DCICT

Introducing Cisco Data Center Technologies

Volume 1

Version 1.0

Student Guide

Text Part Number: 97-3181-01

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

Students, this letter describes important course evaluation access information!

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco Systems is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. We would appreciate a few minutes of your time to complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web.

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,

Cisco Systems Learning

https://t.me/learningnets

https://t.me/learningnets

Table of Contents

Volume 1

Course Introduction ***1***

Overview 1
Learner Skills and Knowledge 2
Course Goal and Objectives 3
Course Flow 4
Additional References 5
Cisco Glossary of Terms 6
Cisco Online Education Resources 7
Cisco NetPro Forums 8
Cisco Learning Network 9
Introductions 10

Cisco Data Center Network Services ***1-1***

Overview 1-1
Module Objectives 1-1

Examining Functional Layers of the Cisco Data Center **1-3**

Overview 1-3
Objectives 1-3
Traditional Isolated LAN and SAN Networks 1-4
Cisco Data Center Infrastructure 1-11
Cisco Data Center Infrastructure Topology Layout 1-13
LAN Core, Aggregation, and Access Layers 1-14
Core and Access Layers in a LAN Collapsed-Core Design 1-18
Example: Collapsed Core in Traditional Network 1-18
Example: Collapsed Core in Routed Network 1-19
Core and Edge Layers in a Data Center SAN Design 1-20
Collapsed-Core SAN Design 1-23
Summary 1-25

Reviewing the Cisco Nexus Product Family **1-27**

Overview 1-27
Objectives 1-27
Cisco Nexus Data Center Product Portfolio 1-28
Cisco Nexus 7000 Series Chassis Options 1-31
Cisco Nexus 7000 Series 9-Slot Switch Chassis 1-32
Cisco Nexus 7000 Series 10-Slot Switch Chassis 1-33
Cisco Nexus 7000 Series 18-Slot Switch Chassis 1-34
Cisco Nexus 7000 Series Supervisor Module 1-37
Supervisor CMP 1-38
Cisco Nexus 7000 Series Licensing Options 1-39
Cisco Nexus 7000 Series Fabric Modules 1-42
Cisco Nexus 7000 Series I/O Modules 1-44
M1-Series 8-Port 10 Gigabit Ethernet Module 1-44
M1-Series 32-Port 10 Gigabit Ethernet Module 1-44
M1-Series 48-Port Gigabit Ethernet Module 1-45
F1-Series 32-Port 1 and 10 Gigabit Ethernet Module 1-46
F2-Series 48-Port 1 and 10 Gigabit Ethernet Module 1-46
Cisco Nexus 7000 Series Power Supply Options 1-48
Cisco Nexus 7000 6.0-kW AC Power Supply Module 1-48
Cisco Nexus 7000 7.5-kW AC Power Supply Module 1-49
Cisco Nexus 7000 6.0-kW DC Power Supply Module 1-49
Cisco Nexus 5000 Series Chassis Options 1-54
Cisco Nexus 5010 and 5020 Switches Features 1-55
Cisco Nexus 5010 and 5020 Expansion Modules 1-59
Cisco Nexus 5500 Platform Switches Features 1-61
Cisco Nexus 5500 Platform Switches Expansion Modules 1-65

https://t.me/learningnets

Cisco Nexus 5000 Series Software Licensing 1-70
Cisco Nexus 2000 Series Fabric Extenders Function in the Cisco Data Center 1-73
Cisco Nexus 2000 Series Fabric Extenders Features 1-80
Summary 1-84

Reviewing the Cisco MDS Product Family 1-87

Overview 1-87
Objectives 1-87
Cisco MDS 9000 Series Product Suite 1-88
Cisco MDS 9500 Series Chassis Options 1-90
Cisco MDS 9506 Chassis 1-92
Cisco MDS 9509 Chassis 1-92
Cisco MDS 9513 Chassis 1-92
Cisco MDS 9500 Series Supervisor Modules 1-94
Cisco MDS 9000 Series Licensing Options 1-98
Common Software Across All Platforms 1-99
Cisco MDS 9000 Series Switching Modules 1-101
Cisco MDS 9000 4-Port 10-Gbps Fibre Channel Switching Module 1-101
Cisco MDS 9000 Family 8-Gbps Fibre Channel Switching Modules 1-101
Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules 1-102
Advanced Fibre Channel Switching Module Features 1-102
Cisco MDS 9000 10-Gbps 8-Port FCoE Module 1-103
Cisco MDS 9500 Series Power Supply Options 1-105
Cisco MDS 9100 Series Switches 1-107
Cisco MDS 9124 Multilayer Fabric Switch 1-107
Cisco MDS 9148 Multilayer Fabric Switch 1-107
Cisco MDS 9222i Switch 1-109
Summary 1-110

Monitoring the Cisco Nexus 7000 and 5000 Series Switches 1-111

Overview 1-111
Objectives 1-111
Connecting to the Console Port 1-112
Running the Initial Setup Script 1-113
Connecting to the Cisco Nexus 7000 CMP 1-116
Connecting to the Switch Using SSH to Connect to the Management VRF 1-120
Reviewing the ISSU on the Cisco Nexus Switches 1-123
Verifying VLANs 1-129
Examining the Operational Plane 1-131
Unified Port Controller 1-133
Reviewing Cisco NX-OS Default Control Plane Policing 1-135
Default Policing Policies 1-136
Using Important CLI Commands 1-137
Summary 1-152

Describing vPCs and Cisco FabricPath in the Data Center 1-153

Overview 1-153
Objectives 1-153
Virtual Port Channels 1-154
Verifying vPCs 1-165
Cisco FabricPath 1-168
Verifying Cisco FabricPath 1-176
Summary 1-178

Using OTV on Cisco Nexus 7000 Series Switches 1-179

Overview 1-179
Objectives 1-179
OTV on the Cisco Nexus 7000 Series Switches 1-180
Verifying OTV on the Cisco Nexus 7000 Series Switches 1-193
Summary 1-197
Module Summary 1-199
References 1-200

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check 1-203
Module Self-Check Answer Key 1-207

Cisco Data Center Virtualization 2-1

Overview 2-1
Module Objectives 2-1

Virtualizing Network Devices 2-3

Overview 2-3
Objectives 2-3
Describing VDCs on the Cisco Nexus 7000 Series Switch 2-4
Verifying VDCs on the Cisco Nexus 7000 Series Switch 2-15
Navigating Between VDCs on the Cisco Nexus 7000 Series Switch 2-19
Describing NIV on Cisco Nexus 7000 and 5000 Series Switches 2-20
Summary 2-36

Virtualizing Storage 2-37

Overview 2-37
Objectives 2-37
LUN Storage Virtualization 2-38
Storage System Virtualization 2-39
Host-Based Virtualization 2-43
Array-Based Virtualization 2-43
Network-Based Virtualization 2-43
Storage Virtualization 2-44
Summary 2-47

Virtualizing Server Solutions 2-49

Overview 2-49
Objectives 2-49
Benefits of Server Virtualization 2-50
VM Partitioning 2-56
VM Isolation 2-56
VM Encapsulation 2-56
VM Hardware Abstraction 2-57
Available Data Center Server Virtualization Solutions 2-58
Overview 2-69
Virtual Partitions 2-69
Cluster Shared Volumes 2-70
Live Migration Feature 2-71
Summary 2-72

Using the Cisco Nexus 1000V Series Switch 2-73

Overview 2-73
Objectives 2-73
Limitations of VMware vSwitch 2-74
Advantages of VMware vDS 2-84
How the Cisco Nexus 1000V Series Switch Brings Network Visibility to the VM Level 2-90
How the VSM and VEM Integrate with VMware ESX or ESXi and vCenter 2-100
Summary 2-110

Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch 2-111

Overview 2-111
Objectives 2-111
Verifying the Initial Configuration and Module Status on the Cisco Nexus 1000V Series Switch 2-112
Verifying the VEM Status on the ESX or ESXi Host 2-118
Validating VM Port Groups 2-123
Summary 2-128
Module Summary 2-129
References 2-130
Module Self-Check 2-133
Module Self-Check Answer Key 2-136

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Course Introduction

Overview

Introducing Cisco Data Center Technologies (DCICT) v1.0 is designed to provide students with foundational knowledge and a broad overview of Cisco data center products and their operation.

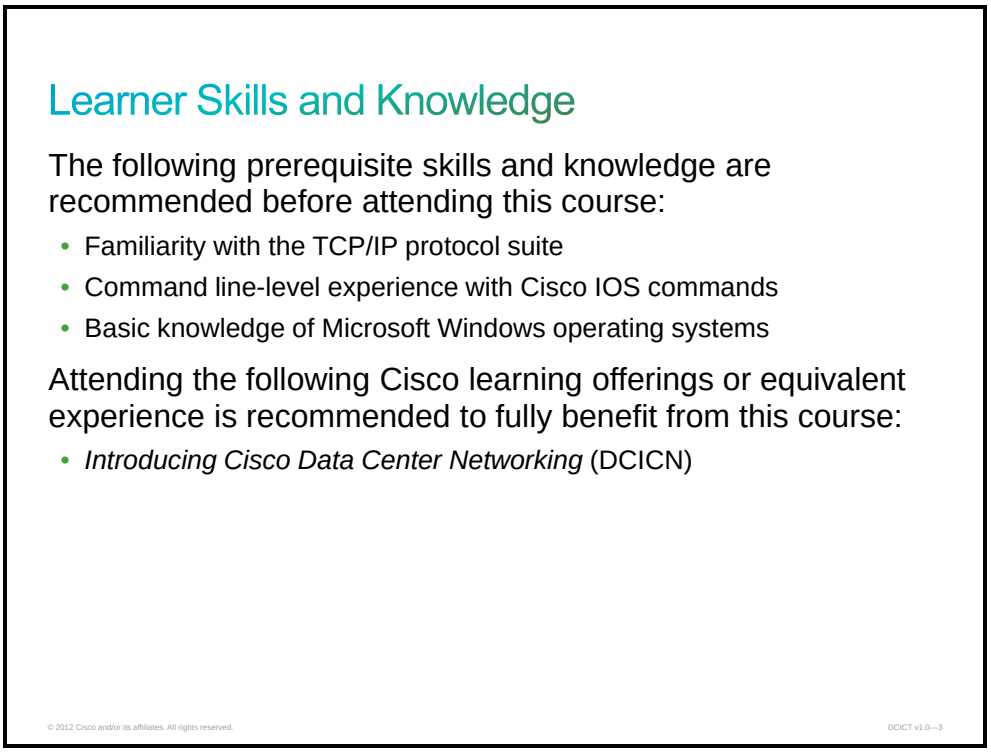
The course covers the architecture, components, connectivity, and features of a Cisco data center network.

The student will gain practical experience configuring the initial setup of Cisco Nexus 7000 9-Slot Switch, Cisco Nexus 5548UP Switch, Cisco Unified Computing System (UCS) 6120XP 20-Port Fabric Interconnect, and Cisco MDS 9124 Multilayer Fabric Switch. Students will also learn to verify the proper operation of a variety of features such as Overlay Transport Virtualization (OTV), Cisco FabricPath, port channels, virtual port channels (vPC), and Cisco Nexus 1000v Distributed Virtual Switch for VMware ESX.

https://t.me/learningnets

Learner Skills and Knowledge

This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.

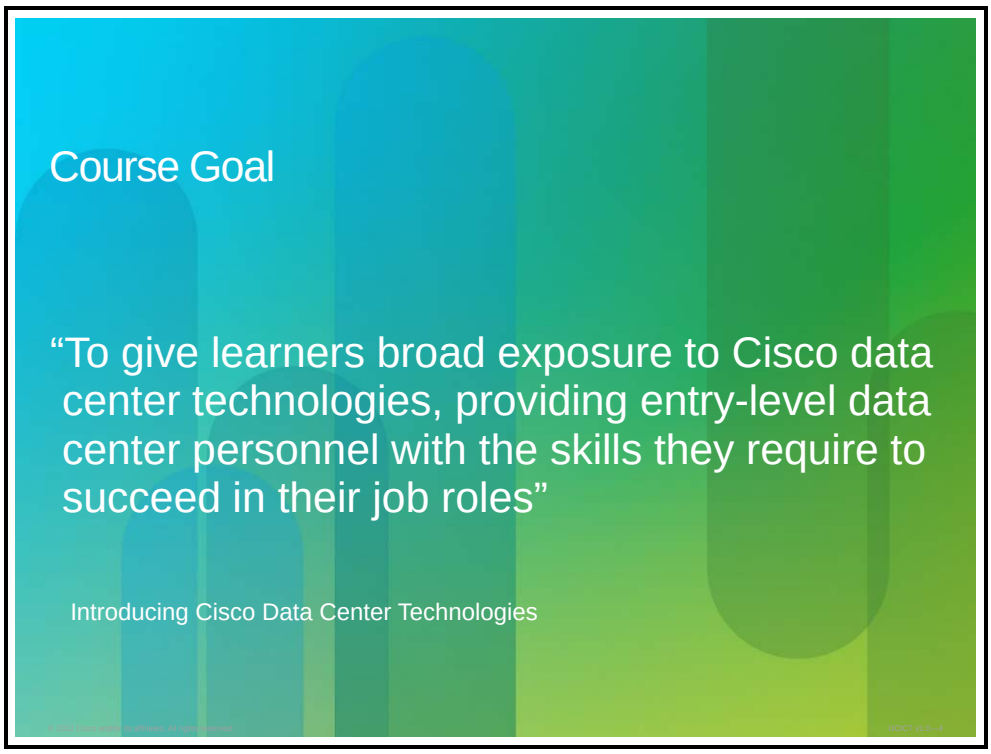


 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Course Goal and Objectives

This topic describes the course goal and objectives.



