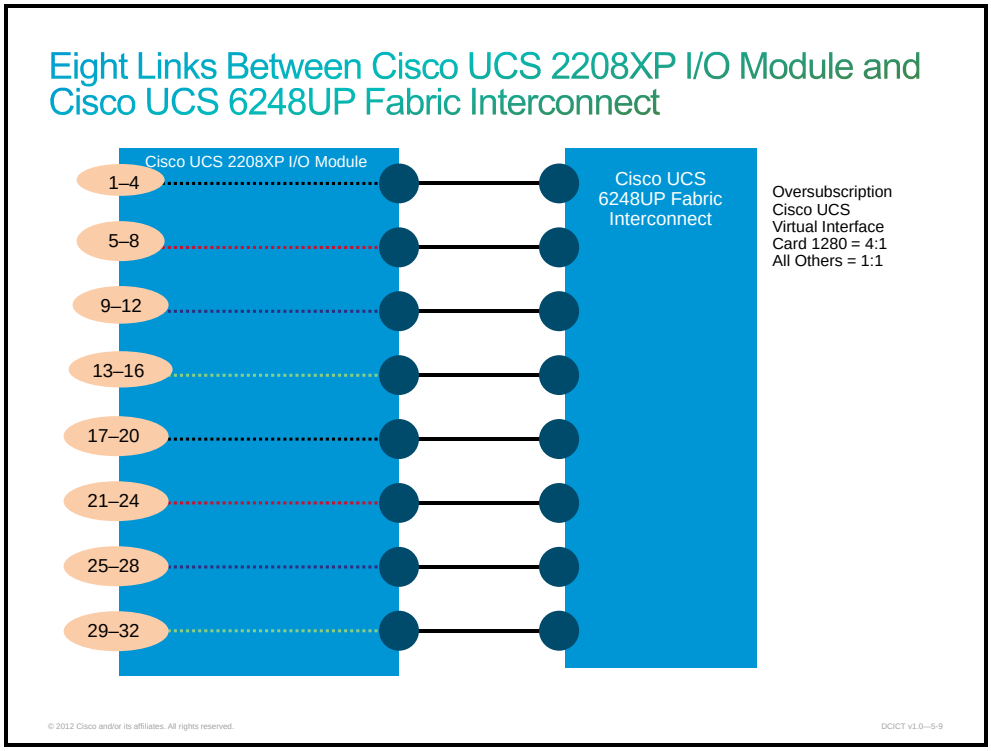
In figure, 16 10-GB connections are pinned to uplink 1, and the other 16 connections are pinned to uplink 2. This configuration results in an oversubscription rate of 16:1 when using Cisco UCS Virtual Interface Card 1280 adapters. If any other mezzanine card is used in the servers, then only one 10 Gigabit Ethernet interface per fabric is available on the adapter. The resulting oversubscription rate would be 4:1.

https://t.me/learningnets



In the figure, eight 10-GB connections are pinned to uplink 1. The next eight 10-GB connections are pinned to uplink 2. The next eight 10-GB connections are pinned to uplink 3, and the remaining eight 10-GB connections are pinned to uplink 4. This configuration results in an oversubscription rate of 8:1 when using Cisco UCS Virtual Interface Card 1280 adapters. If any other mezzanine card is used in the servers, then only one 10 Gigabit Ethernet interface per fabric is available on the adapter. The resulting oversubscription rate would be 2:1.

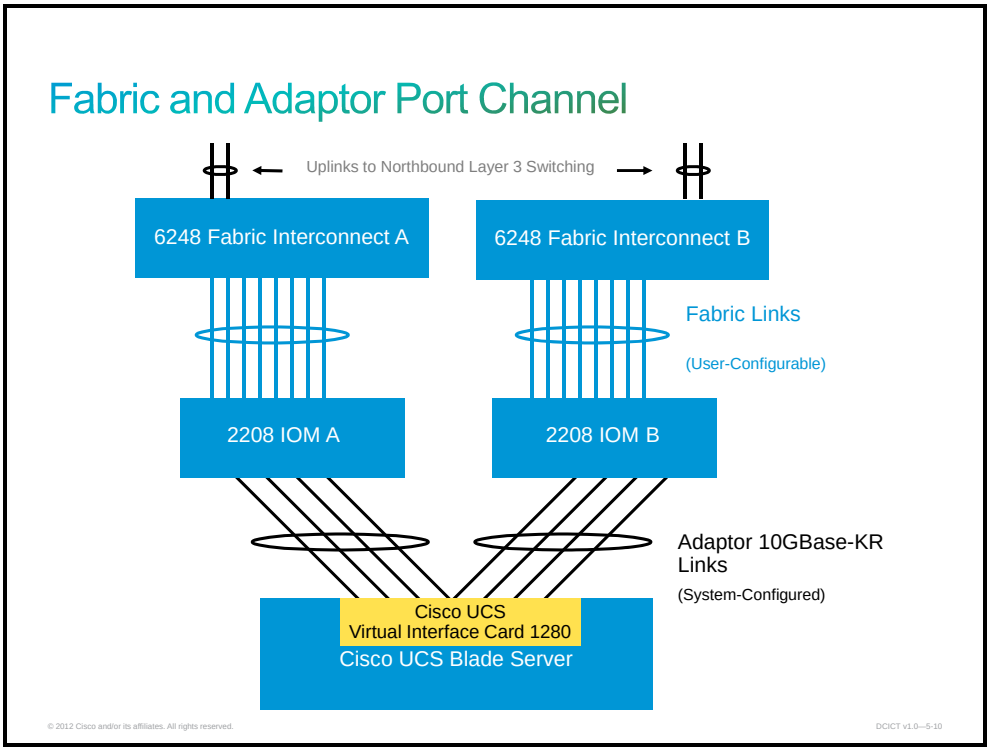
© 2012 Cisco Systems, Inc.
https://t.me/learningnets



In the figure, each Cisco UCS Virtual Interface Card 1280 adapter is assigned to an unshared 10-GB link to the fabric interconnect. This configuration results in an oversubscription rate of 4:1 when using Cisco UCS Virtual Interface Card 1280 adapters. If any other mezzanine card is used in the servers, the only one 10 Gigabit Ethernet interface per fabric is available on the adapter. The resulting oversubscription rate would be 1:1.

Note The number of links must be a power of 2: 1, 2, 4, or 8.

https://t.me/learningnets



One important benefit of upgrading to Cisco UCS second-generation hardware is the ability to configure port channels from the Cisco UCS 6248UP fabric interconnects to the Cisco UCS 2208XP I/O modules and to present a port channel interface to the Cisco UCS Virtual Interface Card 1280 adapter of each server. This configuration offers these benefits:

- Higher bandwidth
  - Maximum bandwidth available to a server can be greater than 10 Gb/s
  - Necessary to fully service the 40G Cisco UCS Virtual Interface Card 1280 adapter
- Redundancy (operational as long as one link is active)
- Server adapter that is pinned to port channel (not to a specific I/O module uplink)
  - Removes complexity of blade placement because it does not matter which slot the blade is in
  - No failover when an active fabric link fails
- Load-balancing hash (source-dest-mac)
  - Better traffic distribution among I/O module uplinks and better link utilization
  - Load balancing implicitly managed by Cisco UCS Manager and cannot be modified

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Fabric Port Channel Support Matrix

Fabric port channels require the Cisco UCS 6248UP or 6296UP fabric interconnects, Cisco UCS 2204XP or 2208XP I/O module, and Cisco UCS Virtual Interface Card 1280 or 1240 on each server in the chassis.

Fabric Interconnect	I/O Module	Fabric Port Channel Supported
Cisco UCS 6120XP and 6140XP	Cisco UCS 2104XP	No
	Cisco UCS 2204XP	No
	Cisco UCS 2208XP	No
6248UP and 6296UP	Cisco UCS 2104XP	No
	Cisco UCS 2204XP	Yes
	Cisco UCS 2208XP	Yes

© 2012 Cisco and/or its affiliates. All rights reserved.

The fabric port channel matrix illustrates that the Cisco UCS implementation must include only second-generation fabric interconnects, I/O modules, and Cisco UCS Virtual Interface Card 1280 or 1240 mezzanine adapters.

© 2012 Cisco Systems, Inc.

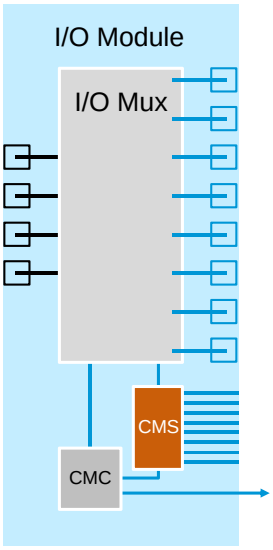
https://t.me/learningnets

I/O Module Architectures

This topic describes how physical connections from the fabric interconnect to the I/O module are routed to the various internal components of the blade servers.

I/O Module Components

- I/O multiplexer multiplexes data from the fabric interconnect to host ports.
- Chassis management switch carries management traffic to the Cisco Integrated Management Controller.
- Chassis management controller is used in chassis and component discovery.
- Chassis management controller also monitors chassis sensors and controls fan speed.



Mux = Multiplexer
CMS = Chassis Management Switch
CMC = Chassis Management Controller

© 2012 Cisco and/or its affiliates. All rights reserved.

Although the I/O module is frequently referred to as a fabric extender (FEX), the I/O module includes additional functionality that is not found in Cisco Nexus 2000 Series Fabric Extenders. These are the three main component blocks of an I/O module:

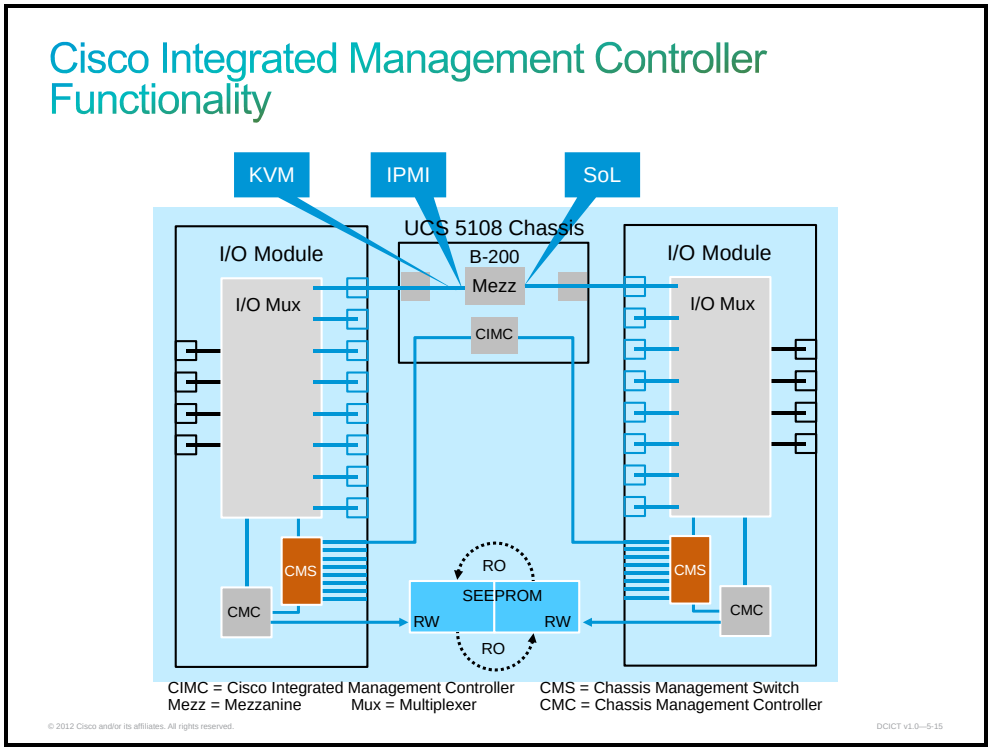
- Chassis management controller
- Chassis management switch
- I/O multiplexer

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Integrated Management Controller Chip on Cisco UCS B-Series Blade Servers

This topic describes the three main features of the Cisco Integrated Management Controller.



This figure illustrates the Cisco Integrated Management Controller chip that is included on the motherboard of every Cisco UCS B-Series and Cisco UCS C-Series server. The Cisco Integrated Management Controller provides these three important functions:

- Keyboard, video, mouse (KVM) over IP
- Serial over LAN (SoL)
- Intelligent Platform Management Interface (IPMI)

Note	The Cisco Integrated Management Controller was formerly known as the Baseboard Management Controller. You might encounter this terminology in older Cisco UCS documentation.

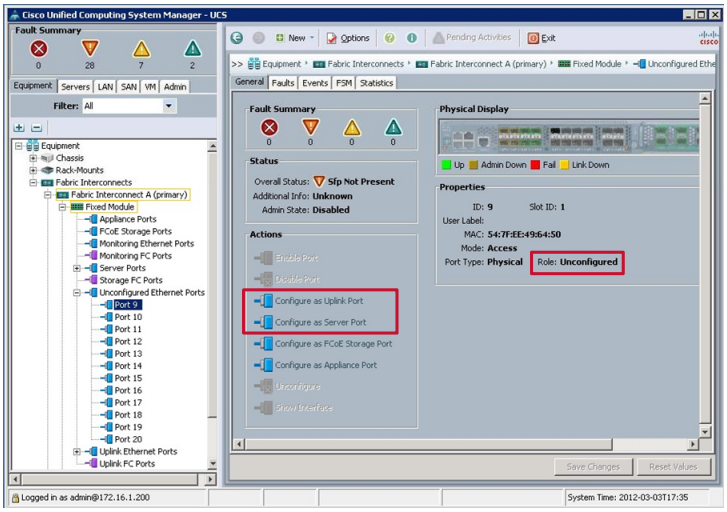
https://t.me/learningnets

Three Basic Port Personalities in the Fabric Interconnect

This topic describes the three basic states in which a 10-GB interfaces can be configured on the fabric interconnect.

Unconfigured Port Personalities

The default state of every Ethernet port is unconfigured.



© 2012 Cisco and/or its affiliates. All rights reserved.

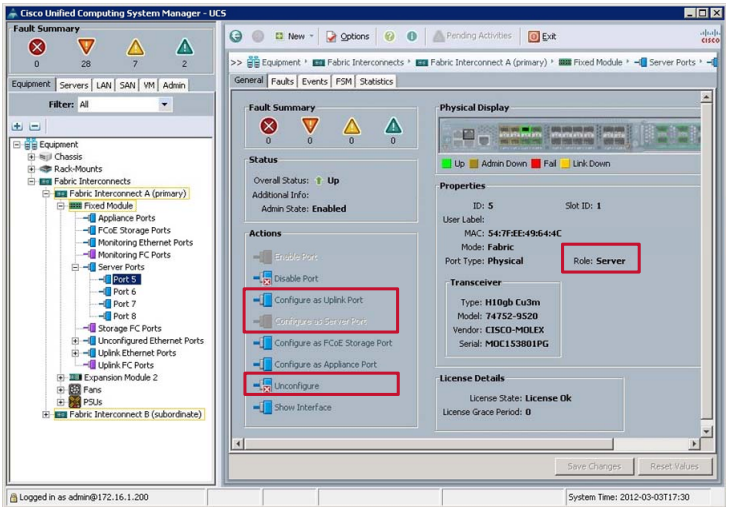
All Ethernet ports that are in the fabric interconnect will default to the unconfigured state. This default state is the same for both the fixed ports and Ethernet ports in expansion modules. This default port personality is consistent with both first- and second-generation fabric interconnects.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Basic Ethernet Port Personalities

Ports can also be configured as server or uplink.



© 2012 Cisco and/or its affiliates. All rights reserved.