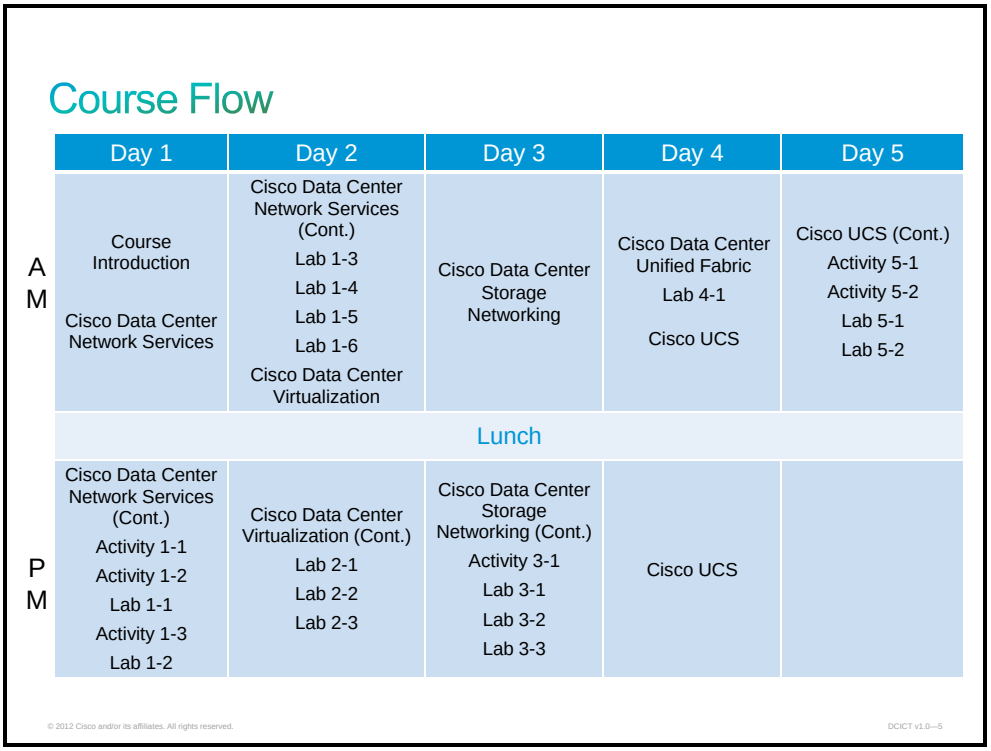
Upon completing this course, you will be able to meet these objectives:

- Describe and verify Cisco data center fundamentals
- Describe Cisco data center virtualization
- Describe Cisco data center storage networking
- Describe Cisco data center unified fabric
- Describe and verify Cisco UCS

https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.



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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

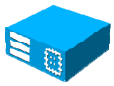
Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.

Cisco Icons and Symbols

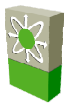

Cisco UCS 6100 Series Fabric Interconnect


Cisco UCS 5100 Series Blade Server Chassis


Cisco UCS C-Series Rack Server


Cisco Nexus 5000 Series Switch


Cisco Nexus 7000 Series Switch


Cisco MDS 9500 Series Multilayer Director


Cisco MDS 9200 Series Multilayer Switch


Cisco MDS 9100 Series Multilayer Fabric Switch

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

Cisco Icons and Symbols (Cont.)

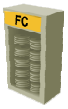

Cisco Nexus 2000 Fabric Extender


Cisco Nexus 1000V Virtual Ethernet Module (VEM)


Workstation


Application Server


Fibre Channel JBOD



Fibre Channel RAID Subsystem



Fibre Channel Tape Subsystem

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

https://t.me/learningnets

Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://docwiki.cisco.com/wiki/Internetworking_Terms_and_Acronyms_%28ITA%29_Guide.

© 2012 Cisco Systems, Inc.

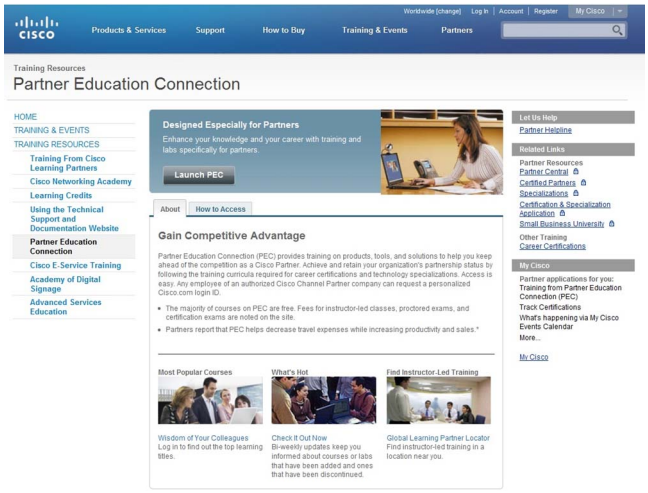
https://t.me/learningnets

Cisco Online Education Resources

This topic presents Cisco online training resources that complement this course.

Cisco Partner Education Connection

http://www.cisco.com/go/pec



Cisco Partner Education Connection (PEC) provides training on products, tools, and solutions to help you keep ahead of the competition as a Cisco Partner. Achieve and advance your partnership status for your organization by following the training curriculum that is required for career certifications and technology specializations. Access is easy. Any employee of an authorized Cisco Channel Partner company can request a personalized Cisco.com login ID.

- Most courses on PEC are free. Fees for instructor-led classes, proctored exams, and certification exams are noted on the site.
- Partners report that PEC helps decrease travel expenses while increasing productivity and sales.

https://t.me/learningnets

Cisco NetPro Forums

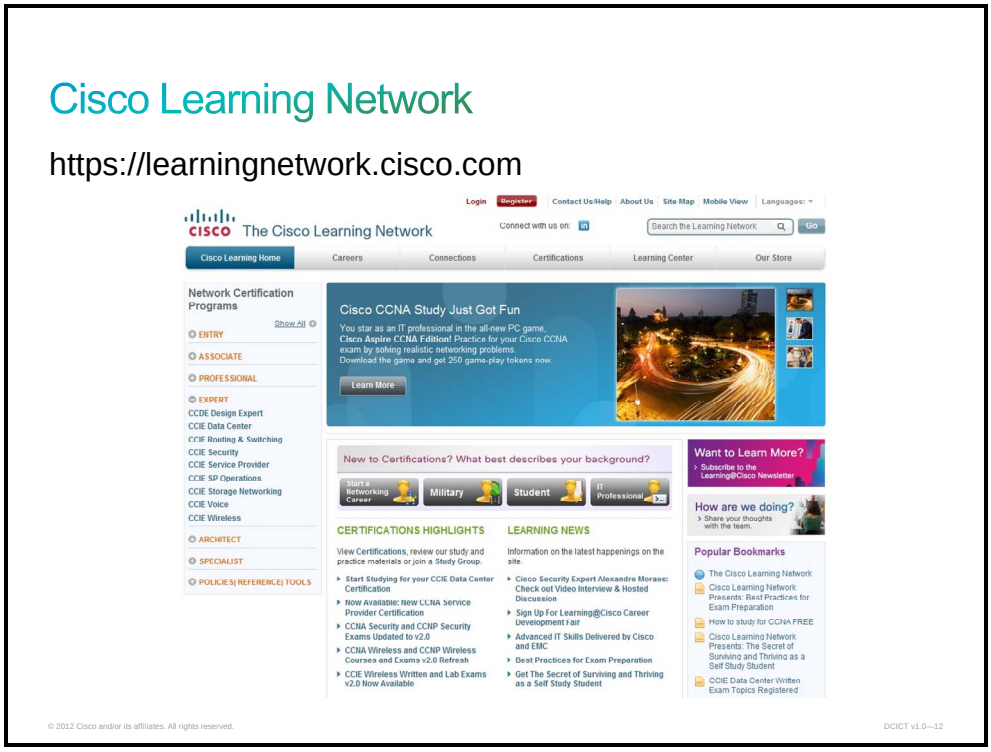


Cisco NetPro forums are part of the online Cisco Support Community. NetPro forums are designed to share configurations, issues, and solutions among a community of experts. The forums are conveniently arranged into distinct subject matter expertise categories to make finding or supplying solutions a simple process.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Learning Network

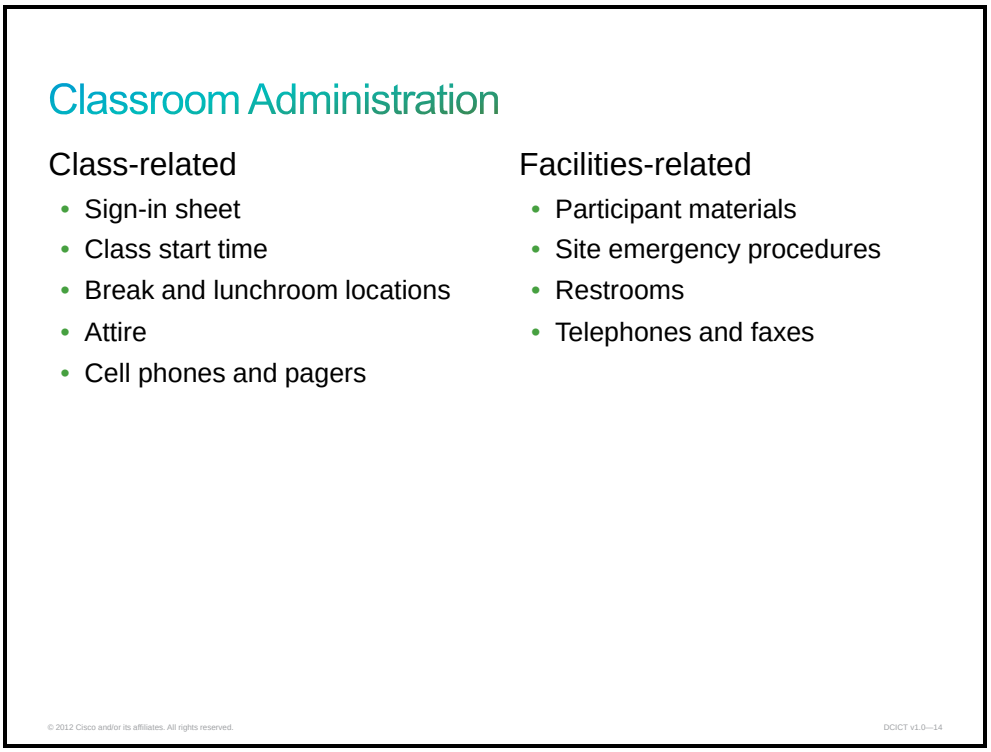


The Cisco Learning Network is a one-stop repository where certification seekers can find the latest information on certification requirements and study resources, and discuss certification with others. Whether you are working toward certification at the Associate, Professional or Expert level, the Cisco Learning Network is always available to assist with reaching your certification goals.

https://t.me/learningnets

Introductions

This topic presents the general administration of the course, and an opportunity for student introductions.



The instructor will tell you about specific site requirements and the location of restrooms, break rooms, and emergency procedures.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Introductions

- Your name
- Your company
- Prerequisite skills
- Brief history
- Objective



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

DCICT v1.0—15

Be prepared to introduce yourself to the class and discuss your experience, environment, and specific learning goals for the course.

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module 1

Cisco Data Center Network Services

Overview

This module describes Cisco data center fundamentals.

Module Objectives

Upon completing this module, you will be able to describe the features of the data center switches for network and SAN connectivity and their relationship to the layered design model. This ability includes being able to meet these objectives:

- Explain the functions of topology layers in Cisco data center LAN and SAN networks
- Describe the features of Cisco Nexus 7000 and 5000 Series Switches, and the Cisco Nexus 2000 Series Fabric Extenders and their relationship to the layered design model
- Describe the features of the Cisco MDS Fibre Channel switches and their relationship to the layered design model
- Perform and initial configuration and validate common features of the Cisco Nexus 7000 and 5000 Series Switches
- Describe vPCs and Cisco FabricPath and how to verify their operation on Cisco Nexus 7000 and 5000 Series Switches
- Describe OTV as a method of data center interconnect

https://t.me/learningnets

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

Lesson 1

Examining Functional Layers of the Cisco Data Center

Overview

The functional layers that are designed to serve the needs of the data center are typically built around three layers—the core layer, the aggregation layer, and the access layer. This topic looks at these layers and explains their use in the data center model. This lesson describes the functions of the topology layers in Cisco data center LAN and SAN networks.

Objectives

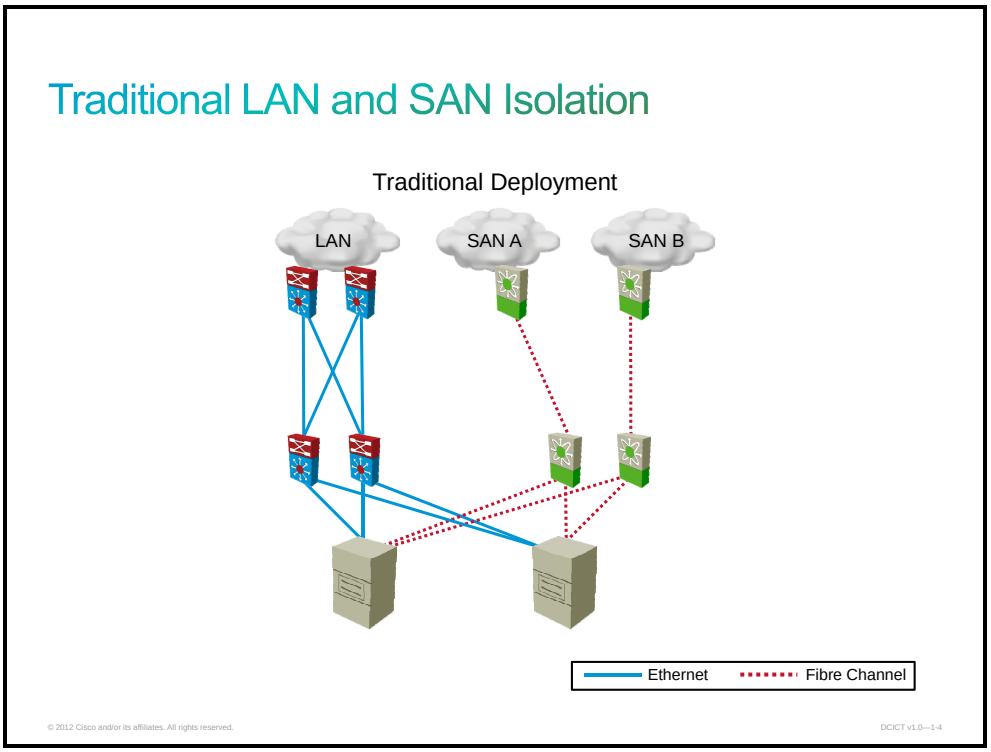
Upon completing this lesson, you will be able to describe the functional layers of the Cisco data center LAN and SAN networks. This ability includes being able to meet these objectives:

- Describe the reasons for segregating LAN and SAN traffic
- Describe the core, aggregation, and access layers in the data center and their functions
- Describe the reasons for combining the core and aggregation layers in a LAN design
- Describe the core and edge layers in a data center SAN design
- Describe the reasons for combining the core and edge layers in a SAN design

https://t.me/learningnets

Traditional Isolated LAN and SAN Networks

This topic describes the reasons for segregating LAN and SAN traffic and the topology of the traditional LAN and SAN infrastructures.



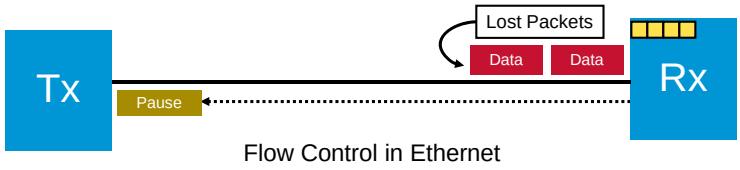
Traditionally, LAN and SAN infrastructures have been separated for various reasons. These issues are some of the reasons:

- **Security:** Separation helps ensure that data that is stored on the SAN appliances cannot be corrupted through normal TCP/IP hacking methodologies.
- **Bandwidth:** Initially, higher bandwidth was more prevalent on SAN infrastructures than on LAN infrastructures.
- **Flow Control:** SAN infrastructures use a buffer-to-buffer flow control that ensures that data is delivered in order and without any loss in transit unlike the TCP/IP flow control methodology.
- **Performance:** SAN infrastructures traditionally provide higher performance.

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

Ethernet Flow Control

- Flow control is how data interchange is controlled in a network.
- The flow control strategy that is used by Ethernet and other data networks can degrade performance.
  - Transmitter does not stop transmitting packets until after the receiver buffers overflow.
  - Lost packets must be retransmitted.
  - Degradation can be severe under heavy traffic loads.



Flow Control in Ethernet

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

Flow control is a mechanism that is used to ensure that frames are sent only when there is somewhere for them to go. Just as traffic lights are used to control the flow of traffic in cities, flow control manages the data flow in a LAN or SAN environment.

Some data networks, such as Ethernet, use a flow control strategy that can result in degraded performance:

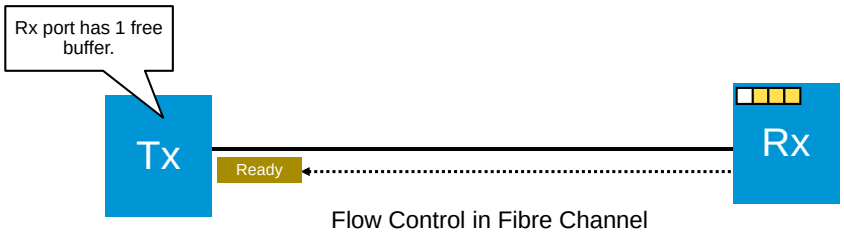
- A Tx port can begin sending data packets at any time.
- When the Rx port buffers are completely filled and cannot accept any more packets, the Rx port tells the Tx port to stop or slow the flow of data.
- After the Rx port has processed some data and has some buffers available to accept more packets, it tells the Tx port to resume sending data.

This strategy results in lost packets when the Rx port is overloaded, because the Rx port tells the Tx port to stop sending data after it has already overflowed. All lost packets must be retransmitted. The retransmissions degrade performance. Performance degradation can become severe under heavy traffic loads.

https://t.me/learningnets

SAN Flow Control

- Fibre Channel uses a credit-based strategy:
  - The transmitter does not send a frame until the receiver tells the transmitter that the receiver can accept another frame.
  - The receiver is always in control.
- Benefits:
  - Prevents loss of frames that are caused by buffer overruns
  - Maximizes performance under high loads



Flow Control in Fibre Channel

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

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

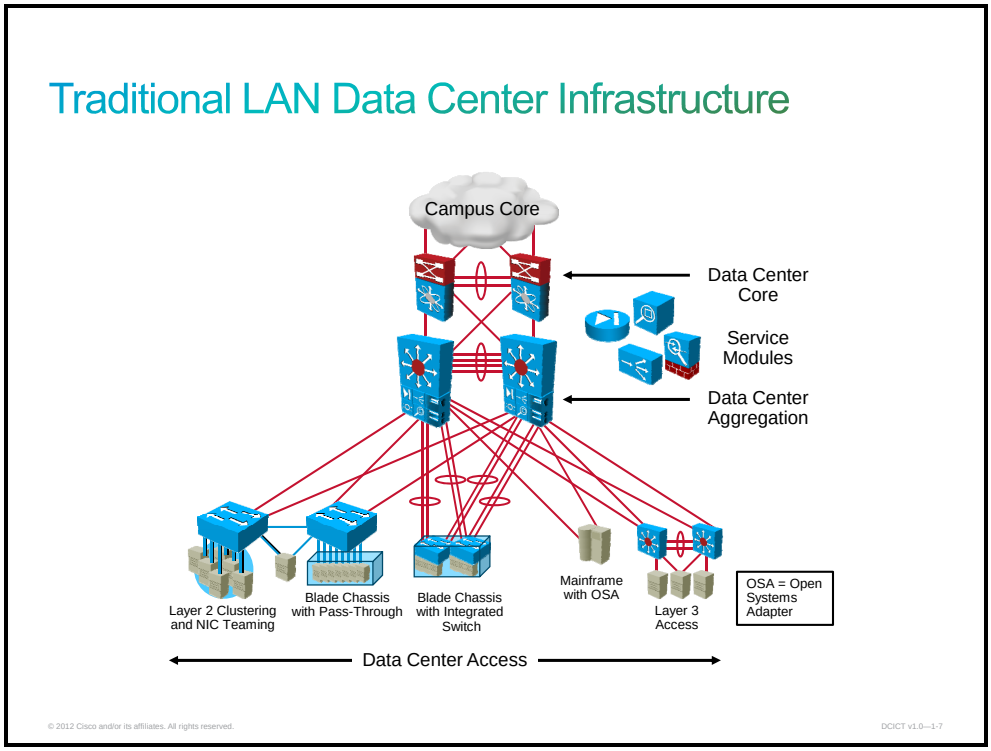
A credit-based strategy ensures that the Rx port is always in control. The Rx port must issue a credit for each frame that is sent by the transmitter. This strategy prevents frames from being lost when the Rx port runs out of free buffers. Preventing lost frames maximizes performance under high-traffic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

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



The figure shows a typical large enterprise data center design. The design follows the Cisco multilayer infrastructure architecture, including core, aggregation (or distribution), and access layers.

The data center infrastructure must provide the following important features and services:

- Port density
- Layer 2 and Layer 3 connectivity for servers at the access layer
- Security services that are provided by access control lists (ACLs), firewalls, and intrusion prevention systems (IPSs) at the data center aggregation layer
- Server farm services such as content switching, caching, and Secure Sockets Layer (SSL) offloading
- Multitier server farms, mainframes, and mainframe services (such as Telnet 3270 [TN3270], load balancing, and SSL offloading)

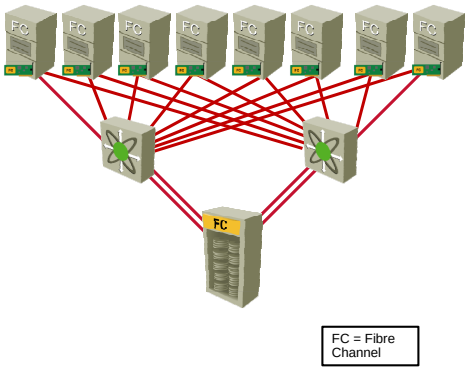
Network devices are often deployed in redundant pairs to avoid a single point of failure.

https://t.me/learningnets

Traditional SAN Data Center Infrastructure

Single Tier Design

- Inexpensive:
  - Small switches
  - Low port count
- Manageable:
  - Small number of switches
- Simple Design:
  - Single hop
  - Low latency
- Limited scalability



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

A single tier design has one layer of switches connecting servers to storage. It is relatively inexpensive and easy to manage because there are usually a few switches.

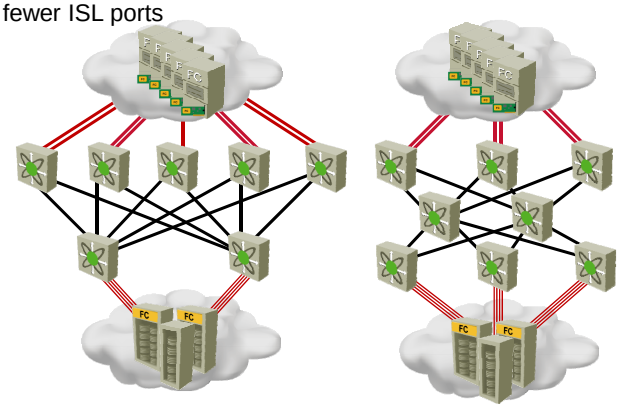
© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Traditional SAN Data Center Infrastructure (Cont.)