DCICT v1.0—5-18

The figure illustrates an Ethernet port that is assigned the Server role. This role indicates that this port connects to an I/O module port on a Cisco UCS 5108 Blade Server Chassis. The personality of the port can be changed dynamically.

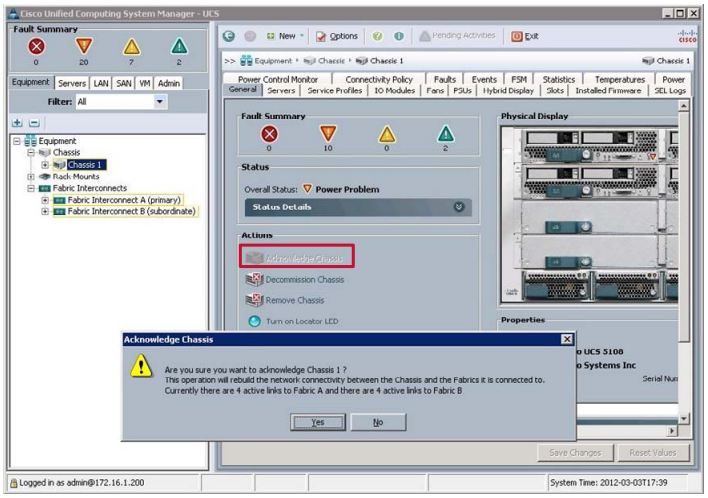
https://t.me/learningnets

Discovery Process

This topic describes the discovery process and how use finite state machine (FSM) output for monitoring.

Acknowledge Chassis

Equipment > Chassis 1 > Actions > Acknowledge Chassis



© 2012 Cisco and/or its affiliates. All rights reserved.

Chassis discovery occurs when a chassis is initially connected to the fabric interconnect and Cisco UCS Manager recognizes that a new server link is active. A connection is made to the chassis management controller, which sends information about the chassis inventory (I/O module, fans, power supplies, serial numbers, part IDs, and so on) to Cisco UCS Manager.

If server presence is detected, then the inventory of the server (or servers) is sent to Cisco UCS Manager. The details of the inventory include serial number, vendors, CPU type, installed DIMMs, adapter cards, hard drives, BIOS, and Cisco Integrated Management Controller.

The Data Management Engine (DME) stores all the inventory details as managed objects in the Cisco UCS Manager database.

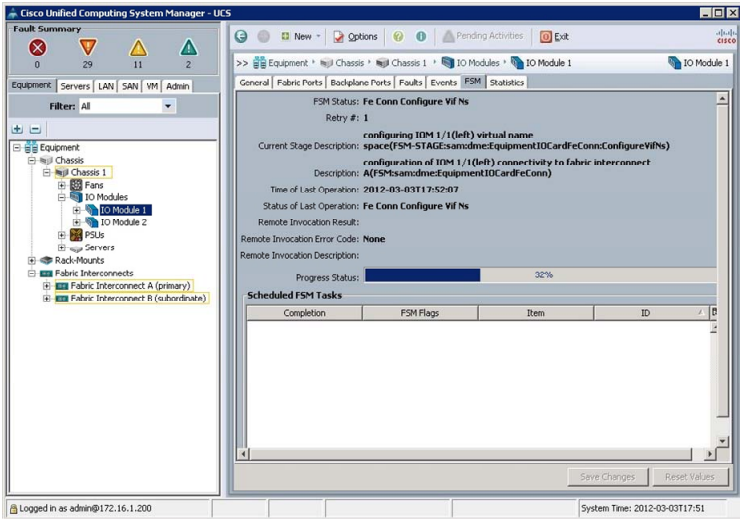
To observe chassis discovery, you can re-acknowledge the chassis and monitor the process.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Monitor I/O Module Discovery in FSM

Equipment > Chassis 1 > IO Module 1 > FSM



© 2012 Cisco and/or its affiliates. All rights reserved.

After the chassis has been re-acknowledged, you can watch the process in the FSM tab of either I/O module.

An FSM is a series of logical transitions that occur in a specific order. The chassis and server discovery processes are examples of the use of an FSM to validate a complex series of events.

https://t.me/learningnets

Summary

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

Summary

- Depending on the deployed Cisco UCS hardware, blade servers are pinned to either an I/O module uplink or an I/O module port channel.
- The I/O module includes the chassis management controller, chassis management switch, and an I/O multiplexer.
- The Cisco Integrated Management Controller provides KVM over IP, SoL, and IPMI.
- A 10 Gigabit Ethernet interface can be in three basic states:
  - Unconfigured
  - Server
  - Uplink
- All devices connected to the fabric interconnects are discovered at the time of connection and are FSM output.

© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

Setting up an Initial Cisco UCS B-Series Cluster

Overview

The purpose of this lesson is to describe the required information to perform initial Cisco Unified Computing System (UCS) cluster setup for Cisco UCS B-Series Blade Servers and to validate cluster operation.

Objectives

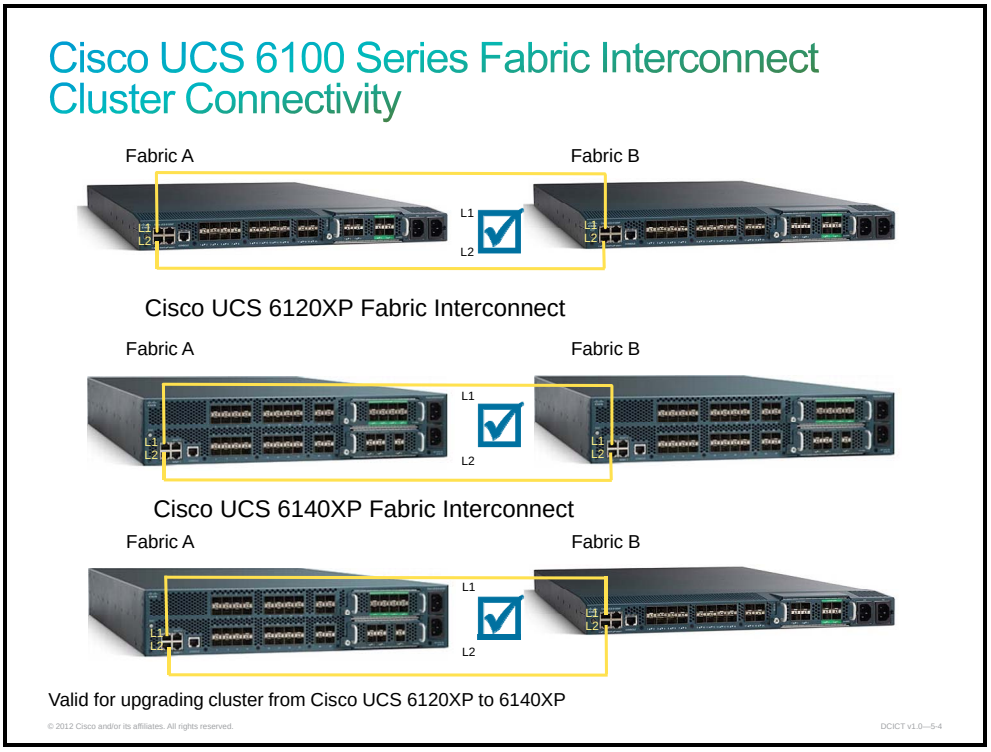
Upon completing this lesson, you will be able to describe the required information to perform initial Cisco UCS B-Series cluster setup and to validate cluster operation. You will be able to meet these objectives:

- Describe the required physical connectivity to create a fabric interconnect cluster
- Describe the required information to configure the primary cluster peer
- Describe the required information to configure the secondary cluster peer
- Verify that normal cluster communications and roles have been established

https://t.me/learningnets

Cabling a Cisco UCS Fabric Interconnect Cluster

This topic describes the required physical connectivity to create a fabric interconnect cluster.



Cisco UCS 6100 Series and Cisco UCS 6200 Series Fabric Interconnects offer the ability to configure a high-availability cluster that can eliminate a single point of failure for the management, control, and data planes.

On a Cisco UCS 6120XP 20-Port Fabric Interconnect or Cisco UCS 6140XP 40-Port Fabric Interconnect, the L1 and L2 links must be cross-connected from the primary to subordinate peers. This cross-connection requires a straight-through Category 6 Ethernet cable.

Although fabric interconnects are not required to be deployed in a clustered configuration, it is a best practice to engineer a solution without single points of failure.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco UCS 6248UP Fabric Interconnect Cluster Connectivity

The cluster connections L1 and L2 for the Cisco UCS 6248UP fabric interconnects are located at the rear of the chassis.



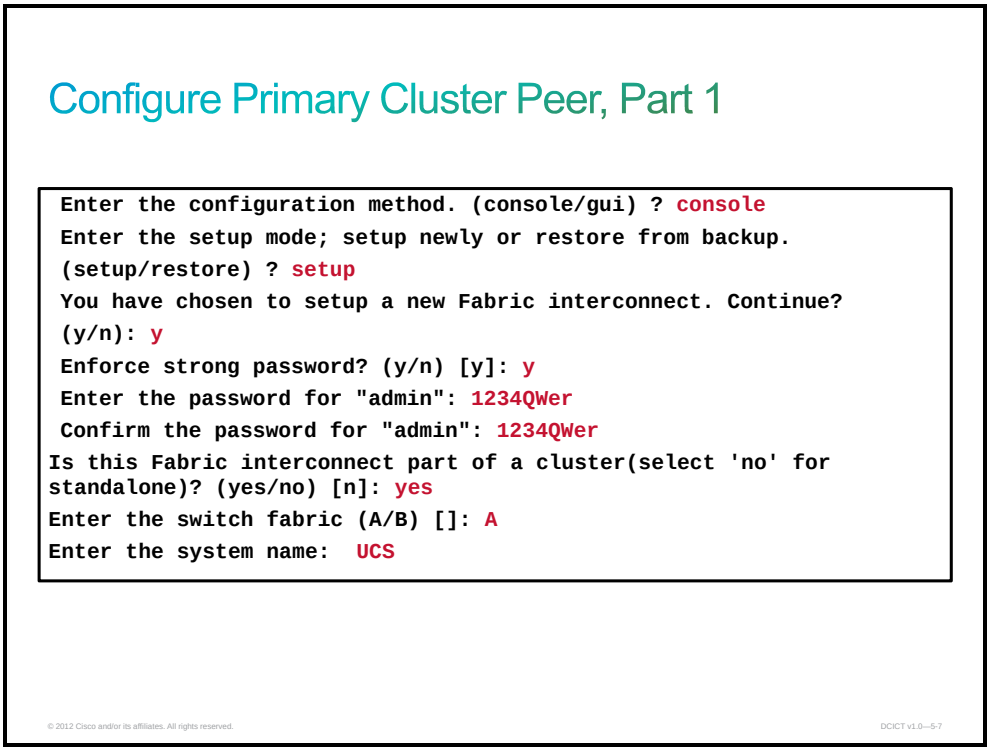
© 2012 Cisco and/or its affiliates. All rights reserved.

With Cisco UCS 6248UP 48-Port fabric interconnects, the L1 and L2 links are at the rear of the chassis. These links must be cross-connected from the primary to subordinate peers. This cross-connection requires a straight-through Category 6 Ethernet cable.

https://t.me/learningnets

Initial Setup Script for the Primary Peer

This topic describes the information that is needed to configure the primary cluster peer.



To begin the cluster configuration on a new fabric interconnect, you must connect to the serial console of the first member of the cluster (fabric A). Although a GUI Express Start option is available after fabric interconnect A is configured with an IP address, only the CLI method is covered in this lesson.

As on most Cisco devices with an RJ-45 serial console port, you can use a standard Cisco DB-9 to RJ-45 console cable. Configure your favorite terminal emulator for 9600 b/s, 8 data bits, 1 stop bit, and no parity (9600-8-1-none).

You will be asked whether you want to enter setup. Enter **yes**.

All cluster configurations are completed from the console, so choose **console**. Enter **yes** when asked whether you want to create a cluster on the switch. If you decide to initially configure the fabric interconnect as standalone, you can convert it to cluster mode later.

The fabric selection for the first fabric interconnect is usually A. The role of fabric A or fabric B is established and permanent at the time of the installation. However, the role of primary or subordinate is assigned by an election process. Fabric B might be elected as the primary peer.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Configure Primary Cluster Peer, Part 2

```
Physical Switch Mgmt0 IPv4 address : 172.16.1.101
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 172.16.1.254
Cluster IPv4 address : 172.16.1.200
Configure the DNS Server IPv4 address? (yes/no) [n]: n
Configure the default domain name? (yes/no) [n]: n
```

© 2012 Cisco and/or its affiliates. All rights reserved.

After entering the hostname and IP information, you must choose a virtual IP address. You will use this virtual IP address to connect to the active management node. Although both nodes can be accessed from the CLI, only the active management node permits a Cisco UCS Manager GUI connection.