Multitier Design

- Scalable:
  - Just add more edge switches when required
- Efficient and well-structured design:
  - More user ports, fewer ISL ports
- Resilient design:
  - Multiple paths



FC = Fibre Channel

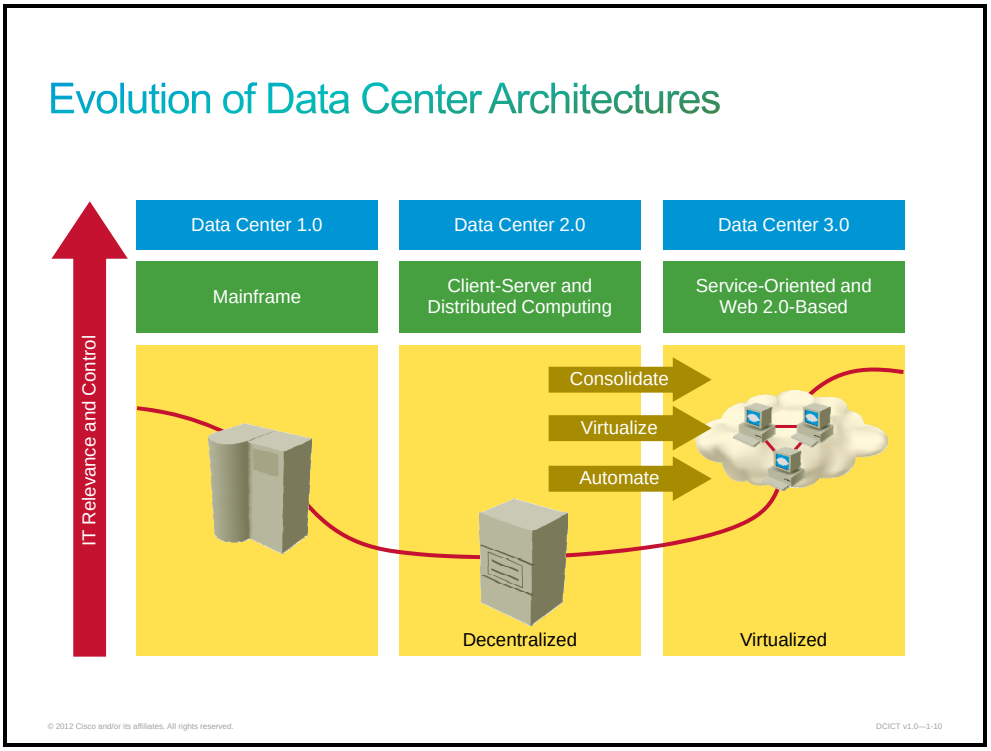
© 2012 Cisco and/or its affiliates. All rights reserved.

In the multitier design that is shown in the figure, the following connections occur:

- Storage targets connect to one set of switches and are often locally attached to the core switches.
- Hosts, often remote, connect into the fabric through a separate layer of switches called edge switches.

The multitier design supports multiple layers of switches and thus can scale significantly, although many ports are used for interswitch links (ISLs). The major advantage of this design is its simplicity, which makes it relatively easy to visualize and troubleshoot. The multitier design is good where storage is centralized, except where edge switches are located at satellite locations. Like core-edge fabrics, this design has a high inherent oversubscription because of the hierarchical design. This oversubscription must be offset by adding ISLs, which decreases the effective port count.

https://t.me/learningnets

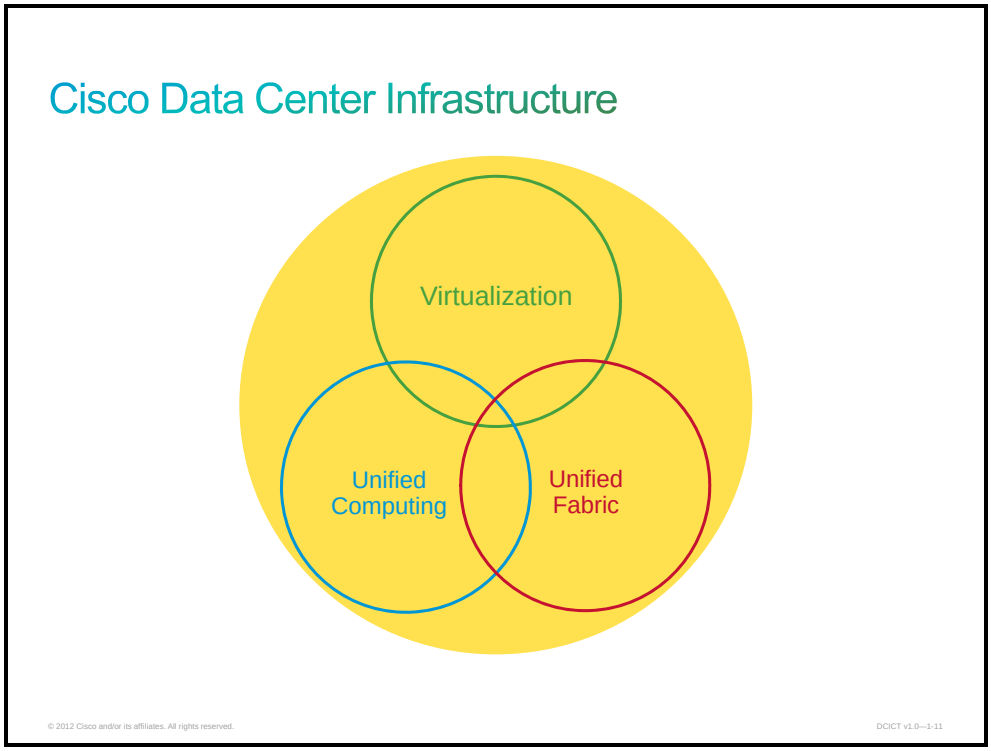


This figure shows how data centers have changed in the last two decades. At first, data centers were monolithic and centralized, employing mainframes and terminals, which the users used to perform their work on the mainframe. The mainframes are still used in the finance sector because they are an advantageous solution in terms of availability, resilience, and service level agreements (SLAs).

The second era emphasized client-server and distributed computing, with applications being designed in such a way that the user used client software to access an application. Also, services were distributed because of poor computing ability and high link cost. Mainframes were too expensive to serve as an alternate solution.

Currently, with communication infrastructure being relatively cheaper, and with an increase in computing capacities, data centers are being consolidated. Consolidation is occurring because the distributed approach is expensive in the long term. The new solution uses equipment virtualization, resulting in a much higher utilization of servers than in the distributed approach. The new solution also brings a significantly higher return on investment (ROI) and lower total cost of ownership (TCO).

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



Cisco Data Center Infrastructure

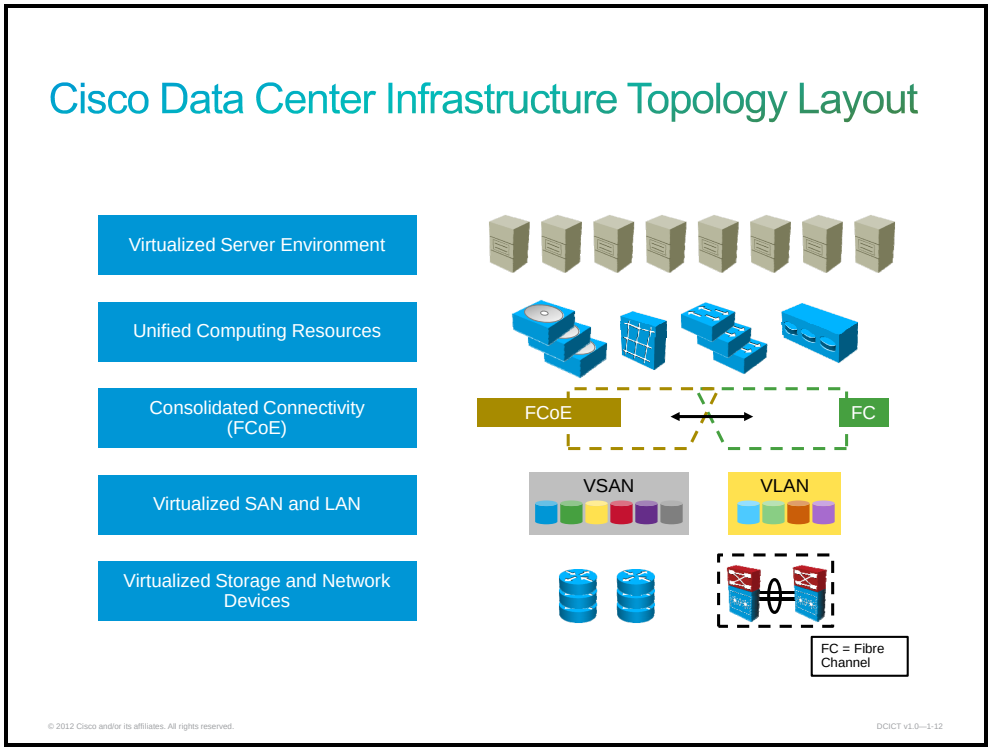
The Cisco data center infrastructure is made up of a comprehensive portfolio of virtualization technologies and services that bring network, computing and storage, and virtualization platforms closer together. These technologies and services provide unparalleled flexibility, visibility, and policy enforcement within virtualized data centers. The following components are the three main components of the data center architectures:

- **Virtualization**
  - Cisco Virtual Network Link (Cisco VN-Link) technologies, including the Cisco Nexus 1000V Distributed Virtual Switch for VMware ESX or ESXi, deliver consistent per-virtual-machine visibility and policy control for SAN, LAN, and unified fabric.
  - Virtual SAN, virtual device contexts (VDCs), and unified fabric help multiple virtual networks converge to simplify and reduce data center infrastructure and TCO.
  - Flexible networking options support all server form factors and vendors, including options for integrated Ethernet and Fibre Channel switches for Dell, IBM, and HP blade servers. These options provide a consistent set of services across the data center to reduce operational complexity.
  - The Cisco Unified Computing System (Cisco UCS) solution unifies network, compute, and virtualization resources into a single system that delivers end-to-end optimization for virtualized environments, while retaining the ability to support traditional operating system and application stacks in physical environments.

https://t.me/learningnets

- **Unified Fabric**
  - There are two primary approaches to deploying a unified data center fabric: Fibre Channel over Ethernet (FCoE) and Internet Small Computer Systems Interface (iSCSI). Both are supported on the unified fabric, which provides a reliable 10 Gigabit Ethernet foundation.
  - Unified fabric lossless operation also improves the performance of iSCSI that is supported by both Cisco Catalyst and Cisco Nexus switches. In addition, the Cisco MDS series of storage switches has hardware and software features that are specifically designed to support iSCSI.
  - The Cisco Nexus Series of switches was designed to support unified fabric. The Cisco Nexus 7000 and Cisco Nexus 5000 Series Switches support Data Center Bridging (DCB) and FCoE, with support for FCoE on the Cisco MDS Series of switches as well.
  - Special host adapters, called converged network adapters (CNAs), are required to support FCoE. Hardware adapters are available from vendors such as Emulex and QLogic, while a software stack is available for certain Intel 10 Gigabit Ethernet network interfaces.
  - FCoE is supported on VMware ESX or ESXi vSphere and higher.
- **Unified computing:** The Cisco UCS platform is a next-generation data center platform that does the following:
  - Unites computing, networking, storage access, and virtualization into a cohesive system
  - Integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers
  - Increases IT staff productivity and business agility through rapid provisioning and mobility support
  - Provides a standards-based unified network fabric that is supported by a partner ecosystem of industry leaders
  - Offers Cisco VN-Link virtualization support
  - Offers Cisco Unified Computing System (UCS) Extended Memory Technology

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



Cisco Data Center Infrastructure Topology Layout

Starting with the top layer of components, virtual machines (VMs) are one of the key components of the Cisco data center infrastructure. VMs are entities that are running an application within the client operating system, which is further virtualized and running on common hardware.

The logical server personality is defined using management software. It defines the properties of the server: amount of memory, percent of total computing power, number of network interface cards (NICs), boot image, and so on.

The network hardware serves as consolidated connectivity, for example, the Cisco Nexus 5000 Series Switches. FCoE is one of the key technologies that provides for fabric unification.