https://t.me/learningnets

Configure Primary Cluster Peer, Part 3

```
Following configurations will be applied:
    Switch Fabric=A
    System Name=UCS
    Enforced Strong Password=yes
    Physical Switch Mgmt0 IP Address=172.16.1.101
    Physical Switch Mgmt0 IP Netmask=255.255.255.0
    Default Gateway=172.16.1.254
    Cluster Enabled=yes
    Cluster IP Address=172.16.1.200
    NOTE: Cluster IP will be configured only after both Fabric
Interconnects are initialized
  Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no): yes
  Type 'reboot' to abort configuration and reboot system
  or hit enter to continue. (reboot/<CR>)? <CR>
Applying configuration. Please wait.
   Configuration file - Ok
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The remainder of the dialog prompts for basic network-connectivity options, including the Domain Name System (DNS) server and domain. The setup wizard then allows you to review your configuration and make corrections if necessary. After review, enter **yes** to apply and save the configuration. The cluster services are configured and enabled on fabric interconnect A.

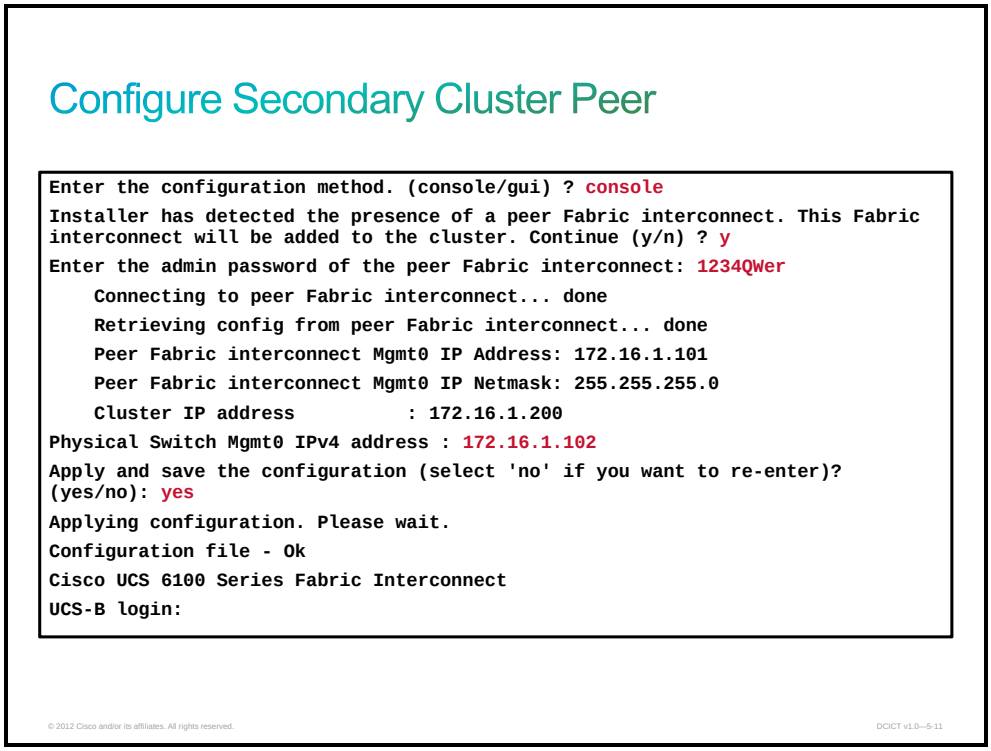
Note	Hostname and domain name need to be correctly configured to enroll with a certificate authority (CA). The fully qualified domain name (FQDN) is part of the certificate signing request (CSR). Be certain that the FQDN is unique within your domain.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Initial Setup Script for the Secondary Peer

This topic describes the information that is needed to configure the secondary cluster peer.

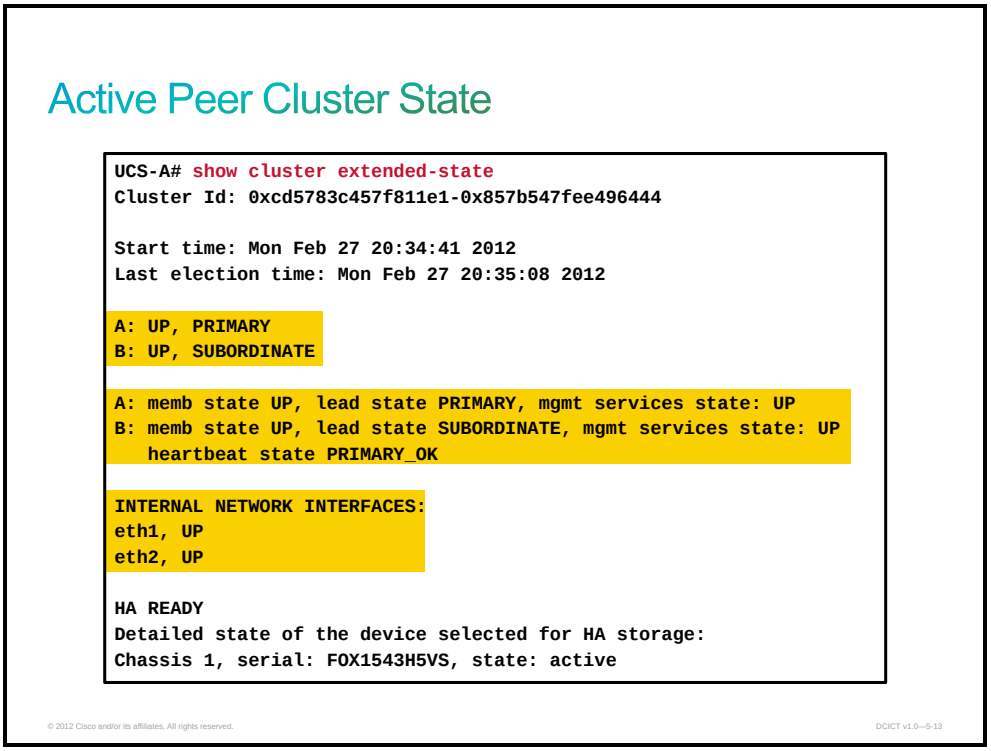


Assuming that the L1 and L2 links are correctly connected, the second peer detects the primary peer and asks to join the cluster. Enter the admin password and choose **yes** after reviewing the configuration. The primary peer performs an initial synchronization from active to subordinate node. This synchronization includes all elements of the Cisco UCS Manager Configuration database and images in the firmware store.

https://t.me/learningnets

Verifying a Fabric Interconnect Cluster

This topic describes how Cisco UCS Manager discovers Cisco UCS B-Series Blade Servers.

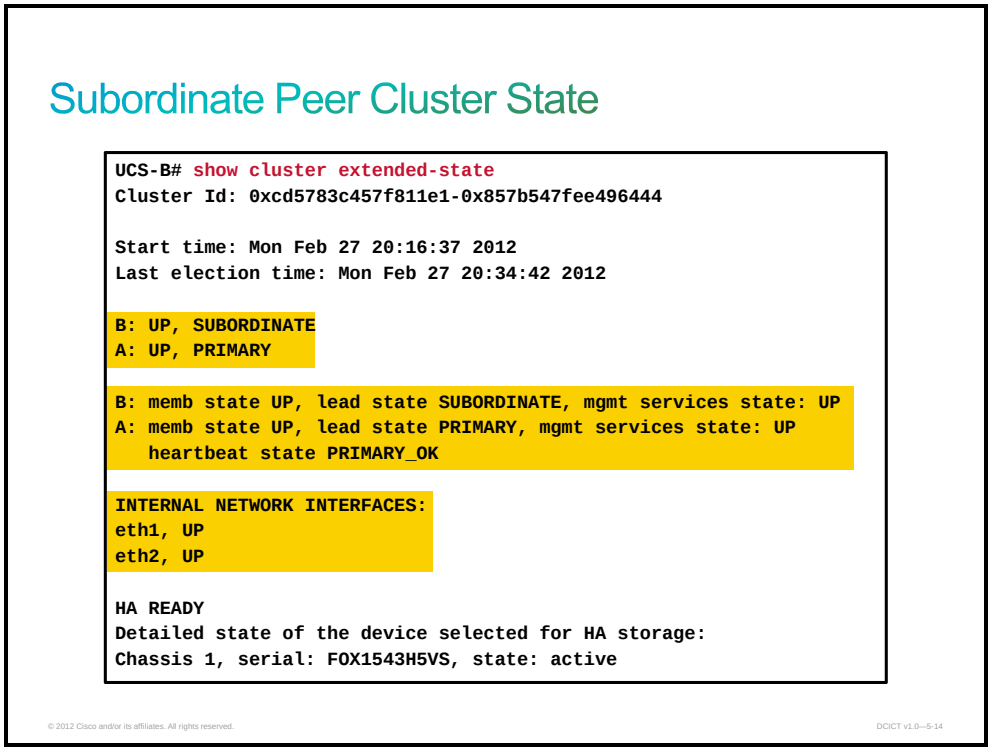


The **show cluster extended-state** command includes information that administrators need to validate cluster configuration and troubleshoot cluster operation. The output in the figure indicates that this fabric interconnect is operating fabric A and is the primary management node. All configuration is performed on the active management node. Before the active node commits a configuration transaction to the data management engine (DME), the node first replicates the transaction to the subordinate node. After the subordinate acknowledges that its DME has committed the transaction, the primary node commits the transaction.

On the active management node, the management plane and data planes are both active.

© 2012 Cisco Systems, Inc.

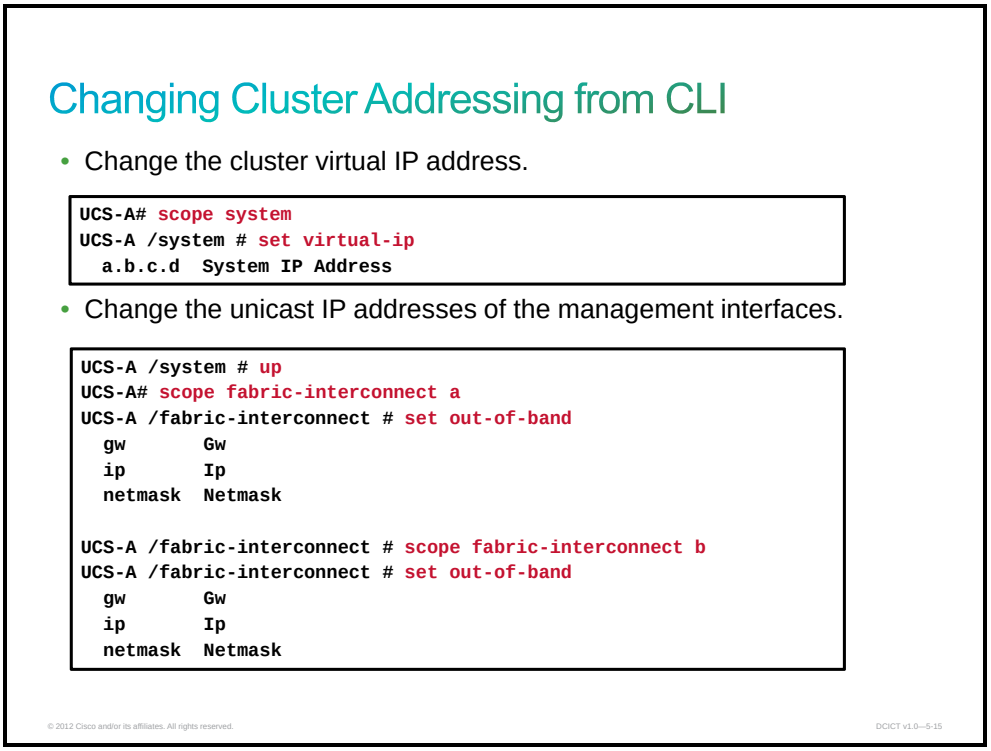
https://t.me/learningnets



The subordinate node is not in a passive standby mode. The DME is active and accepts data that is synchronized from the active management node. Therefore, a distinction is made that this node is subordinate active. If the active management node fails, the subordinate will wait a predetermined time and then declare the peer dead. At that time, the subordinate becomes the active management node.

Even when in subordinate mode for the management plane, the data plane is active and forwarding.

https://t.me/learningnets



The ability to modify the IP settings of the cluster from the serial console is useful. For example, if the wrong subnet mask is set on a cluster peer during setup, the cluster will not synchronize. This lack of synchronization can cause an inability to access the virtual IP address for management.

The commands in the figure, shown here with sample IP addresses, are used to modify the IP settings:

```
s6100-A # scope fabric-interconnect a
s6100-A /fabric-interconnect # set out-of-band ip 172.16.1.101
s6100-A /fabric-interconnect # set out-of-band netmask 255.255.255.0
s6100-A /fabric-interconnect # set out-of-band gw 172.16.1.254
```

The address of the cluster virtual-IP can be changed only in the Cisco UCS CLI, by using the commands in this way:

```
s6100-A# scope system
s6100-A /system # set virtual-ip 172.16.1.200
```


© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Summary

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

Summary

- Cluster peers require 1 Gigabit Ethernet connections from L1 to L1 and L2 to L2.
- The primary cluster peer is configured with a hostname, admin password, fabric affiliation, and IP addressing.
- The secondary cluster peer requires only IP addressing information. The rest of its configuration is synchronized by the primary cluster peer.
- Verify that normal cluster communications and roles have been established.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 5

Describing Cisco UCS Manager Operations

Overview

The purpose of this lesson is to describe how Cisco Unified Computing System (UCS) Manager discovers new hardware and manages and monitors all Cisco UCS components.

Objectives

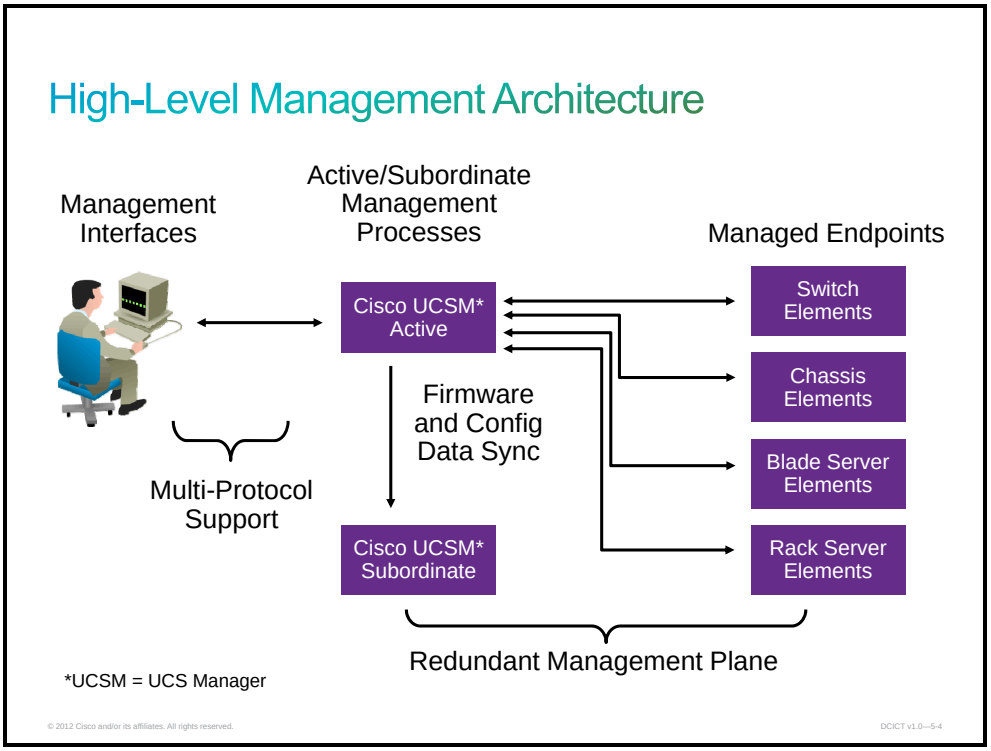
Upon completing this lesson, you will be able to navigate the Cisco UCS Manager GUI and access common administrative functions. You will be able to meet these objectives:

- Describe the functions of Cisco UCS Manager
- Describe the basic layout of Cisco UCS Manager
- Describe which configuration elements are configured under the six tabs of the navigation pane
- Describe how Cisco UCS Manager discovers components
- Describe how to use FSMs to monitor device discovery

https://t.me/learningnets

Cisco UCS Manager

This topic describes the functions of Cisco UCS Manager.

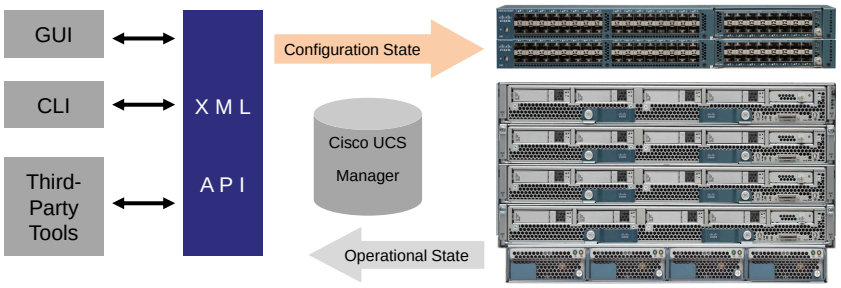


Cisco UCS Manager is the management service for all the Cisco UCS components. As such, users can manage the entire Cisco UCS system through one GUI window. This GUI interface reduces management complexity. Cisco UCS Manager runs on the management nodes (Cisco UCS fabric interconnects). In a high-availability configuration, a redundant instance of Cisco UCS Manager runs on the second fabric interconnect. Managed switch elements include Generic Expansion Modules (GEMs) and ports. Managed chassis elements include the chassis management controller, fan modules, and power supplies. Managed server elements include disks, mezzanine cards, Cisco Integrated Management Controller, BIOS, and so on.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Cisco UCS Manager and XML

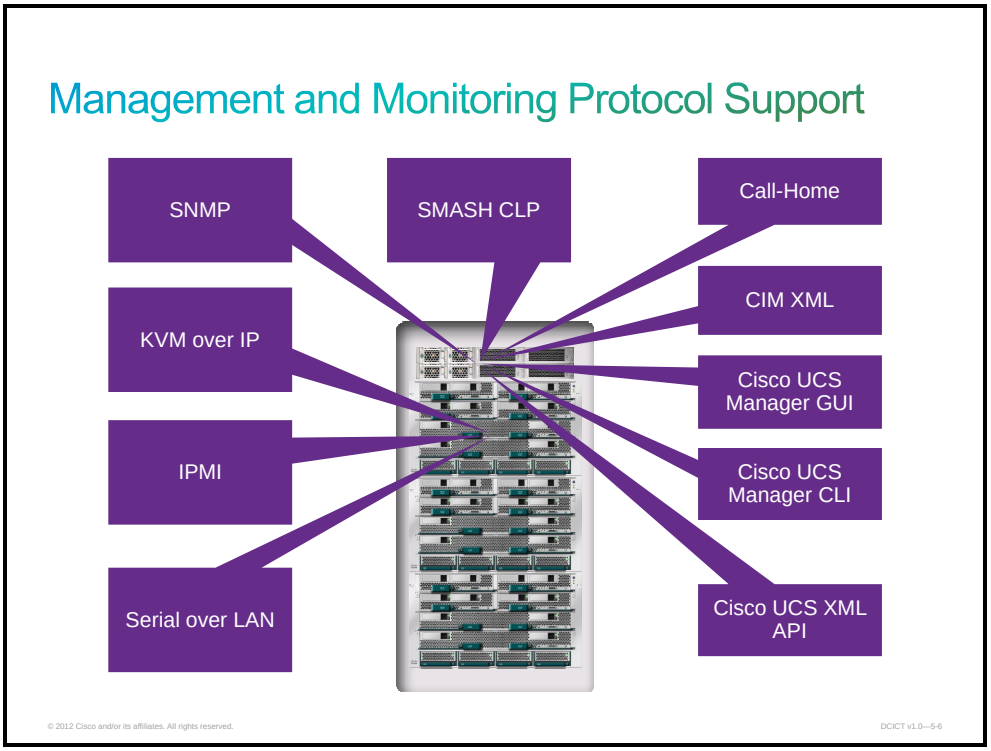
- The Cisco UCS Manager database stores all configuration data in XML format.
- Whether configuration or monitoring is performed in the Cisco UCS Manager GUI, CLI, or a third-party application, all communication is in XML format.



© 2012 Cisco and/or its affiliates. All rights reserved.

The native language of Cisco UCS Manager is XML. Both the GUI and CLI are written by using the Cisco UCS Manager XML application programming interface (API). Third-party tools can also leverage the Cisco UCS Manager XML API to integrate with Cisco UCS Manager. When administrators submit configuration requests for a managed device to Cisco UCS Manager, Cisco UCS Manager will call the appropriate device manager to deploy the management request. The operational state of a device is communicated back to Cisco UCS Manager and reflected in the user interfaces (GUI and CLI). The Cisco UCS Manager database also stores the running state of components. As an example, if the fabric interconnects are rebooted, Cisco UCS Manager knows which service profile is associated with a given server.

https://t.me/learningnets



Cisco UCS Manager supports a variety of industry-standard protocols for management, monitoring, and logging. This includes but is not limited to the following protocols:

- Simple Network Management Protocol (SNMP)
- Keyboard, video, mouse (KVM) over IP
- Intelligent Platform Management Interface Specification (IPMI)
- Server Hardware Command Line Protocol (SMASH CLP)
- Common Information Model (CIM) XML
- Serial over LAN
- Call Home
- Cisco UCS Manager GUI
- Cisco UCS Manager CLI
- Cisco UCS Manager XML API

The Cisco UCS Manager XML API (and corresponding software development kit [SDK]) allows you to perform custom wrappers around Cisco UCS configuration and monitoring.

These are possible uses:

- Creating a multitenant portal with your own presentation and access control interface
- Integrating with customized automation tools
- Scripting and integrating into custom management solutions
- Extending the base functionality that Cisco UCS Manager provides

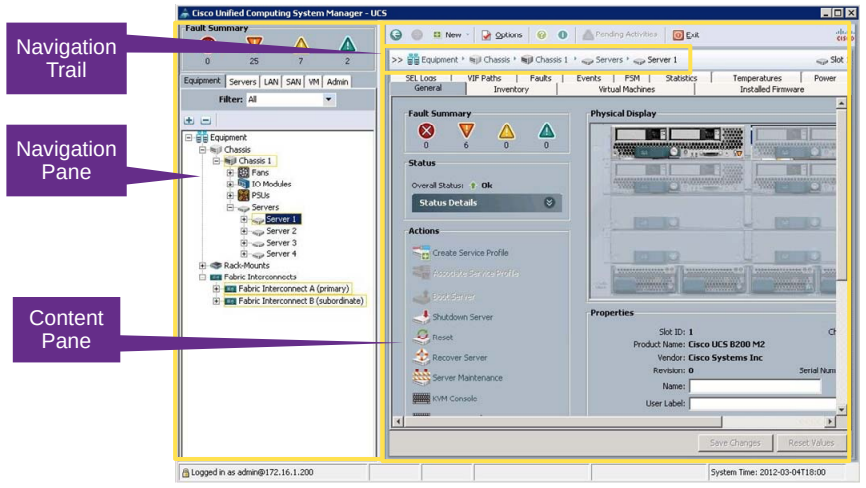
© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Layout of the Cisco UCS Manager GUI

This topic describes the basic layout of Cisco UCS Manager.

Primary Components of Cisco UCS Manager GUI

The Cisco UCS Manager GUI is divided into two major panes.



The navigation pane consists of a fault summary bar across the top view and a series of six tabs. These tabs offer differing views of the various managed Cisco UCS components. The fault summary has four conditions: critical, major, minor, and warning. The fault summary contains the cumulative totals for all Cisco UCS components.

The expandable branch (or tree) functionality allows the administrator to traverse the various components in the six tabs.

The content pane consists of a top toolbar with a back button, new object creation drop-down list, question button, and information button.

The second toolbar in the content pane offers the operator a navigation trail of object hierarchies that have already been traversed. This toolbar includes the ability to rapidly return to a previous location along the trail. The current location is at the right-most portion of this bar.

The largest part of the content pane offers granular detail that is associated with varying objects that have been selected in the navigation pane. At the bottom of the content pane are function buttons that are associated with committing, saving, and discarding configuration change.

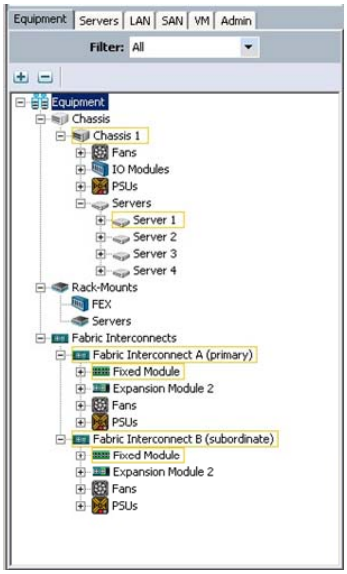
https://t.me/learningnets

Navigation Window Tabs

This topic describes which configuration elements are configured under the six tabs of the navigation pane.

Exploring the Equipment Tab

- Use the + and - to move up and down the equipment hierarchy.
- The Equipment tab allows you to interact with your equipment as physical elements.



© 2012 Cisco and/or its affiliates. All rights reserved.

The Equipment tab allows you to select and interact with Cisco UCS components at the physical level. As you expand into the equipment hierarchy, any element that you select is featured in the content pane.

These major categories are on the Equipment tab:

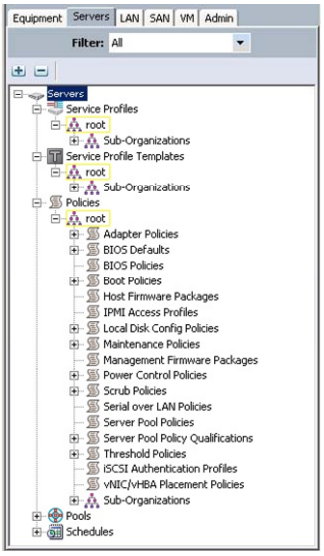
- Chassis
- Rack Mounts
- Fabric Interconnects

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Exploring the Servers Tab

- On the Servers tab, create service profiles, templates, policies, resource, and identity pools.
- Items on the Servers tab are logical configuration elements applied to physical servers.



© 2012 Cisco and/or its affiliates. All rights reserved.

On the Servers tab, Cisco UCS administrators create, modify, and delete logical server components and apply them to physical servers.

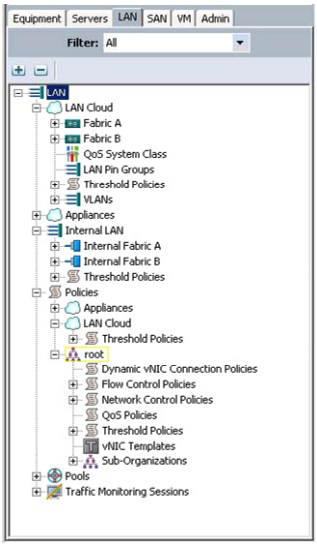
These major categories are on the Servers tab:

- Service Profiles
- Service Profiles Templates
- Policies
- Pools
- Schedules

https://t.me/learningnets

Exploring the LAN Tab

- On the LAN tab, define quality of service (QoS), create VLANs, policies, and identity pools.
- Items on the LAN tab are logical configuration elements applied to physical servers.



© 2012 Cisco and/or its affiliates. All rights reserved.

On the LAN tab, Cisco UCS administrators create, modify, and delete configuration elements that are associated with the Ethernet network.

These major categories are on the LAN tab:

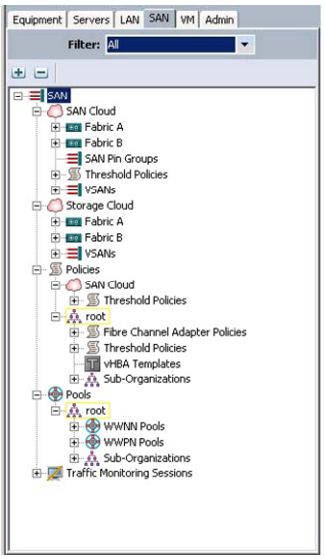
- LAN Cloud
- Appliances
- Policies
- Pools
- Internal LAN
- Traffic Monitoring Sessions

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Exploring the SAN Tab

- On the SAN tab, define QoS, create VSANs, policies, and identity pools.
- Items on the SAN tab are logical configuration elements applied to physical servers.



© 2012 Cisco and/or its affiliates. All rights reserved.

On the SAN tab, Cisco UCS administrators create, modify, and delete configuration elements that are associated with the Fibre Channel SAN and Fibre Channel over Ethernet (FCoE) communications.

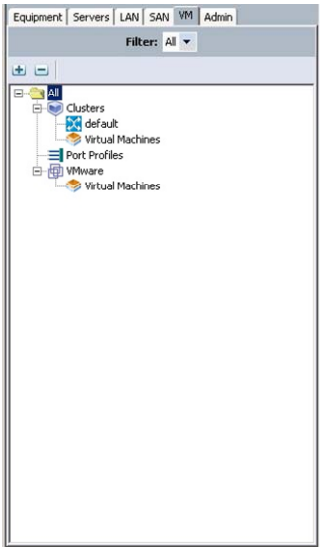
These major categories are on the SAN tab:

- SAN Cloud
- Storage Cloud
- Policies
- Pools
- Traffic Monitoring Sessions

https://t.me/learningnets

Exploring the VM Tab

- On the VM tab, define vCenter server and port profiles for dynamic vNICs that can be vMotioned from host to host by using VMware pass-through switching.
- Items on the VM tab are logical configuration elements applied to physical network adapters and VMware vCenter.



© 2012 Cisco and/or its affiliates. All rights reserved.

On the VM tab, Cisco UCS administrators create, modify, and delete configuration elements that are associated with VMware vSphere. Beginning with Cisco UCS release 1.2, the Cisco UCS M81KR Virtual Interface Card supports dynamic virtual network interface cards (vNICs) and VMware pass-through switching.