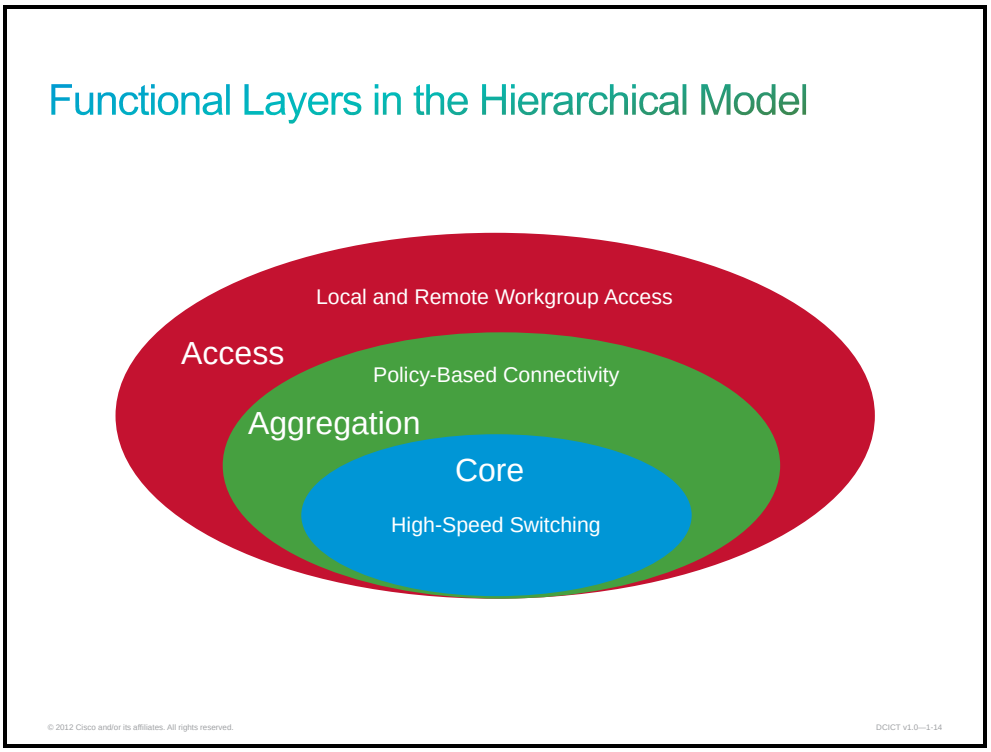
VLANs and virtual storage area networks (VSANs) provide for virtualized LAN and SAN connectivity, separating physical networks and equipment into virtual entities.

On the lowest layers, there is virtualized hardware. Storage devices can be virtualized into storage pools, and network devices are virtualized using device contexts.

https://t.me/learningnets

LAN Core, Aggregation, and Access Layers

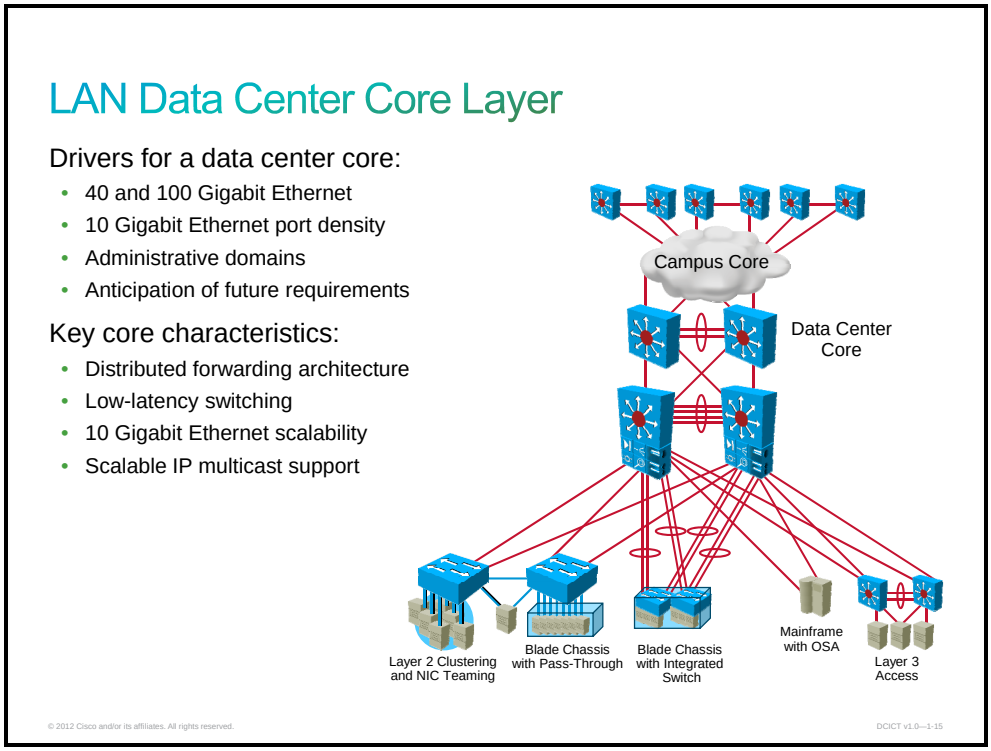
When you design a data center, the network is broken down into functional layers. There are three functional layers. These layers can be physical or logical. This topic describes the functions of the core, aggregation, and access layers in the data center.



The hierarchical network model provides a modular framework that allows flexibility in network design and facilitates ease of implementation and troubleshooting. The hierarchical model divides networks or their modular blocks into three layers: the access, aggregation (or distribution), and core layers. These layers consist of the following features:

- **Access layer:** A layer that is used to grant user access to network devices. In a network campus, the access layer generally incorporates switched LAN devices with ports that provide connectivity to workstations, IP phones, and servers. In the WAN environment, the access layer for teleworkers or remote sites may provide access to the corporate network across WAN technology. In the data center, the access layer provides connectivity for servers.
- **Aggregation (or distribution) layer:** A layer that aggregates the wiring closets, using switches to segment workgroups and isolate network problems in a data center environment. Similarly, in the campus environment, the aggregation layer aggregates WAN connections at the edge of the campus and provides policy-based connectivity.
- **Core layer (also referred to as the backbone):** A high-speed backbone that is designed to switch packets as fast as possible. Because the core layer is critical for connectivity, it must provide a high level of availability and must adapt to changes very quickly. The core layer also provides scalability and fast convergence.

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



Implementing a data center core is a best practice for large data centers. When the core is implemented in an initial data center design, it helps ease network expansion and avoid disruption to the data center environment. The following drivers are used to determine if a core solution is appropriate:

- **40 and 100 Gigabit Ethernet density:** Are there requirements for higher bandwidth connectivity such as 40 or 100 Gigabit Ethernet? With the introduction of the Cisco Nexus 7000 M-2 Series of modules, the Cisco Nexus 7000 Series Switches can now support much higher bandwidth densities.
- **10 Gigabit Ethernet density:** Without a data center core, will there be enough 10 Gigabit Ethernet ports on the campus core switch pair to support both the campus distribution and the data center aggregation modules?
- **Administrative domains and policies:** Separate cores help to isolate campus distribution layers from data center aggregation layers for troubleshooting, administration, and implementation of policies (such as quality of service [QoS], ACLs, troubleshooting, and maintenance).
- **Anticipation of future development:** The impact that may result from implementing a separate data center core layer at a later date may make it worthwhile to install the core layer at the beginning.

The data center typically connects to the campus core using Layer 3 links. The data center network is summarized, and the core injects a default route into the data center network.

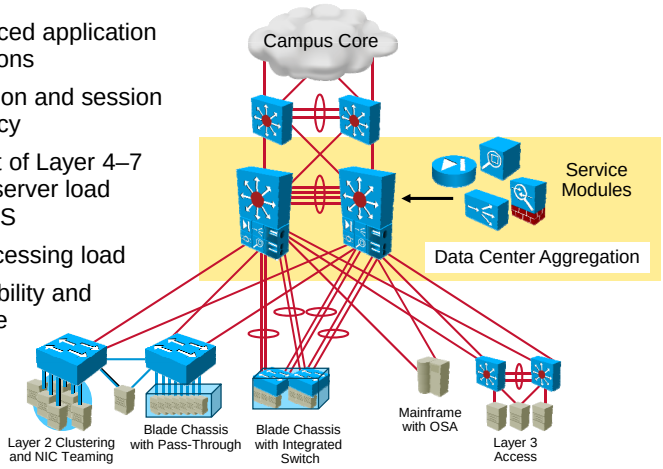
Key core characteristics include the following:

- Distributed forwarding architecture
- Low-latency switching
- 10 Gigabit Ethernet scalability
- Scalable IP multicast support

https://t.me/learningnets

LAN Data Center Aggregation Layer

- Aggregates traffic to data center core
- Aggregates advanced application and security functions
- Maintains connection and session state for redundancy
- Allows deployment of Layer 4–7 services: firewall, server load balancing, SSL, IPS
- Supports STP processing load
- Provides high flexibility and economies of scale



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

The aggregation (or distribution) layer aggregates the uplinks from the access layer to the data center core. This layer is the critical point for control and application services. Security and application service devices (such as load-balancing devices, SSL offloading devices, firewalls, and IPS devices) are often deployed as a module in the aggregation layer. This design lowers TCO and reduces complexity by reducing the number of components that you need to configure and manage.

Note Service devices that are deployed at the aggregation layer are shared among all the servers. Service devices that are deployed at the access layer provide benefit only to the servers that are directly attached to the specific access switch.

The aggregation layer typically provides Layer 3 connectivity from the data center to the core. Depending on the requirements, the boundary between Layer 2 and Layer 3 at the aggregation layer can be in the multilayer switches, the firewalls, or the content switching devices. Depending on the data center applications, the aggregation layer may also need to support a large Spanning Tree Protocol (STP) processing load.

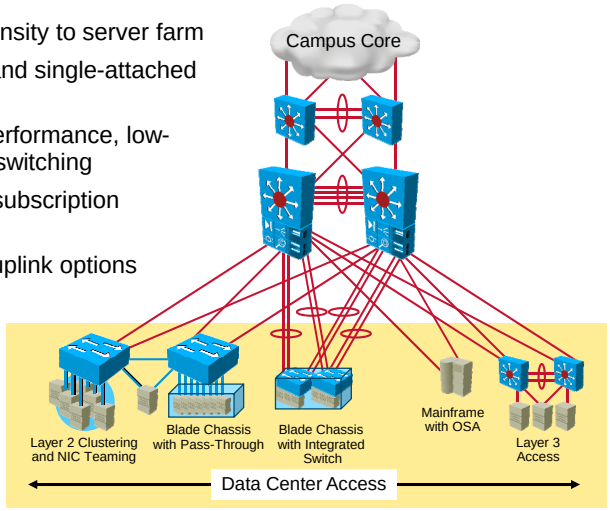
Note OSA is an Open Systems Adapter. The OSA is a network controller that can be installed in a Mainframe I/O cage. It integrates several hardware features and supports many networking transport protocols. The OSA card is the communications device for the mainframe architecture.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

LAN Data Center Access Layer

- Can support Layer 2 or Layer 3 access
- Provides port density to server farm
- Supports dual- and single-attached servers
- Provides high-performance, low-latency Layer 2 switching
- Has mix of oversubscription requirements
- Provides many uplink options



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

The data center access layer provides Layer 2, Layer 3, and mainframe connectivity. The design of the access layer varies, depending on whether you use Layer 2 or Layer 3 access. The access layer in the data center is typically built at Layer 2, which allows better sharing of service devices across multiple servers. This design also allows the use of Layer 2 clustering, which requires the servers to be Layer 2-adjacent. With Layer 2 access, the default gateway for the servers can be configured at the access layer or the aggregation layer.

With a dual-homing network interface card (NIC), you need a VLAN or trunk between the two access switches. The VLAN or trunk supports the single IP address on the two server links to two separate switches. The default gateway would be implemented at the access layer as well.