These major categories are in the VM tab:

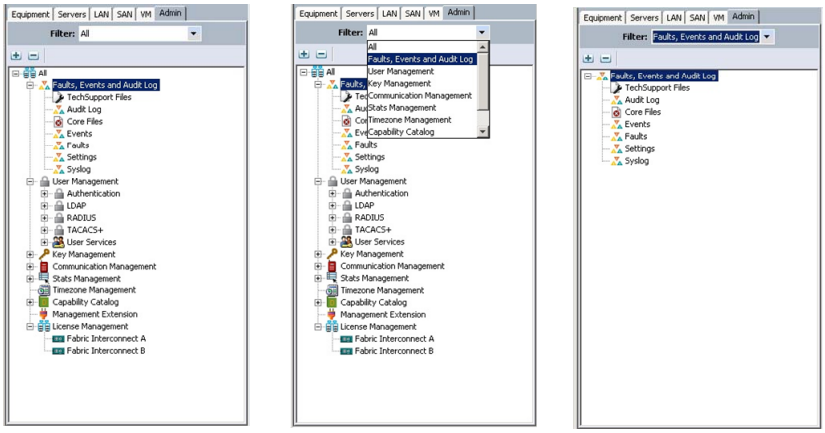
- Clusters
- Port Profiles
- Virtual Machines

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Exploring the Admin Tab

- On the Admin tab, perform administrative tasks.
- As with the other navigation tabs, you can apply filters to limit the scope of the display.



© 2012 Cisco and/or its affiliates. All rights reserved.

On the Admin tab, Cisco UCS administrators create, modify, and delete configuration elements that are associated with general Cisco UCS administration tasks. With so many major categories in the Admin tab, making use of category filtering to display only the elements of interest is convenient. Filtering is available in all navigation tabs, but is demonstrated only on this tab in this lesson.

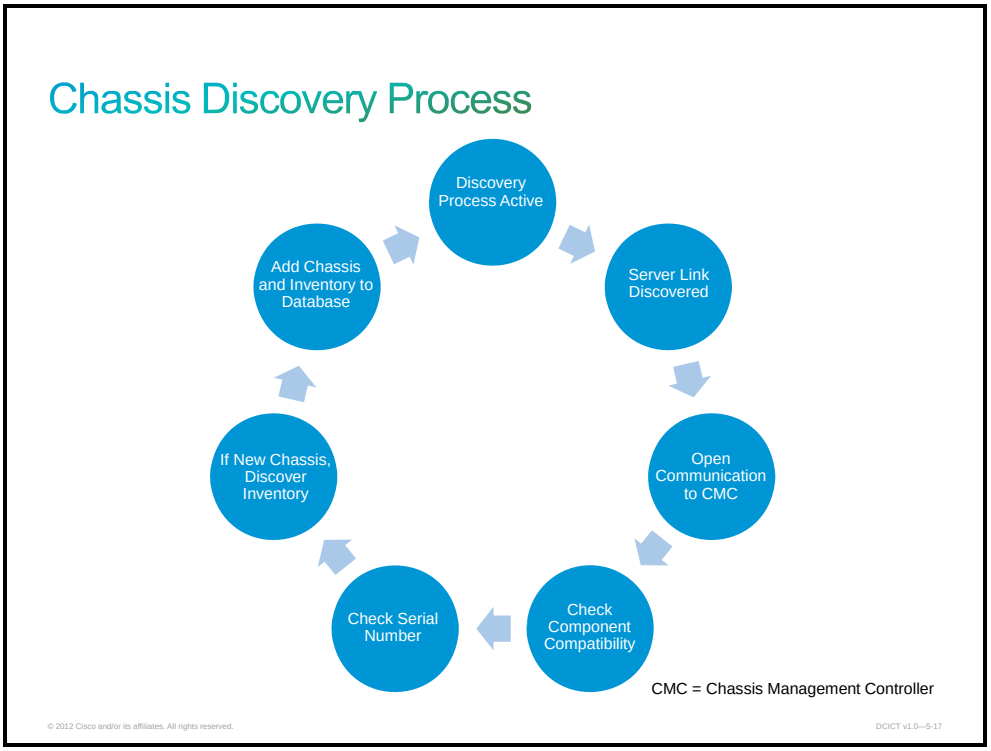
These major categories are on the Admin tab:

- Faults, Events, and Audit Log
- User Management
- Key Management
- Communications Management
- Stats Management
- Time zone Management
- Capability Catalog
- License Management

https://t.me/learningnets

Device Discovery in Cisco UCS Manager

This topic describes how Cisco UCS Manager discovers components.



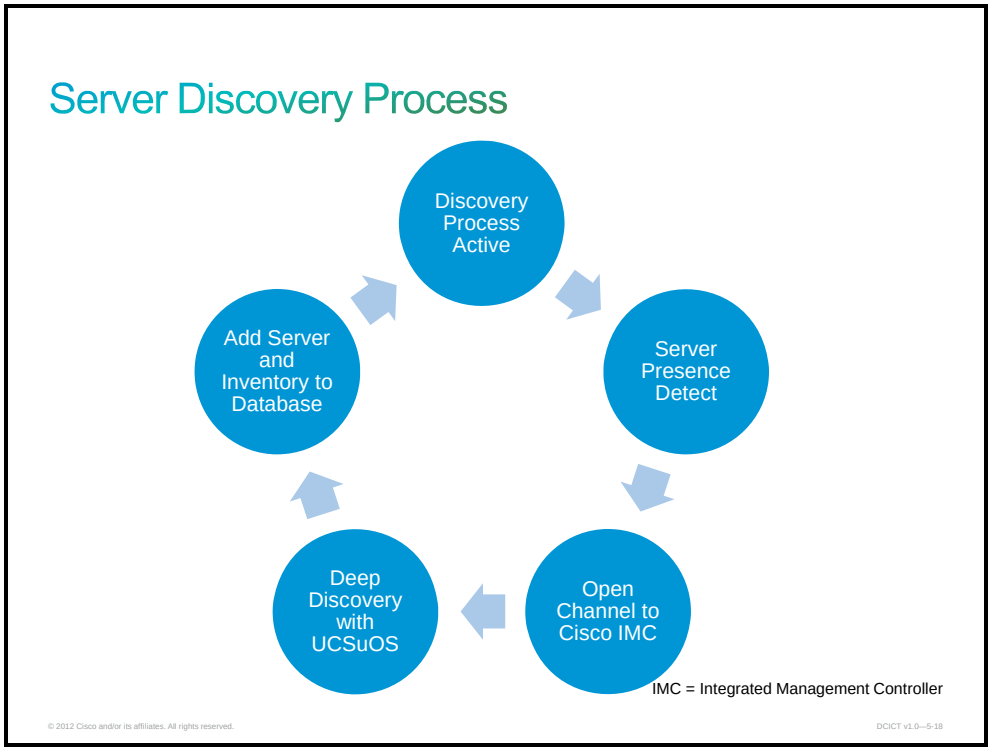
Cisco UCS Manager is always in discovery mode. Through a series of presence sensors and voltage indications, Cisco UCS Manager knows when new components are added or removed.

Chassis Discovery

During the chassis discovery process, these components are identified and, if necessary, verified:

- The presence sensor discovers a link on a server-facing port.
- Open basic communication to the chassis management controller in the I/O module.
- Check component compatibility.
- Check that this device is supported.
- Check the serial number on the chassis.
  - Check whether the administrator has removed this device from management.
- This new link is to an existing chassis:
  - If the chassis is new, accept the chassis.
  - Perform discovery of components such as model, firmware, fans, power supply units (PSUs), and slots.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets



Discovery of Server in Chassis

During the discovery of a server in the chassis, these steps are performed:

1. If a slot raises a state, then the slot has presence (a server is present).
2. Open communication with the Cisco Integrated Management Controller on the server in that slot.
3. Discover server information:

   Vendor, model, serial number, BIOS, Cisco Integrated Management Controller, CPUs, memory, host bus adapters (HBAs), network interface cards (NICs), and local storage controller.
4. Boot into Cisco UCS Utility Operating System
5. Perform deep discovery:

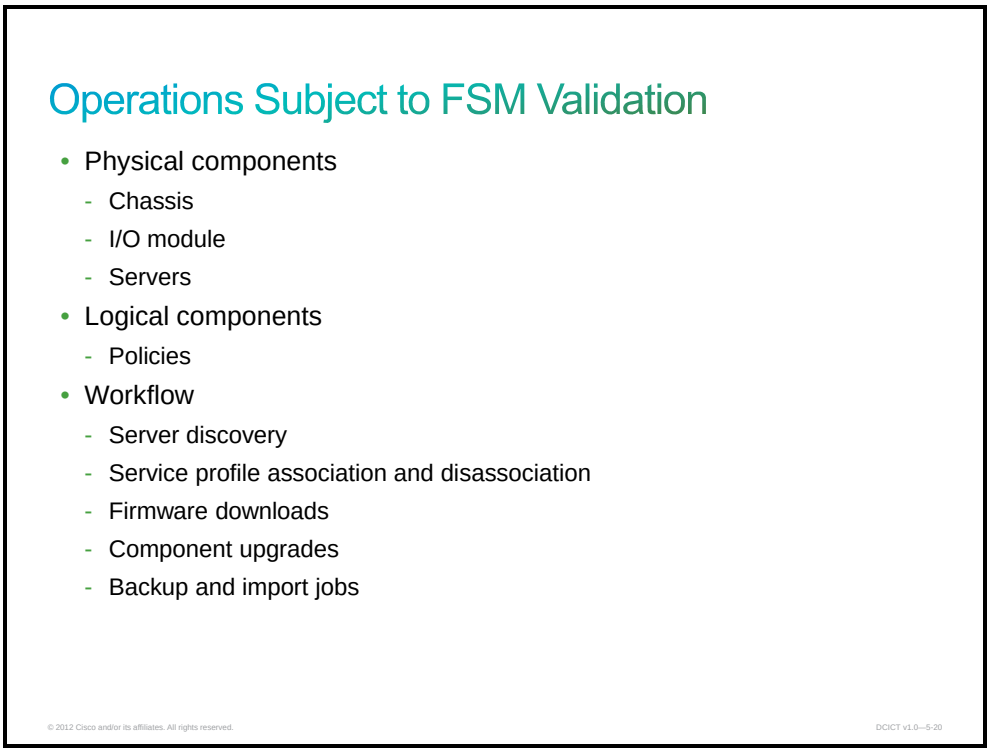
   Local disk information, vendor-specific HBA, and NIC information

Note	The Cisco UCS Utility Operating System is triggered at device discovery, service profile association, and service profile disassociation. This small utility operating system resides in fabric interconnect flash memory and is booted in a Preboot Execution Environment (PXE) when needed by a blade server.

https://t.me/learningnets

Verifying Device Discovery in Cisco UCS Manager

This topic describes how Cisco UCS Manager discovers components by monitoring finite state machines (FSMs).



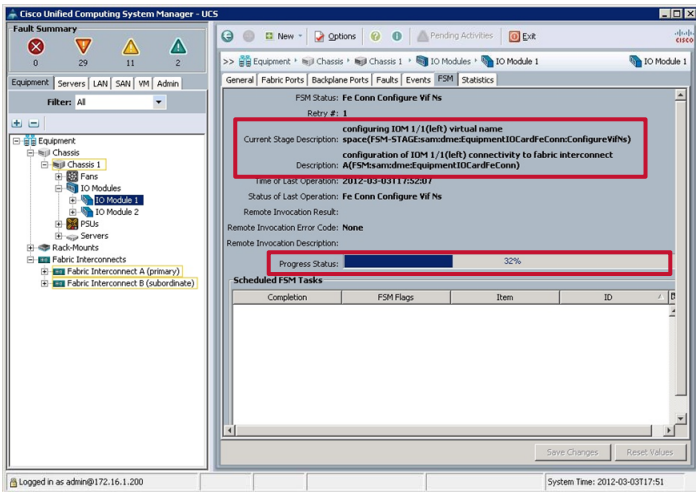
Many components and processes within Cisco UCS are characterized by highly complex state transitions. FSMs are assigned to audit the state transitions and validate correct operation.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Reset I/O Module: Discovery Process in FSM

FSM tracking the transition states of I/O module discovery.



© 2012 Cisco and/or its affiliates. All rights reserved.

DCICT v1.0—5-21

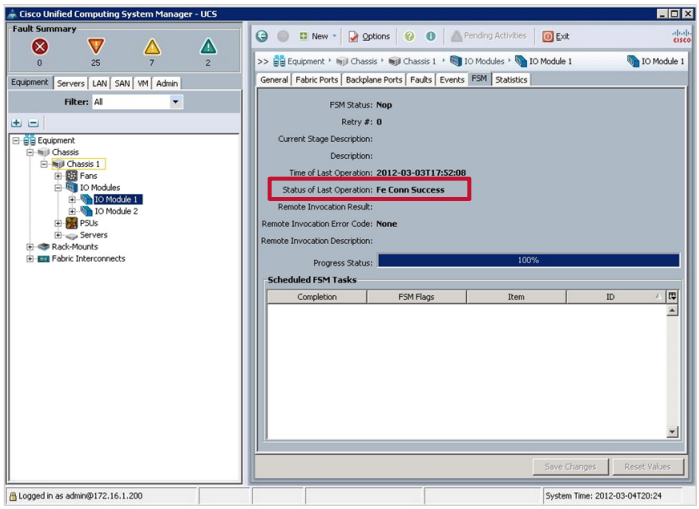
In the figure, the I/O module for fabric A was manually reset. The current stage description clearly indicates that chassis discovery is underway. The description field below this description indicates that I/O module connectivity is being restored.

The progress indicator provides a graphical representation of how far the FSM processed tree has proceeded. In many cases, there might be a long pause at a particular percentage point. This pause is process-specific and usually nothing to worry about. If the FSM stops responding for more than a certain amount of time, a timeout will occur and the operation fails.

https://t.me/learningnets

Reset I/O Module: Discovery Successful

FSM reaches 100% with no errors.



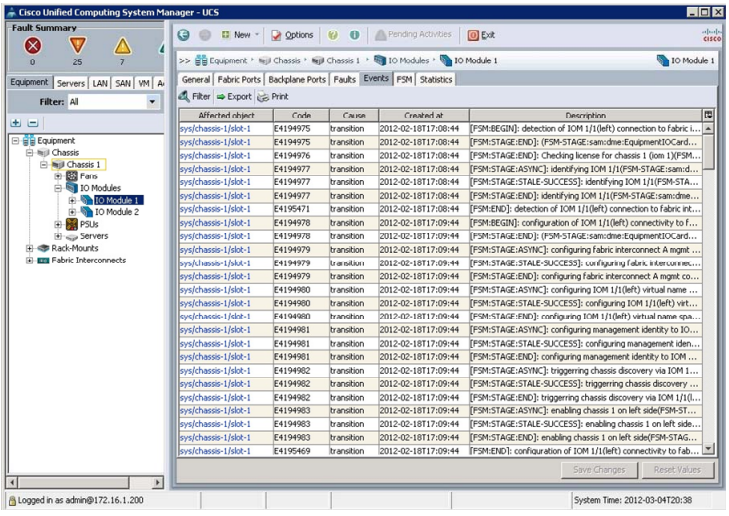
© 2012 Cisco and/or its affiliates. All rights reserved.

DCICT v1.0—5-22

The figure illustrates a successful reset of the I/O module and the resulting discovery process. The result code, feConnSuccess, might seem cryptic, but it means that the fabric extender (FEX) connection was a success. FEX is a term that is sometimes applied to the I/O module. As you continue to work with FSM output, it will be apparent that the Cisco personnel who wrote the messages took great care to make them readable by humans.

FSM Discovery Details

Click the Event tab to review the log of FSM state transitions.



© 2012 Cisco and/or its affiliates. All rights reserved.

DCICT v1.0—5-23

In many cases, the FSM state transitions occur too fast for a human operator to read. Simply click the Events tab to see a detailed list of all the transitions.

https://t.me/learningnets

Summary

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

Summary

- The Cisco UCS Manager GUI is the interface from which customers provision, maintain, and monitor Cisco UCS B-Series blade servers and Cisco UCS C-Series integrated rack servers.
- The Cisco UCS Manager GUI is divided between the navigation pane and the content pane.
- The navigation pane contains six tabs: Equipment, Servers, LAN, SAN, VM, and Admin.
- Cisco UCS Manager discovers new Cisco B-Series components by interacting with the chassis management controller on the I/O module.
- Components are discovered by Cisco UCS Manager and subject to validation by FSMs.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Lesson 6

Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles

Overview

The purpose of this lesson is to describe features of Cisco Unified Computing System (UCS) Manager that allow rapid provisioning and consistent application of policy.

Objectives

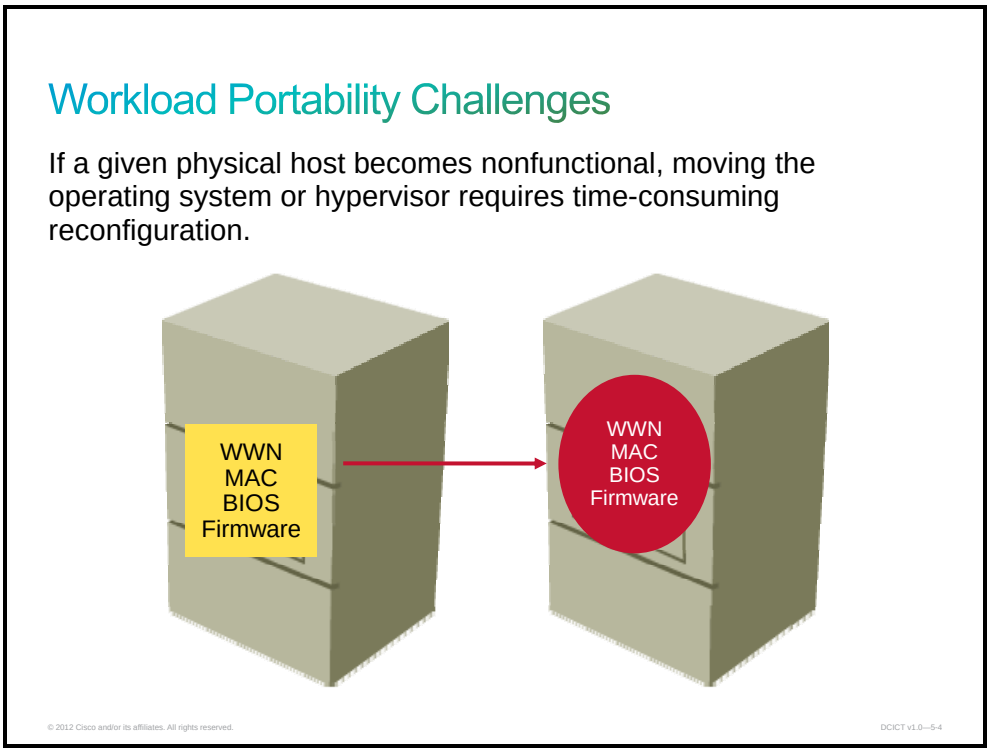
Upon completing this lesson, you will be able to describe features of Cisco UCS Manager that allow rapid provisioning and consistent application of policy. You will be able to meet these objectives:

- Describe how the abstraction of hardware identities allows fast server mobility from one compute node to another
- Describe how pooled resources ensure uniqueness of identities that are assigned to service profiles
- Describe how service profile templates allow rapid scaling of Cisco UCS
- Describe the creation of polices that map to service profiles and templates
- Describe the creation of chassis and blade power-capping policies

https://t.me/learningnets

Benefits of Stateless Computing

This topic describes how the abstraction of hardware identities allows fast server mobility from one compute node to another.

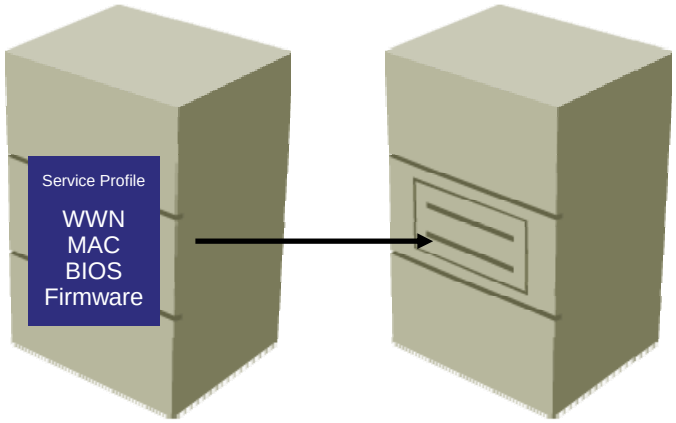


In a traditional data center computing environment, the operating system or hypervisor is installed on a compute node. Burned-in values for universally unique identifiers (UUIDs), MAC addresses, world wide nodes (WWNs), and port names are consumed during the installation process. If the underlying compute node becomes nonfunctional, then the replacement compute node will have different values for UUID, MAC addresses, and WWNs. This difference can lead to time-consuming and expensive downtime for a server. In a virtualized server, this difference can also affect many hosts in virtual machines (VMs).

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

Stateless Computing

In a Cisco UCS stateless computing environment, hardware identifiers and firmware versions are abstracted into a service profile that can rapidly be moved to a replacement compute node.

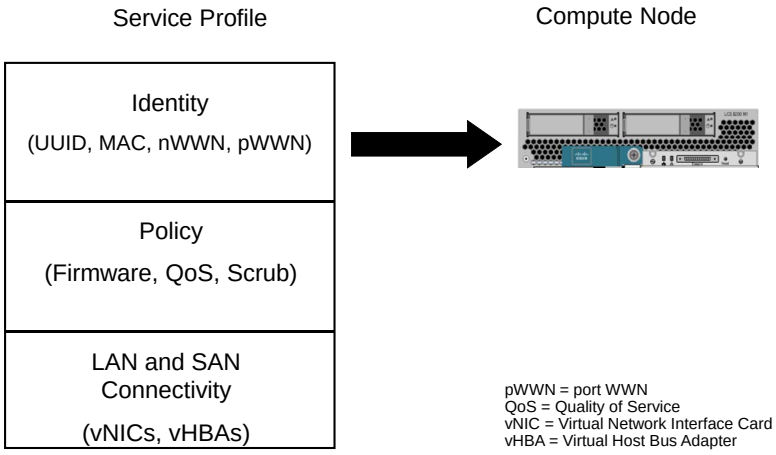


© 2012 Cisco and/or its affiliates. All rights reserved.

Cisco UCS offers a new paradigm for data center computing. Because all the identity properties (MAC, UUID, and WWNs) can be administered locally, a SAN-booted server can be moved to a replacement compute node with minimal downtime. This move is referred to as stateless computing.

Contents of a Service Profile

Each compute node requires a unique service profile.



© 2012 Cisco and/or its affiliates. All rights reserved.

The service profile is the software construct that contains the identity and operational policies that make up a compute node. Each compute node in a Cisco UCS B-Series or C-Series integrated server requires a unique service profile to operate.

https://t.me/learningnets

Describe the Benefits of Service Profiles

- Plan and preconfigure once.
  - Collaborative server, storage, and network planning and configuration
- Incremental deployment is nondisruptive.
- Ease of server replacement: Simply replace the physical blade and re-associate the service profile.
- Ease of server upgrades: Move the service profile to a server with more capabilities.

© 2012 Cisco and/or its affiliates. All rights reserved.

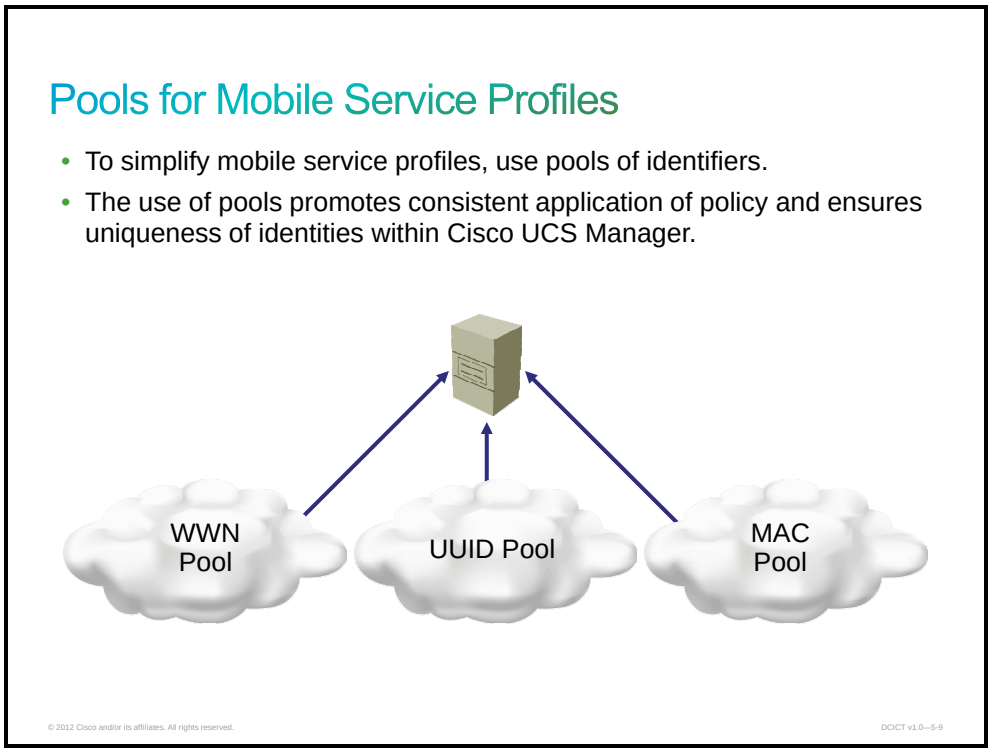
The primary benefit of service profiles is server mobility (stateless computing).

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Using Identity Pools in Service Profiles

This topic describes how pooled resources ensure the uniqueness of the identities that are assigned to service profiles.

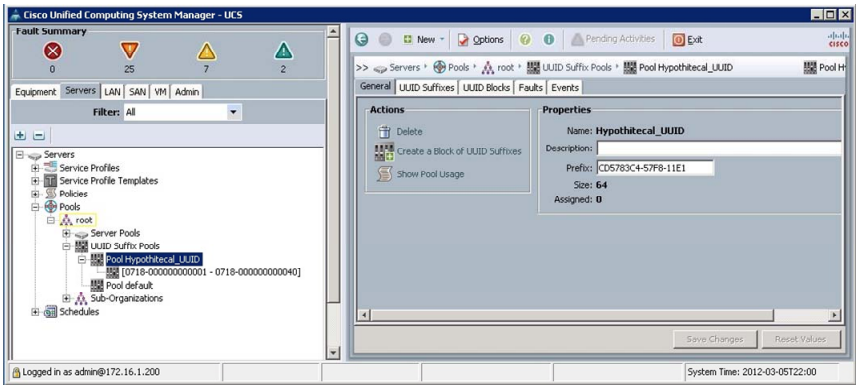


Stateless computing requires unique identity resources for UUID, MAC address, and WWNs for Fibre Channel. Using pooled resources ensures consistent application of policy and reasonable assurances that identities are unique within Cisco UCS Manager.

https://t.me/learningnets

Logical Resource Pools

Identity pools supply abstracted identities to service profiles in service profile templates.

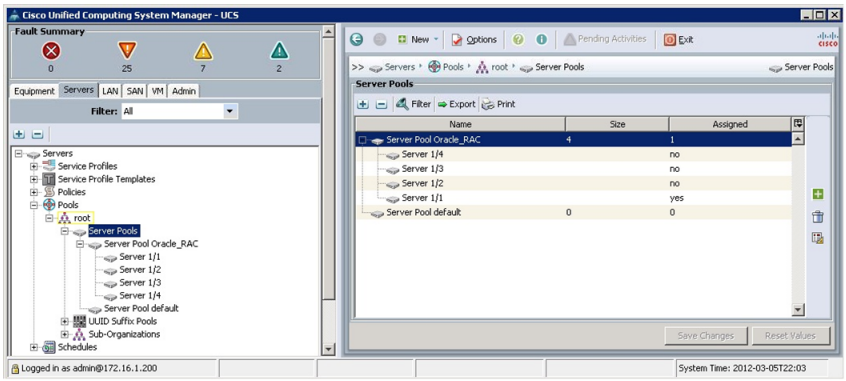


© 2012 Cisco and/or its affiliates. All rights reserved.

Logical resource pools (identity pools) provide abstracted identities that service profiles consume in service profile templates, to facilitate stateless computing.

Physical Resource Pools

Physical resource pools provide blade server resources to service profiles to facilitate rapid provisioning.



© 2012 Cisco and/or its affiliates. All rights reserved.

Physical resource pools are used to create blade-server groupings that are based on arbitrary administrative criteria. These pools can be used in conjunction with service profile templates to provide rapid provisioning of compute resources.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

UUID Use

- UUIDs are essentially standardized serial numbers that identify a particular server.
- Traditional servers have a hardware UUID stored in the system BIOS.
- Operating systems and software licensing schemes may use the UUID to detect whether they have been moved between physical servers.
- Cisco UCS allows for the manual or automatic assignment of UUIDs to enhance mobility of operating systems and applications.

© 2012 Cisco and/or its affiliates. All rights reserved.

UUIDs are designed as globally unique identifiers (GUIDs) for each compute node on a network. UUIDs are used in a number of different ways, but in the context of Cisco UCS, the UUID refers to a 128-bit identifier that is coded into the compute node BIOS.

Operating systems, hypervisors, and applications can leverage the UUID for processes such as activation, internal disk labels, and so on. Some applications use the UUID as an internal root value that is propagated tightly within data structures. For these reasons, UUIDs should be administered locally in the service profile instead of being derived from the BIOS. UUIDs within a service profile allow for mobile computing. If the underlying compute node fails, then the service profile carries the UUID to the replacement compute node.

https://t.me/learningnets

UUID Format

- UUIDs are 128-bit numbers (represented in hexadecimal), expected to be globally unique.
- Many schemes exist to define or generate the UUID.
- Cisco UCS Manager uses a configurable 64-bit prefix and allows administrators to specify a range of suffixes for use by compute nodes.
- Cisco recommends that prefixes be set to the same 24-bit organizationally unique identifier (OUI) as used in WWN pools, and padded as necessary.

Sample UUID in traditional colon-separated hexadecimal:
32:1F:66:B4:B7:FF:1E:D4:56:38:07:17:65:10:06:0C

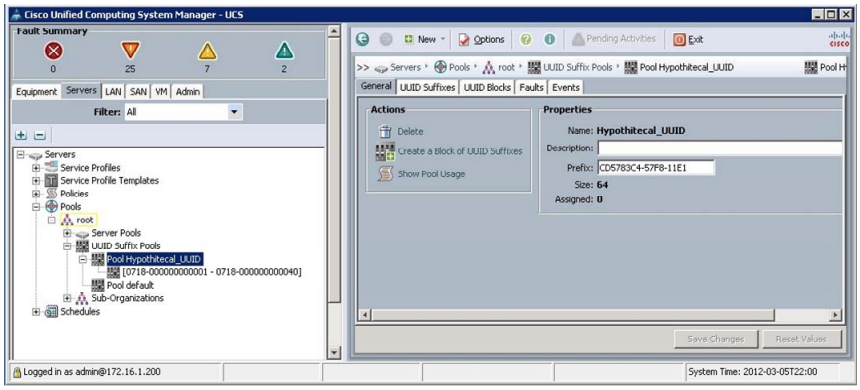
Sample UUID as displayed in Cisco UCS Manager:
321F66B4-B7FF-1ED4-5638-07176510060C

© 2012 Cisco and/or its affiliates. All rights reserved.

There are many schemas to deploy and format UUIDs. It is the responsibility of Cisco UCS to determine which values to encode in the UUID prefix and suffix.

UUID Pool

UUIDs are now available for assignment in a service profile or service profile template.



© 2012 Cisco and/or its affiliates. All rights reserved.

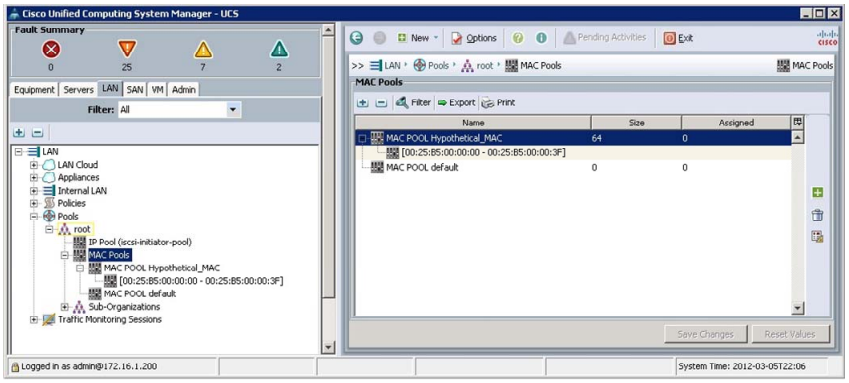
The figure shows that a pool of UUIDs was created and is immediately available for consumption by service profiles or service profile templates.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

MAC Address Block

A MAC pool is now available for assignment.



© 2012 Cisco and/or its affiliates. All rights reserved.

The figure shows that a MAC address pool was created to supply MAC addresses to compute node Ethernet communication. Service profiles or service profile templates can now consume this pool.

WWN Format

- WWN is a 64-bit address.
- WWNs use an extended format: 2X:XX:YY:YY:YY:ZZ:ZZ:ZZ
  - **Example:** 20:00:00:25:B5:20:20:00
- X = organizationally assigned
- YY:YY:YY = OUI
- ZZ:ZZ:ZZ = organizationally assigned
- A unique WWNN is assigned to each HBA.
- Each port on the HBA requires a unique pWWN.

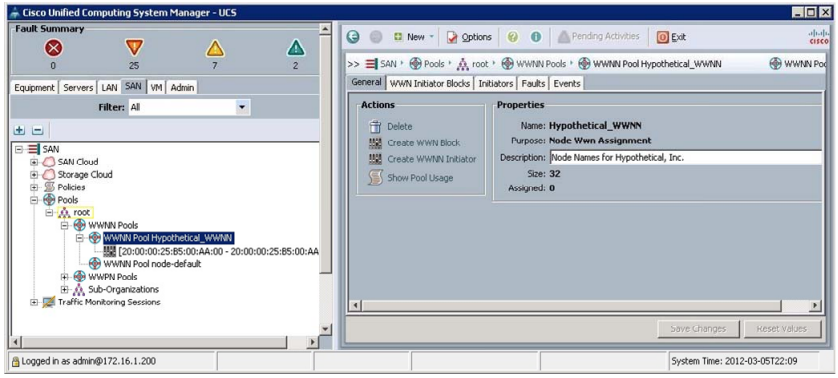
© 2012 Cisco and/or its affiliates. All rights reserved.

WWNs are 64-bit addresses that have many possible formats. The example in the figure is for reference only. Cisco UCS manager enforces a specific format in all WWN pools.

https://t.me/learningnets

nWWN Block

The WWNN Pool is now available for assignment.

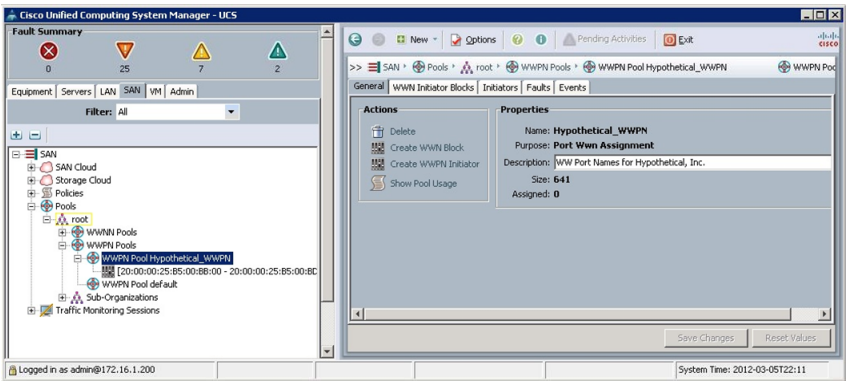


© 2012 Cisco and/or its affiliates. All rights reserved.

The figure shows that a node WWN (nWWN) pool, called a WWNN pool in the GUI, was created. These addresses are immediately available for assignment to a service profile or service profile template.

pWWN Block

The WWPN pool is now available for assignment.



© 2012 Cisco and/or its affiliates. All rights reserved.

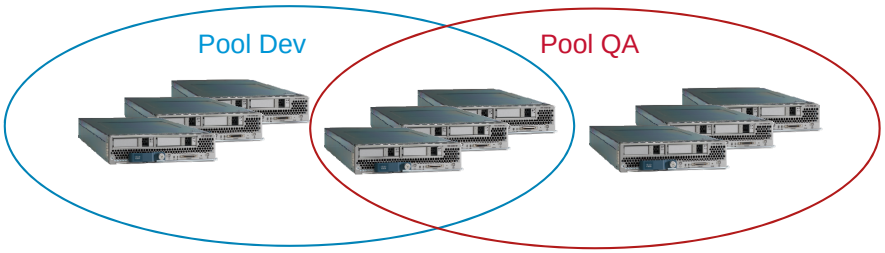
The figure shows that a port WWN (pWWN) pool, called a WWPN pool in the GUI, was created. These addresses are immediately available for assignment to a service profile or service profile template.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Server Pools

- Server pools can be manually populated or autopopulated.
- Blade server can be in multiple pools at the same time.
- Associate a service profile with a pool:
  - A compute node is selected automatically from the pool.
  - Cisco UCS Manager selects only a blade server that is neither associated with another logical server nor in the process of being disassociated.



© 2012 Cisco and/or its affiliates. All rights reserved.

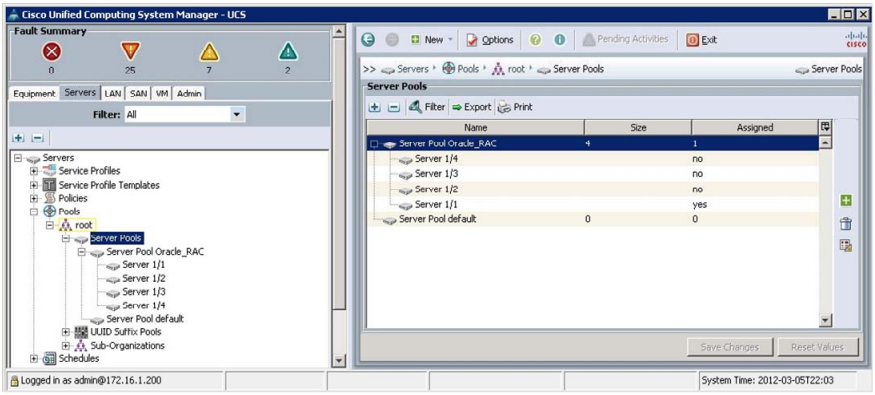
A compute node can be in multiple pools at the same time. The profile that is associated to a specific compute node owns the node, regardless of the number of pools in which the blade server resides.

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 that is discovered but not associated with any profile or in the process of being associated or disassociated.

https://t.me/learningnets

Server Pool Created

The servers in the Oracle_RAC pool can now be consumed by service profiles or service profile templates.



© 2012 Cisco and/or its affiliates. All rights reserved.

The content pane displays the IDs of the servers that are added to the pool. This pane includes information about whether the server has been assigned. If so, a link to the service profile of the server is displayed in the Assigned To column.

© 2012 Cisco Systems, Inc.

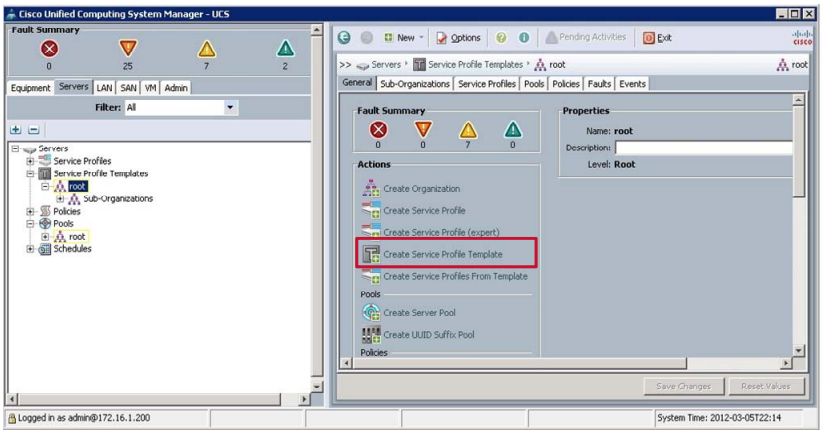
https://t.me/learningnets

Using Service Profile Templates to Enable Rapid Provisioning and Consistent Application of Policy

This topic describes how service profile templates allow rapid scaling of Cisco UCS.

Creating a Service Profile Template

- Similar to service profile creation
- Requires server pool (cannot select individual blade)
- Requires pooled identities (UUID, MAC, WWNN, WWPN)



© 2012 Cisco and/or its affiliates. All rights reserved.

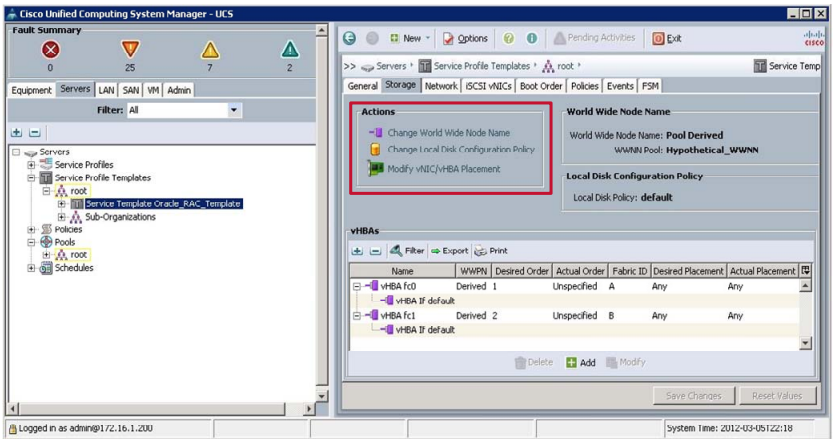
The process of creating a service profile template is nearly identical to creating a service profile manually. The principal differences are that service profile templates cannot be directly applied to a compute node and that no hardware elements can use a derived value.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Modify Template

Pool assignments and policies can be changed after template creation.



© 2012 Cisco and/or its affiliates. All rights reserved.

After a service profile template has been generated, it can be modified like a manually created service profile. Certain changes that are made to an updating template are propagated to every service profile that the template generates.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Using Templates for Differentiated Policy

The use of templates allows the consistent application of policy to meet application requirements in heterogeneous computing environments.

Oracle RAC	SAP	Web Server
B250-M2	B230-M2	B200-M3
256-Gb/s RAM	128-Gb/s RAM	8-Gb/s RAM
Local HD for swap	No local HD	Local HD
M81-KR VIC x2	M72-KR-Q	M72-KR-E
Xeon 5660 x2	Xeon E7-2860 x2	Xeon E5-2620 x1
RSS support	RSS support	RSS support
Intel VT-d BIOS on	Intel VT-d BIOS off	Intel VT-d BIOS off
Jumbo frames	Jumbo frames	Standard MTU

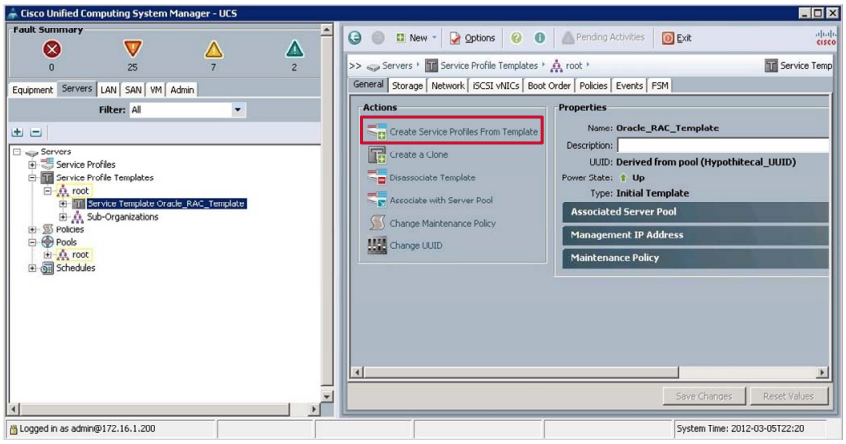
© 2012 Cisco and/or its affiliates. All rights reserved.