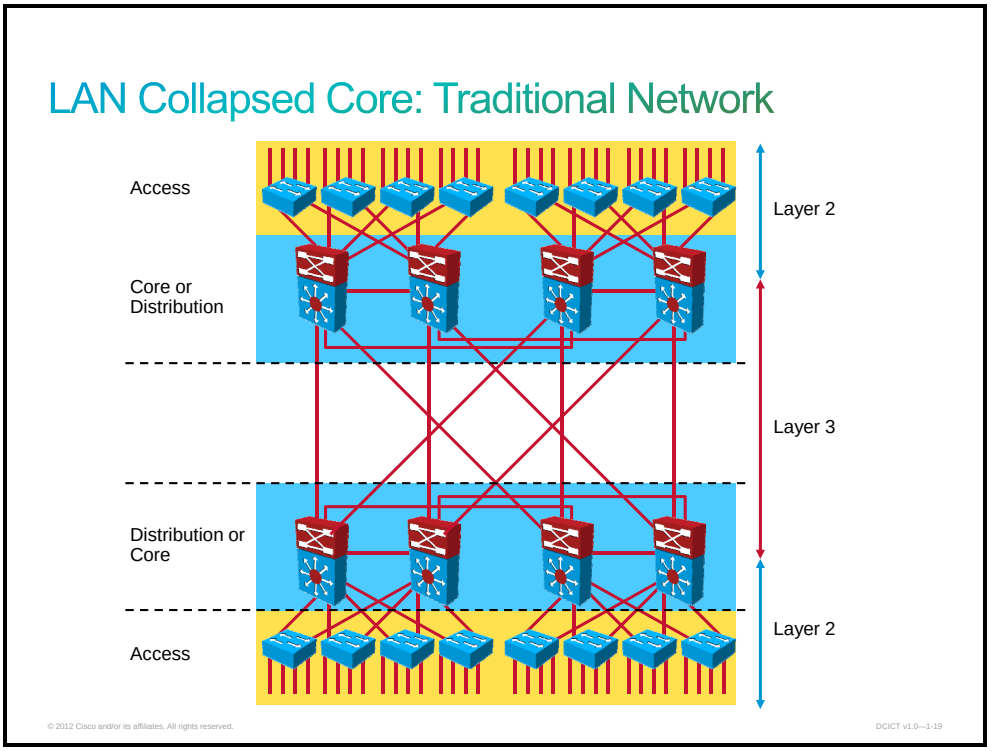
Although Layer 2 at the aggregation layer is tolerated for traditional designs, new designs try to confine Layer 2 to the access layer. With Layer 2 at the aggregation layer, there are physical loops in the topology that must be managed by the STP. Rapid Per VLAN Spanning Tree Plus (Rapid PVST+) is a recommended best practice to ensure a logically loop-free topology over the physical topology.

A mix of both Layer 2 and Layer 3 access models permits a flexible solution and allows application environments to be optimally positioned.

https://t.me/learningnets

Core and Access Layers in a LAN Collapsed-Core Design

Although there are three functional layers in a network design, these layers do not have to be physical layers, but can be logical layers. This topic describes the reasons for combining the core and aggregation layers.



Example: Collapsed Core in Traditional Network

The figure shows Layer 2 switching with a collapsed core. A typical packet flow between Layer 2 access switches would follow these steps:

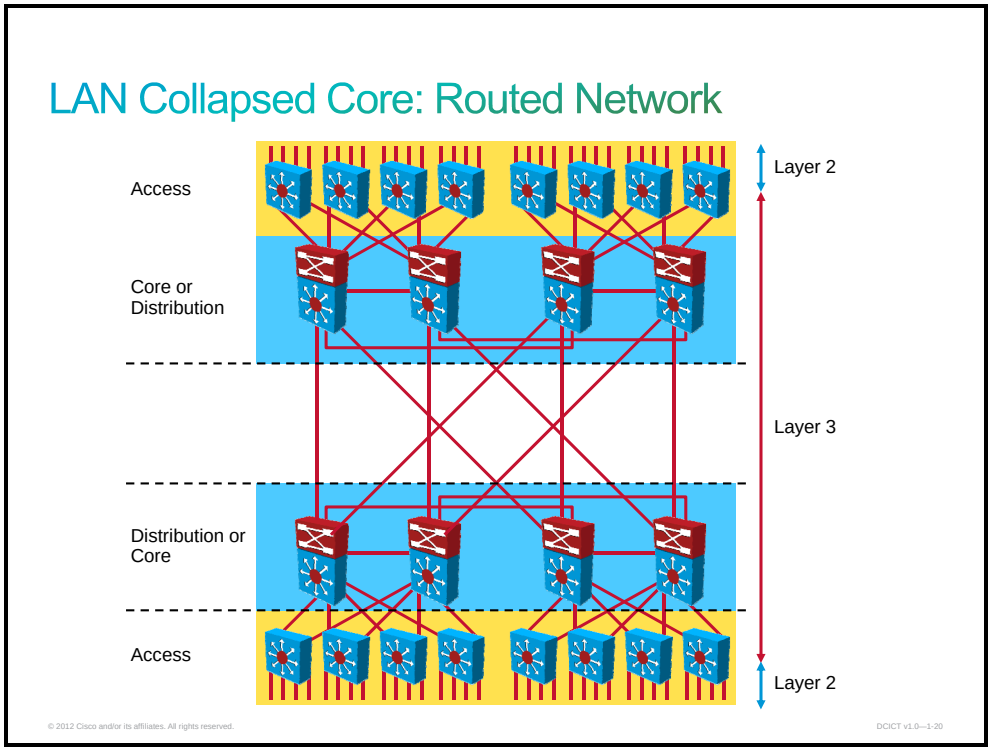
Step 1 A packet is Layer 2-switched toward the distribution switch.

Step 2 The aggregation (or distribution) switch performs Layer 3 switching.

Step 3 The receiving aggregation (or distribution) switch performs Layer 3 switching toward an access LAN.

Step 4 The packet is Layer 2-switched across the access LAN to the destination host.

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



The three-tier architecture that was previously defined is traditionally used for a larger enterprise campus. For the campus of a small- or medium-sized business, the three-tier architecture may be excessive. An alternative would be to use a two-tier architecture, in which the distribution and the core layer are combined. With this option, the small- or medium-sized business can reduce costs, because both distribution and core functions are performed by the same switch. In the collapsed core the switches would provide direct connections for the access-layer switches, server farm, and edge modules.

However, one disadvantage with two-tier architecture is scalability. As the small campus begins to scale, a three-tier architecture should be considered.

Example: Collapsed Core in Routed Network

The figure shows an example of Layer 3 switching with a collapsed core. A typical packet flow between Layer 3 access switches would follow these steps:

Step 1 A packet is Layer 3-switched toward the aggregation (or distribution)/core switch.

Step 2 The aggregation (or distribution) or core switch (or both) performs Layer 3 switching.

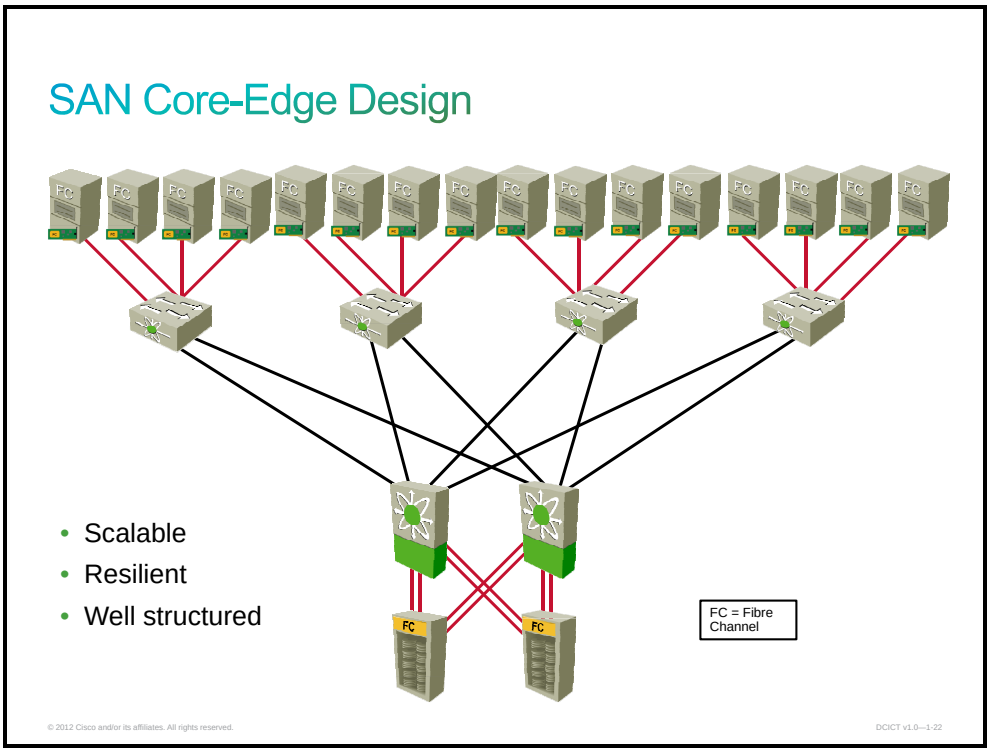
Step 3 The receiving aggregation (or distribution) or core switch (or both) performs Layer 3 switching toward an access LAN.

Step 4 The packet is Layer 3-switched across the access LAN to the destination host.

https://t.me/learningnets

Core and Edge Layers in a Data Center SAN Design

The SAN design is also broken down into functional layers, which are the core and edge Layers. This topic discusses the core and edge layers of the SAN design in a data center.



The core-edge fabric topology has all the necessary features of a SAN in terms of resiliency and performance. It provides resiliency and predictable recovery, scalable performance, and scalable port density in a well-structured design. The architecture of the core-edge model has two or more core switches in the center of the fabric. These core switches interconnect two or more edge switches in the periphery of the fabric. The core itself can be a complex core consisting of a ring or mesh of directors or a simple pair of separate core switches.

There are many choices in terms of which switch models can be used for the two different layers. For example:

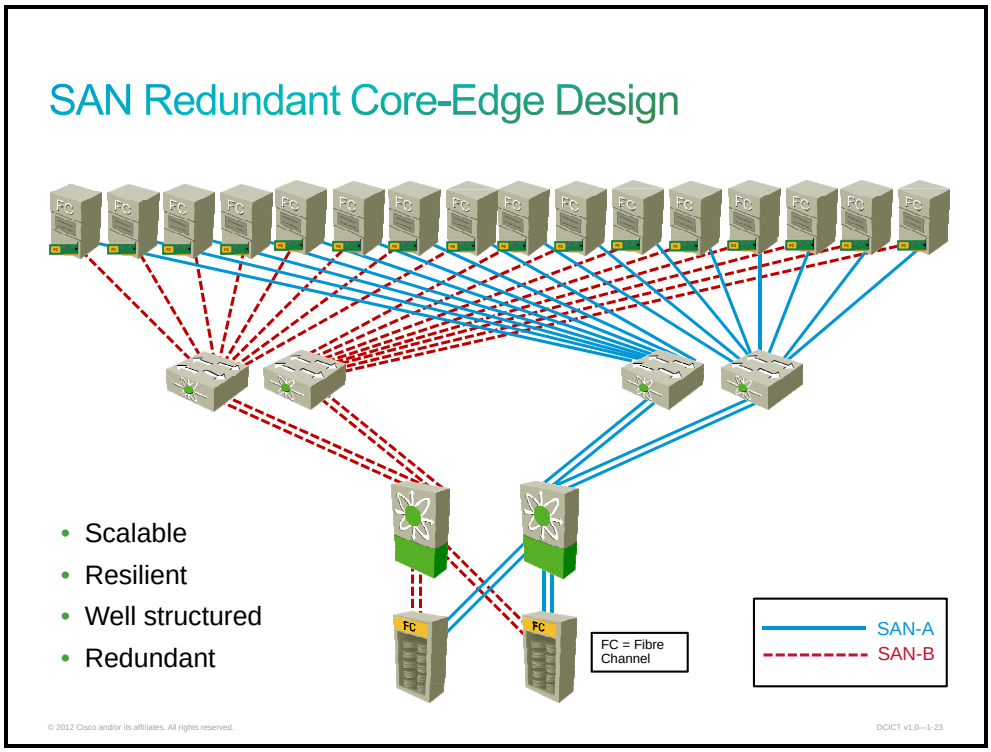
- **Core switches:** Director-class switches, like the Cisco MDS 9500 Series switches, are recommended at the core of the topology because of the excessive high availability requirement.
- **Edge switches:** You can use either use smaller fabric switches or more director-class switches depending on the overall required fabric port density.

You have to make another choice about where to place the storage devices. In smaller core-edge port density solutions, the storage devices are commonly connected to the core. This connection eliminates one hop in terms of the path from hosts to storage devices and also eliminates one point of potential ISL congestion to access storage resources. However, with larger port density fabrics that have more switches at the edge layer, storage devices themselves might need to be deployed in the edge layer. It is important that if storage devices are deployed across a common set of edge switches, these edge switches must be connected with a higher ISL aggregate bandwidth. The reason the switches must be connected this way is that this

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

connection represents a point of traffic consolidation and a potential point of detrimental congestion.