The example in the figure illustrates the power and flexibility of using templates for differentiated policy. Because groups of applications share similar or identical requirements, service profile templates can be created to seamlessly provide the identity of server resources that are needed to serve the application. One of the important operational benefits to this approach is consistency of policy across the entire class of application.

https://t.me/learningnets

Creating Service Profiles from Template

Creating from template allows simultaneous provisioning of multiple service profiles.



© 2012 Cisco and/or its affiliates. All rights reserved.

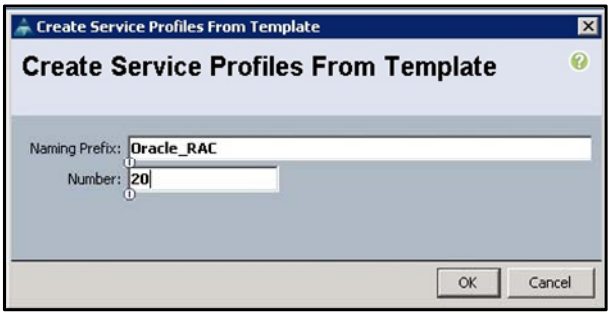
Another benefit to using service profile templates is that the Cisco UCS administrator can automate the provisioning of one to hundreds of compute nodes into simple operations. After a service profile template is built and points to identity and resource pools with sufficient resources, automation can begin. In the navigation pane, choose the service profile template and the organization in which the new service profiles are to be created. In the content pane, click **Create Service Profiles from Template**.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Select Prefix and the Number of Profiles

- Provide naming prefix and number of service profiles, based on the service profile template.
- Example creates service profiles Oracle_RAC1 through Oracle_RAC20.



© 2012 Cisco and/or its affiliates. All rights reserved.

The dialog box that appears prompts you for a naming prefix and the number of service profiles to be generated with that prefix. Immediately after you click OK, service profiles appear under the organization. If the server assignment in the template points to a server pool, then a new service profile immediately begins to associate with the next available server in the pool.

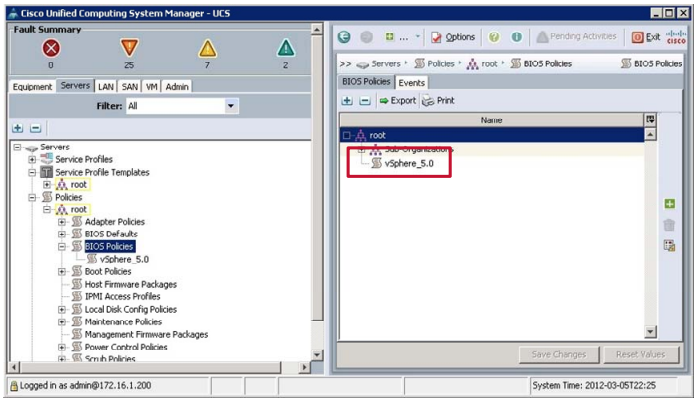
https://t.me/learningnets

Creation of Policies for Service Profiles and Service Profile Templates

This topic describes creating polices that map to service profiles and service profile templates.

Locate BIOS Policies

Expand **Policies** on the Servers tab and locate **BIOS Policies**.



© 2012 Cisco and/or its affiliates. All rights reserved.

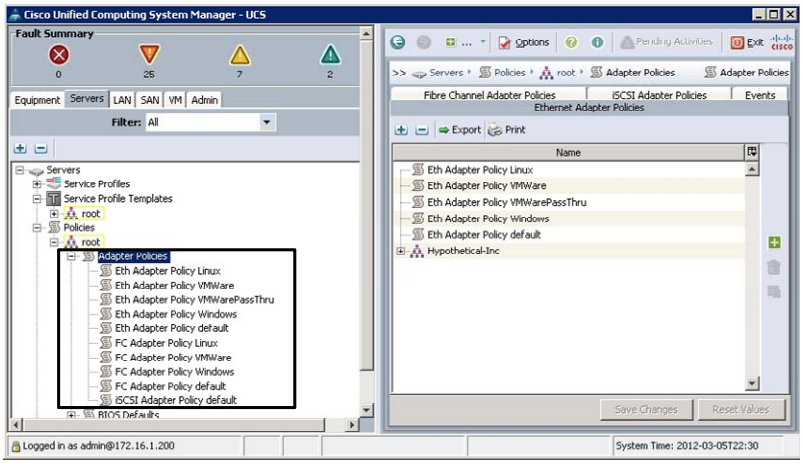
In the Servers tab of the navigation pane, expand the policies to locate BIOS policies. The BIOS policy in the figure is available for assignment to a service profile or service profile template.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Locate Adapter Policies

- Expand **Policies** on the Servers tab and locate **Adapter Policies**.
- Right-click the organization in which the new policy is to be created.



© 2012 Cisco and/or its affiliates. All rights reserved.

Expand policies in the Servers tab, and locate Adapter Policies.

https://t.me/learningnets

Chassis and Blade Power Capping

This topic describes creating chassis and blade power-capping policies.

Need for Power Capping with Cisco UCS B-Series

- With the introduction of the Cisco UCS B230 M2 and B440 M2 Blade Servers, power can be oversubscribed.
- Power capping allows the chassis to dynamically limit power to blades, based on policy and priority.
- When a power cap is applied to a blade, the CPU steps down to a lower power or performance level.



© 2012 Cisco and/or its affiliates. All rights reserved.

The current Cisco UCS chassis is designed with a maximum power limit of 550 W per half-slot and 5000 W per chassis. Newer, denser blades such as the Cisco UCS B230 M2 and Cisco UCS B440 M2 Blade Servers can draw 600 to 800 W, depending on the configuration of the blades. Because of higher power requirements for these newer blades, power capping might be necessary to maintain operation on all the blades in a chassis.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Strategies for Power Capping

There are three methods of power capping.

Type	Details
Explicit group power cap	• Single power cap for multiple chassis in Cisco UCS • Set by customer • Policies for multi-chassis cap
Explicit blade power cap	Blade-level power cap
Implicit chassis power cap	Cisco UCS Manager automatically sets cap on blade or chassis

© 2012 Cisco and/or its affiliates. All rights reserved.

There are three types of power capping:

- **Explicit group:** This power cap is applied to a group through the use of policies. The purpose is to make sure that you never draw more power than you have budgeted for.
- **Explicit blade:** The type caps the power at a specific blade slot.
- **Implicit chassis:** This type is the default power cap (basically, the maximum capacity) if no other power cap is defined.

There are significant advantages to using group power capping:

- **Scalable:** Power caps are set globally rather than individually for each blade.
- **Dynamic reallocation:** The chassis puts the power where it is needed.
- **Safe caps:** Cisco UCS Manager makes sure that there is enough power for booting.

https://t.me/learningnets

Summary

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

Summary

- The abstraction of hardware identities allows fast server mobility from one compute node to another because the identities are locally administered.
- The use of pooled identity resources ensures the uniqueness of identities assigned to service profiles.
- Service profile templates allow rapid provisioning of blade and rack servers.
- Policy objects are consumed by service profiles and service profile templates and allow for differentiated policy.
- Chassis and blade power-capping policies allow the Cisco UCS administrator to manage power consumption of Cisco UCS 5108 chassis and Cisco UCS B-Series blade servers.

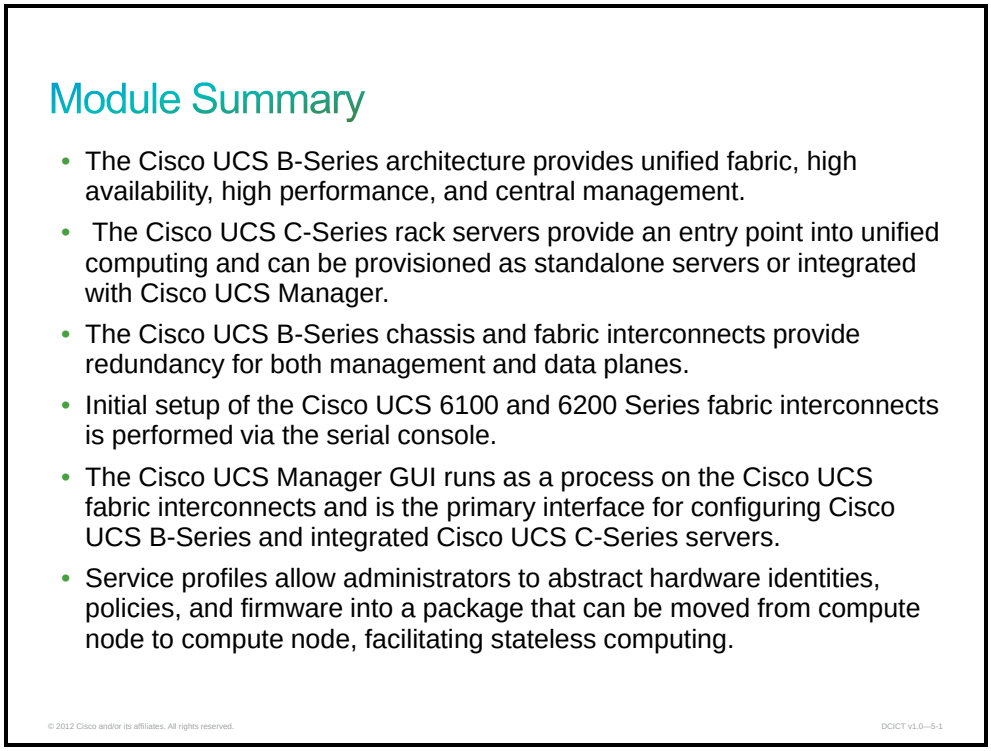
© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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



In addition to providing innovative, data center routing and switching infrastructures, Cisco offers a lineup of data center blade and rack servers that are designed for increased performance, manageability, and superior return on investment (ROI).

The initial Cisco Unified Computing System (UCS) product portfolio includes the Cisco UCS 6100 Series Fabric Interconnects, 5108 Blade Server Chassis, 2204XP I/O module, and B-Series Blade Servers. Cisco UCS B-Series Blade Servers can be configured with local storage and a comprehensive portfolio of mezzanine adapters for LAN and SAN communication. Second-generation hardware allows Cisco UCS B-Series servers to scale to 80- or 160-Gb/s bandwidth, depending on whether one or two Cisco UCS Virtual Interface Card 1280 adapters are installed.

Cisco UCS C-Series Rack Servers provide an entry point into the Cisco UCS product portfolio. Cisco UCS C-Series Rack Servers can take advantage of features such as unified fabric to reduce cable counts and power and cooling requirements in the data center. If integrated with Cisco UCS Manager, Cisco UCS C-Series Rack Servers can take advantage of additional advanced capabilities, including centralized management and stateless computing.

Blade-server connectivity from the Cisco UCS 5108 chassis to the fabric interconnects is made possible by the links on the Cisco UCS 2104XP or 2208XP I/O modules. The number of links and their specific configuration directly influence the bandwidth oversubscription rate and the resulting bandwidth that is available to blade servers.

To ensure continuous operation, Cisco UCS 6100 Series and 6200 Series Fabric Interconnects can be configured in a high-availability cluster. Initial configuration of the cluster begins with a connection to the serial console of each cluster peer. A setup script guides the administrator through the configuration tasks of provisioning the cluster on each fabric interconnect.

https://t.me/learningnets

Cisco UCS Manager is an application that runs on the fabric interconnects. Cisco UCS Manager is implemented as an active and subordinate system, with the configuration database active on both cluster peers. Cisco UCS Manager is a single point of management for Cisco UCS B-Series Blade Servers and C-Series Rack Servers with which it is integrated. As devices are discovered and enumerated, Cisco UCS Manager adds them to the configuration database as managed objects. All configuration data is stored in XML. Cisco offers a software development kit (SDK) and full documentation to the Cisco UCS Manager XML application programming interface (API).

Stateless computing is one of the unique signature features of Cisco UCS. Operating systems and hypervisors are not consuming burned-in identities for resources such as MAC addresses, world wide names (WWNs), and universally unique identifiers (UUIDs). These identities are locally administered and pooled to facilitate uniqueness of identities within Cisco UCS. Identity and policy are stored in a construct that is called the service profile. If a compute node becomes nonfunctional, its service profile can simply be re-associated to a new node and rebooted with its identity, policies, and even firmware versions maintained.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. *Cisco UCS B-Series Blade Servers Data Sheets.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/install.html.
- Cisco Systems, Inc. *Cisco UCS B-Series Blade Servers Adapter Data Sheets.* http://www.cisco.com/en/US/products/ps10280/products_data_sheets_list.html.
- Cisco Systems, Inc. *Cisco UCS 5108 Server Chassis Installation Overview – Installing the Cisco UCS 5108 Server Chassis.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/install.html.
- Cisco Systems, Inc. *Cisco UCS B200 Blade Server Installation and Service Note.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/blade.html.
- Cisco Systems, Inc. *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/fullblade.html.
- Cisco Systems, Inc. *Cisco UCS B230 Blade Server Installation and Service Note.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/B230.html.
- Cisco Systems, Inc. *Cisco UCS B440 High Performance Blade Server Installation and Service Note.* http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/quadblade.html.
- Cisco Systems, Inc. ESD Training Program. http://www.cisco.com/web/learning/le31/esd/WelcomeP.html.
- Cisco Systems, Inc. *Cisco UCS Manager B-Series Troubleshooting Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/ts/guide/UCSTroubleshooting.html.
- Cisco Systems, Inc. *UCS C-Series Servers Troubleshooting Tips.* http://www.cisco.com/en/US/products/ps10493/products_configuration_example09186a0080b10d5e.shtml.
- Cisco Systems, Inc. *Cisco UCS C-Series Rack Servers Data Sheets.* http://www.cisco.com/en/US/products/ps10493/products_data_sheets_list.html.

© 2012 Cisco Systems, Inc.
https://t.me/learningnets

- Cisco Systems, Inc. *Cisco UCS C200 Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C200M1/install/c200M1.html.
- Cisco Systems, Inc. *Cisco UCS C210 Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C210M1/install/C210M1.html.
- Cisco Systems, Inc. *Cisco UCS C220 Server Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/C220.html.
- Cisco Systems, Inc. *Cisco UCS C240 Server Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C240/install/C240.html.
- Cisco Systems, Inc. *Cisco UCS C250 Extended Memory Blade Server Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C250M1/install/c250M1.html.
- Cisco Systems, Inc. *Cisco UCS C460 Server Installation and Service Guide.* http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C460/install/C460.html.
- Cisco Systems, Inc. *UCS C-Series Servers Troubleshooting Tips.* http://www.cisco.com/en/US/products/ps10493/products_configuration_example09186a0080b10d5e.shtml.

https://t.me/learningnets

© 2012 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 two items are valid Cisco UCS B-Series blade server models? (Choose two.) (Source: Describing the Cisco UCS B-Series Product Family)

A) Cisco UCS B210 M2
B) Cisco UCS B260 M2
C) Cisco UCS B250 M3
D) Cisco UCS B250 M2
E) Cisco UCS B440 M2

Q2) What is the maximum number of fixed SFP ports on any Cisco UCS 6100 or 6200 Series fabric interconnect? (Source: Describing the Cisco UCS B-Series Product Family)

A) 64
B) 48
C) 40
D) 32
E) 20

Q3) Which three components are available in the second-generation of Cisco UCS hardware? (Source: Describing the Cisco UCS B-Series Product Family)

Q4) How many fixed interfaces are available on the Cisco UCS 6248UP fabric interconnect? (Source: Cisco UCS B-Series Product Family)

A) 48
B) 24
C) 32
D) 16

Q5) How many hot-swappable fan modules does the Cisco UCS 5108 Blade Server Chassis include? (Source: Describing the Cisco UCS B-Series Product Family)

A) 4
B) 6
C) 8
D) 12

Q6) Which two items are valid Cisco UCS C-Series rack server models? (Choose two.) (Source: Describing the Cisco UCS C-Series Product Family)

A) Cisco UCS C230 M2
B) Cisco UCS C260 M2
C) Cisco UCS C250 M3
D) Cisco UCS C250 M2
E) Cisco UCS C440 M2

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Q7) What is the maximum amount of RAM that can be installed in a Cisco UCS C210 M2 server? (Source: Describing the Cisco UCS C-Series Product Family)

A) 96 GB
B) 128 GB
C) 192 GB
D) 256 GB
E) 512 GB
F) 1 TB

Q8) If you were building a server to host 500 virtual desktops, which Cisco UCS C-Series rack server provides the best performance to memory- and CPU-bound applications? (Source: Describing the Cisco UCS C-Series Product Family)

Q9) How many hard disk slots for local storage are included in the Cisco UCS C210 M2 rack server? (Source: Cisco UCS C-Series Product Family)

A) 12
B) 8
C) 20
D) 16

Q10) How many hot-swappable power-supply modules can be installed in a Cisco UCS C460 M2 High Performance Rack Mount Server chassis? (Source: Describing the Cisco UCS C-Series Product Family)

A) 8
B) 4
C) 3
D) 2

Q11) Which two models are valid Cisco UCS I/O modules? (Choose two.) (Source: Connecting Cisco UCS B-Series Blade Servers)

A) Cisco UCS 2004XP
B) Cisco UCS 2108XP
C) Cisco UCS 2104XP
D) Cisco UCS 2204XP
E) Cisco UCS 2208XP

Q12) What is the maximum number of connections from a single I/O module to the fabric interconnect, using first-generation Cisco UCS hardware? (Source: Connecting Cisco UCS B-Series Blade Servers)

A) 1
B) 2
C) 3
D) 4
E) 6
F) 8

Q13) If you have a requirement for 40 Gb/s bandwidth to a Cisco UCS B-Series blade server, what is required? (Source: Connecting Cisco UCS B-Series Blade Servers)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Q14) How many 10-Gb/s SFP+ slots are included in the Cisco UCS 2208XP I/O module? (Source: Cisco UCS B-Series Connectivity)

A) 10
B) 4
C) 32
D) 8

Q15) If all the servers in a Cisco UCS 5108 Blade Server Chassis have Cisco UCS Virtual Interface Card 1280 mezzanine cards installed, what is the best possible bandwidth oversubscription ratio with all the interfaces between a Cisco UCS 2208XP I/O module and Cisco UCS 6348UP fabric interconnects? (Source: Connecting Cisco UCS B-Series Blade Servers)

A) 16:1
B) 8:1
C) 4:1
D) 2:1
E) 1:1

Q16) Which two choices are configuration methods during initial Cisco UCS fabric interconnect configuration? (Choose two.) (Source: Setting up an Initial Cisco UCS B-Series Cluster)

A) setup
B) config
C) console
D) restore
E) serial
F) GUI
G) Management 0

Q17) What is the maximum number of peers in a Cisco UCS cluster? (Source: Setting up an Initial Cisco UCS B-Series Cluster)

A) 1
B) 2
C) 4
D) 6
E) 8

Q18) What must be done before two Cisco UCS fabric interconnects can form a peer relationship? (Source: Configuring Cisco UCS B-Series Initial Cluster Setup)

Q19) The _____ and _____ must be configured correctly to enroll with a CA. (Source: Setting up an Initial Cisco UCS B-Series Cluster)

Q20) Which command shows the cluster link state? (Source: Setting up an Initial Cisco UCS B-Series Cluster)

A) **show cluster state**
B) **show cluster state links**
C) **show cluster state extended**
D) **show cluster extended-state**

https://t.me/learningnets

Q21) Which two choices are main configuration elements in the Cisco UCS Manager GUI? (Choose two.) (Source: Describing Cisco UCS Manager Operations)

A) breadcrumb trail
B) configuration pane
C) navigation pane
D) action pane
E) content pane
F) management pane

Q22) In which format is configuration data stored in the Cisco UCS Manager database? (Source: Describing Cisco UCS Manager Operations)

A) SQL
B) Sybase
C) Oracle
D) HTML
E) XML

Q23) Where in the Cisco UCS manager GUI would you configure VLANs? (Source: Describing Cisco UCS Manager Operations)

Q24) Physical components subject to FSM validation include _____ and _____. (Source: Describing Cisco UCS Manager Operations)

Q25) On which tab can you find a list of FSM state transitions in chronological order? (Source: Describing Cisco UCS Manager Operations)

A) FSM tab
B) Admin tab
C) Event tab
D) Equipment tab

Q26) Which two selections represent identity resources? (Choose two.) (Source: Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles)

A) server name
B) slot number
C) MAC address
D) vNIC
E) vHBA
F) UUID

Q27) What is the purpose of a UUID? (Source: Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles)

A) SAN adapter identity
B) SAN port identity
C) Ethernet identity
D) server serial number
E) BIOS identity

Q28) Where in the Cisco UCS Manager GUI can you configure a UUID? (Source: Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles)

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Q29) MAC addresses are _____ bits and pWWNs are _____ bits. (Source: Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles)

Q30) What would you use to create service profiles that use consistent policy and unique hardware identities? (Source: Describing Cisco UCS Manager Pools, Policies, Templates, and Service Profiles)

A) pools
B) Service Profile wizard
C) Servers tab
D) service profile template

https://t.me/learningnets

Module Self-Check Answer Key

Q1) D, E

Q2) C

Q3) Cisco UCS 5648UP Fabric Interconnect, Cisco UCS 2208XP I/O module, and Cisco UCS Virtual Interface Card 1280

Q4) C

Q5) C

Q6) B, D

Q7) C

Q8) Cisco UCS C460 M2 High-Performance Rack Server

Q9) D

Q10) B

Q11) C, E

Q12) D

Q13) Cisco UCS 6248UP Fabric Interconnect, Cisco UCS 2208XP I/O module, and Cisco UCS Virtual Interface Card 1280 mezzanine card

Q14) D

Q15) C

Q16) A, D

Q17) B

Q18) Cross-connect links L1 to L1 and L2 to L2

Q19) hostname and domain name

Q20) D

Q21) C, E

Q22) E

Q23) on the LAN tab in the navigation pane

Q24) chassis and I/O module

Q25) C

Q26) C, F

Q27) D

Q28) on the Servers tab in the navigation pane

Q29) 48 and 64

Q30) D

© 2012 Cisco Systems, Inc.

https://t.me/learningnets