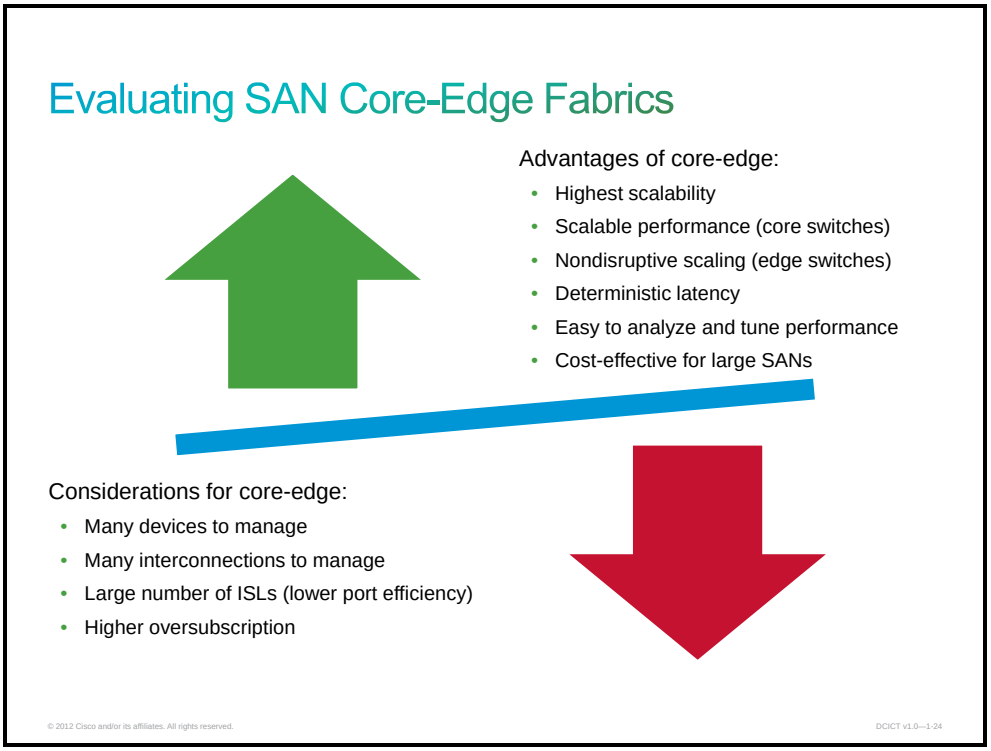
The core-edge fabric topology has all the necessary features of a SAN in terms of resiliency and performance, but it is still a single fabric. That is, all switches are connected together through ISLs. A redundant core-edge design must employ two completely isolated fabrics, which are shown here with red and blue links. Each server now has two separate Fibre Channel ports, one connecting to each fabric, and the storage arrays are connected to both fabrics.

Most Fibre Channel switches can only support a single fabric, so redundant SAN designs require separate physical switches. Cisco MDS 9000 Series Multilayer Switches support VSANs or virtual fabrics, that allow you to configure separate logical fabrics on a smaller number of physical switches.

The redundant SAN design uses multipath software on every server that identifies two separate paths to the primary and secondary storage. The primary path might be active while the secondary path is passive. If a failure occurs in the primary path, the multipath software switches over to the secondary path. Most multipath software can concurrently use both the primary path and secondary path and perform load balancing across both active paths.

This topology provides resiliency and predictable recovery, scalable performance, and scalable port density in a well-structured SAN.

https://t.me/learningnets



The advantages of the core-edge fabric design are as follows:

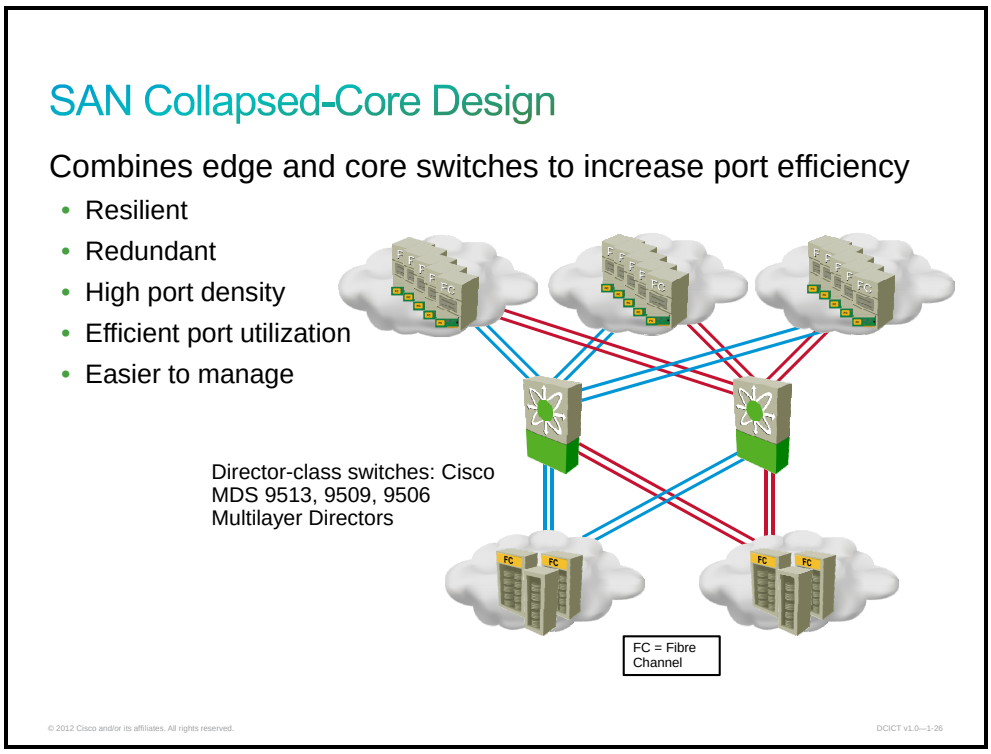
- **Highest scalability:** You can easily add more edge switches.
- **Scalable port counts:** If more ports are needed, you can add more edge switches without disrupting ongoing SAN traffic.
- **Scalable performance:** If more performance is needed, you can add more core switches or core ISLs.
- **FSPF:** When there are at least two core switches, with each edge switch connected to at least two core switches, there are two equal-cost paths from any edge switch to any other edge switch. Therefore, Fabric Shortest Path First (FSPF) can use both data paths, effectively doubling the ISL bandwidth.
- **Deterministic latency:** With a single redundant pair of core switches, the fabric maintains a two-hop count from any device to any other device even if one ISL goes down.
- **Easy to analyze and tune performance:** The simple, symmetric design simplifies throughput calculations.
- **Relatively cost-effective**: Core edge fabrics use fewer ports for ISLs than in other designs and provide similar levels of availability and performance. Smaller switches (or even hubs) can be used at the edge, and larger switches or high-performance director-class switches can be used at the core.

The core-edge design has only one notable consideration for a large SAN; it involves many switches and interconnections. While the symmetrical nature of the core-edge design simplifies performance analysis and tuning, there are still many switches to manage.

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

Collapsed-Core SAN Design

This topic describes the reasons for combining the core and edge layers.



The collapsed-core topology is a topology with the features of the core-edge topology but delivers required port densities in a more efficient manner. The collapsed-core topology is enabled by high port density that is offered by the Cisco MDS 9500 Series of director switches.

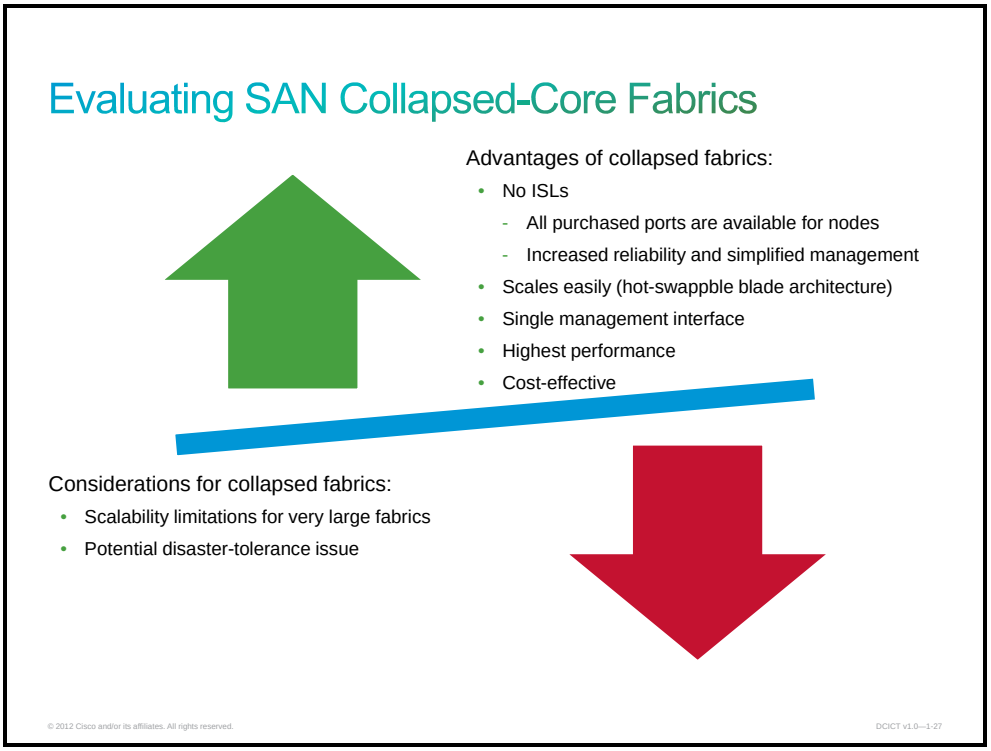
The salient features of the collapsed-core topology are as follows:

- The resiliency of this topology is sufficiently high because of the redundant structure; however, it does not have the excessive ISLs that a mesh topology may have.
- The performance of this topology is predictable, because paths between two communicating devices do not vary in length and hops; the direct path between two switches is always chosen. Using ISL port channeling between the switches provides for faster recovery time with no fabric-wide disruption. The port channels between the switches provide load balancing over a logical ISL link.
- The port density of this topology can scale quite high, but not as high as the core-edge topology.
- The topology is simple to design, implement, maintain, and troubleshoot.

The main advantage of this topology is the degree of scalability that is offered at a very efficient effective port usage. The collapsed-core design aims to offer very high port density while eliminating a separate physical layer of switches and their associated ISLs.

The only disadvantage of the collapsed-core topology is its scalability limit relative to the core-edge topology. While the collapsed-core topology can scale quite large, the core-edge topology should be used for the largest of fabrics. However, to continually scale the collapsed-core design, one could convert the core to a core-edge design and add another layer of switches.

https://t.me/learningnets



The advantages of the collapsed architecture are as follows:

- This architecture has the most efficient use of ports because no ports are consumed for ISLs.
- The absence of ISLs significantly increases reliability and manageability.
- Ports can be scaled easily by adding hot-swappable blades, without disrupting traffic.
- Highest performance is achieved by director-class switches because the high-speed backplane provides low fixed latency between any two ports.
- Single management interface for all ports simplifies performance analysis and tuning. Fewer switches to manage.

The considerations for the collapsed architecture are as follows:

- A collapsed-core fabric runs into scalability limitations at very large port counts because of the increasing number of ISLs that are required. However, with large director switches like the Cisco MDS 9513 Multilayer Director switch, a collapsed-core fabric can easily scale to meet the demands of all but the largest SAN fabrics.
- Customers might not want to locate the entire fabric core in a single location. In a core-edge design, the core switches can be located on separate floors or even at separate facilities for increased disaster tolerance.

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

Summary

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

Summary

- Traditionally, LAN and SAN traffic has been segregated, making use of its own infrastructures. This method has been primarily for security and protection of data.
- In any design, the network is broken down into functional layers, making it easier to design, build, manage, and troubleshoot the infrastructure.
- Although there are three layers within the design, sometimes the core and aggregation layers are collapsed into a single physical layer, but still retain the two functional layers.
- The SAN design is also broken down into layers, primarily a core and edge layer.
- For efficiency of port utilization, the core and edge layer can be collapsed into fewer physical switches, while still retaining the logical separation of the layers themselves.

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

https://t.me/learningnets

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

Lesson 2

Reviewing the Cisco Nexus Product Family

Overview

The Cisco Nexus 7000, 5000, and 2000 Family of products create the network foundation for a unified fabric data center and high-end performance campus core. This lesson describes the features of the Cisco Nexus Family and its relationship to the layer design model

Objectives

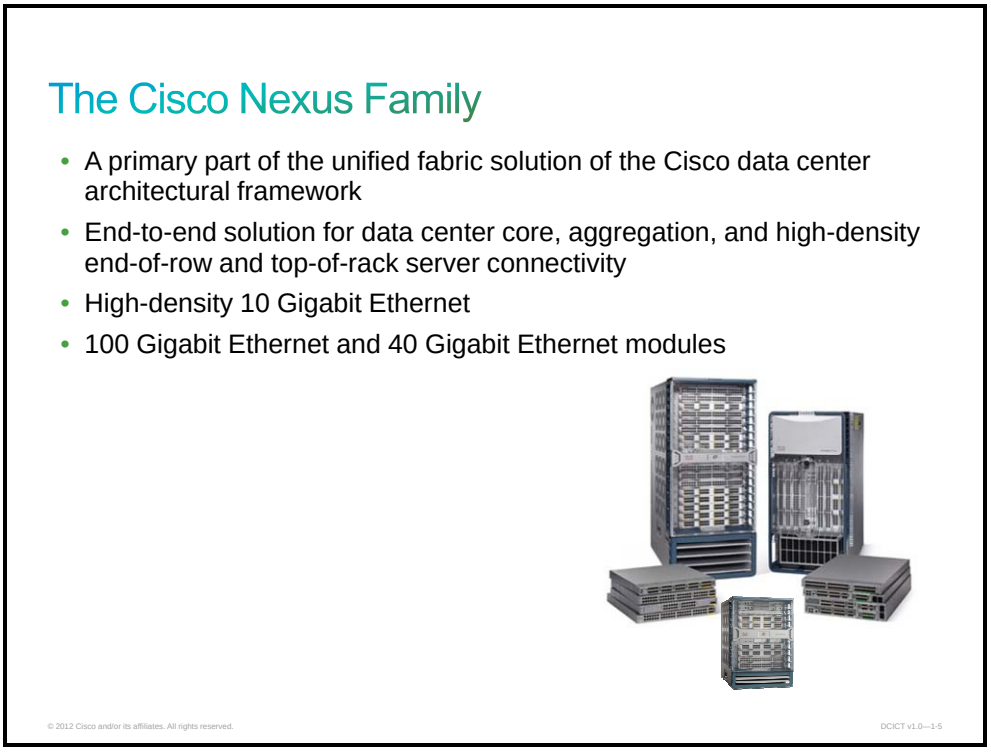
Upon completing this lesson, you will be able to describe the features of Cisco Nexus 7000 Series, 5000 Series, and 2000 Series Fabric Extenders products and their relationship to the layered design model. You will be able to meet these objectives:

- Describe data center transformation with Cisco Nexus products
- Describe the capacities of the Cisco Nexus 7000 Series chassis
- Describe the capabilities of the Cisco Nexus 7000 Series Supervisor Module
- Describe the licensing options for the Cisco Nexus 7000 Series Switches
- Describe the capacities of the Cisco Nexus 7000 Series Fabric Modules
- Describe the capabilities of the Cisco Nexus 7000 Series I/O modules
- Describe the capabilities of the Cisco Nexus 7000 Series power supply modules
- Compare the chassis of the Cisco Nexus 5000 Series Switches
- Describe the Cisco Nexus 5010 and 5020 Switches
- Describe the expansion modules that are available for Cisco Nexus 5010 and 5020 Switches
- Describe the Cisco Nexus 5548P, 5548UP, and 5596UP Switches
- Describe the expansion modules that are available for Cisco Nexus 5548P, 5548UP, and 5596UP Switches
- Describe the software licensing for the Cisco Nexus 5000 Series Switches
- Describe how the Cisco Nexus 2000 Series Fabric Extenders act as remote line modules
- Describe the features of the Cisco Nexus 2000 Series Fabric Extenders

https://t.me/learningnets

Cisco Nexus Data Center Product Portfolio

The Cisco Nexus Family of products has been designed with the data center in mind. They have innovative features and scalability to support the data center requirements both today and in the future. This topic describes the goals of data center transformation with the Cisco Nexus Family of products.



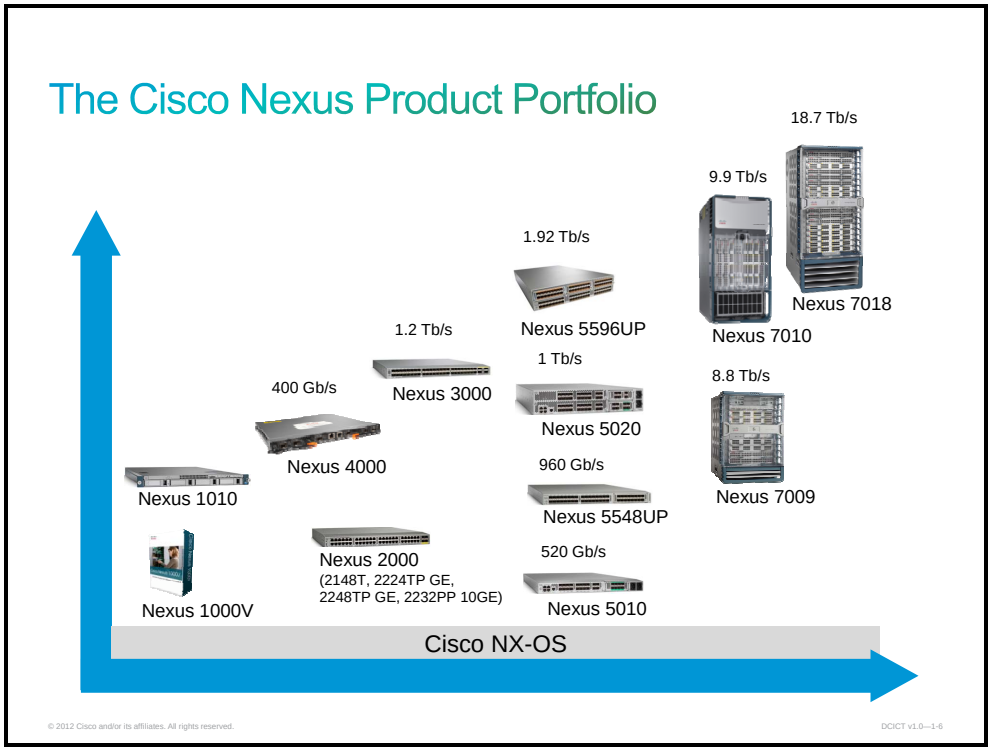
Today, data center design is mostly influenced by the following needs:

- The need for a higher level of reliability, with minimized downtime for updates and configuration changes. Once a consolidated architecture is built, it is critical to keep it operating with minimal disruption.
- The need to optimize the use of the data center network infrastructure by moving toward a topology where no link is kept idle. Traditional topologies that are based on Spanning Tree Protocol (STP) are known to be inefficient because of STP blocking links or because of active and standby network interface card (NIC) teaming. This need is addressed by Layer 2 multipathing (L2M) technologies such as virtual port channels (vPCs).
- The need to optimize computing resources by reducing the rate of growth of physical computing nodes. This need is addressed by server virtualization.
- The need to reduce the time that it takes to provision new servers. This need is addressed by the ability to configure server profiles, which can be easily applied to hardware.
- The need to reduce overall power consumption in the data center. This need can be addressed with various technologies. These technologies include unified fabric (which reduces the number of adapters on a given server), server virtualization, and more power-efficient hardware.
- The need to increase computing power at a lower cost. More and higher-performance computing clouds are being built to provide a competitive edge to various enterprises.

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

The needs that are influencing data center design call for capabilities like the following:

- Architectures capable of supporting a SAN and a LAN on the same network (for power use reduction and server consolidation).
- Architectures that provide an intrinsic lower latency than traditional LAN networks. This architecture means that a computing cloud can be built on the same LAN infrastructure as regular transactional applications.
- Architectures that can distribute Layer 2 traffic on all available links.
- Simplified cabling, which provides more efficient airflow, lower power consumption, and lower cost of deployment of high-bandwidth networks.
- Reduction of management points, which limits the impact of the sprawl of switching points (software switches in the servers, multiple blade switches, and so on).



The Cisco Nexus Family of products includes the following switches:

- **Cisco Nexus 1000V Series Switches**: A virtual machine (VM) access switch that is an intelligent software switch implementation for VMware vSphere environments running the Cisco Nexus Operating System (Cisco NX-OS) Software. The Cisco Nexus 1000V Series Switches operate inside the VMware ESX hypervisor, and support the Cisco Virtual Network Link (Cisco VN-Link) server virtualization technology to provide the following:
  - Policy-based VM connectivity
  - Mobile VM security and network policy
  - Nondisruptive operational model for server virtualization and networking teams
- **Cisco Nexus 1010 Virtual Services Appliance:** This appliance is a member of the Cisco Nexus 1000V Series Switches and hosts the Cisco Nexus 1000V Virtual Supervisor Module (VSM). It also supports the Cisco Nexus 1000V Network Analysis Module (NAM) Virtual Service Blade and provides a comprehensive solution for virtual access switching. The Cisco Nexus 1010 provides dedicated hardware for the VSM, making access switch deployment much easier for the network administrator.

https://t.me/learningnets

- **Cisco Nexus 2000 Series Fabric Extenders**: A category of data center products that are designed to simplify data center access architecture and operations. The Cisco Nexus 2000 Series Fabric Extenders use the Cisco Fabric Extender Link (Cisco FEX-link) architecture to provide a highly scalable unified server-access platform across a range of 100-Mb/s Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, unified fabric, copper and fibre connectivity, and rack and blade server environments. The Cisco Nexus 2000 Series Fabric Extenders act as remote line cards for the Cisco Nexus 5000 and 7000 Series Switches. Some of the models included in the Cisco Nexus 2000 Series Fabric Extenders include the Cisco Nexus 2148T, 2224TP GE, 2248TP GE, and 2232PP 10GE Fabric Extenders that are noted in the figure.
- **Cisco Nexus 3000 Series Switches:** The Cisco Nexus 3000 Series Switches extend the comprehensive, proven innovations of the Cisco Data Center Business Advantage architecture into the High Frequency Trading (HFT) market. The products in this range are the Cisco Nexus 3064, 3048 and 3016 Switches. The Cisco Nexus 3064 Switch supports 48 fixed 1 and 10 Gb/s enhanced small form factor pluggable (SFP+) ports and four fixed quad SFP+ (QSFP+) ports, which allow a smooth transition from 10 Gigabit Ethernet to 40 Gigabit Ethernet. The Cisco Nexus 3000 switches are well suited for financial colocation deployments, delivering features such as latency of less than a microsecond, line-rate Layers 2 and 3 unicast and multicast switching, and the support for 40 Gigabit Ethernet standards technologies.
- **Cisco Nexus 4000 Series Blade Switches:** The Cisco Nexus 4000 switch module is a blade switch solution for the IBM BladeCenter H and HT chassis. It provides the server I/O that is required for high-performance, scale-out, virtualized and nonvirtualized x86 computing architectures.
- **Cisco Nexus 5000 Series Switches including the Cisco Nexus 5000 Platform and 5500 Platform switches**: A family of line-rate, low-latency, lossless 10 Gigabit Ethernet, and Fibre Channel over Ethernet (FCoE) switches for data center applications. The Cisco Nexus 5000 Series Switches are designed for data centers that are transitioning to 10 Gigabit Ethernet as well as data centers that are ready to deploy a unified fabric that can manage LAN, SAN, and server clusters. This capability provides networking over a single link, with dual links used for redundancy. Some of the switches included in this series are the Cisco Nexus 5010, the Cisco Nexus 5020, the Cisco Nexus 5548UP and 5548P, and the Cisco Nexus 5596UP and 5596T.
- **Cisco Nexus 7000 Series Switches**: A modular data center-class switch that is designed for highly scalable 10 Gigabit Ethernet networks with a fabric architecture that scales beyond 15 terabits per second (Tb/s). The switch is designed to deliver continuous system operation and virtualized services. The Cisco Nexus 7000 Series Switches incorporate significant enhancements in design, power, airflow, cooling, and cabling. The 10-slot chassis has front-to-back airflow, making it a good solution for hot aisle and cold aisle deployments. The 9- and 18-slot chassis use side-to-side airflow to deliver high density in a compact form factor. The chassis in this series include Cisco Nexus 7000 9-Slot, 10-Slot, and 18-Slot Switch chassis, sometimes referred to as Cisco Nexus 7009, 7010, and 7018 chassis as seen in the figure.

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

Cisco Nexus 7000 Series Chassis Options

There are three chassis models in the Cisco Nexus 7000 Family of products. This topic describes the performance capacities and key features of the Cisco Nexus 7000 9-Slot, 10-Slot, and 18-Slot Switch chassis (also referred to as 7009, 7010, and 7018 chassis.)

Cisco Nexus 7000 Platform

- 15+ Tb/s system
- DCB*- and FCoE-ready
- Modular operating system
- Device virtualization
- Cisco TrustSec
- Continuous operations





	Nexus 7009	Nexus 7010	Nexus 7018
	Shipping	Shipping	Shipping
Slots	7 I/O + 2 sup	8 I/O + 2 sup	16 I/O + 2 sup
Height	14 RU	21 RU	25 RU
BW / Slot Fab 1	N/A	230 Gb/s per slot	230 Gb/s per slot
BW / Slot Fab 2	550 Gig / Slot	550 Gb/s per slot	550 Gb/s per slot

*DCB = Data Center Bridging

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

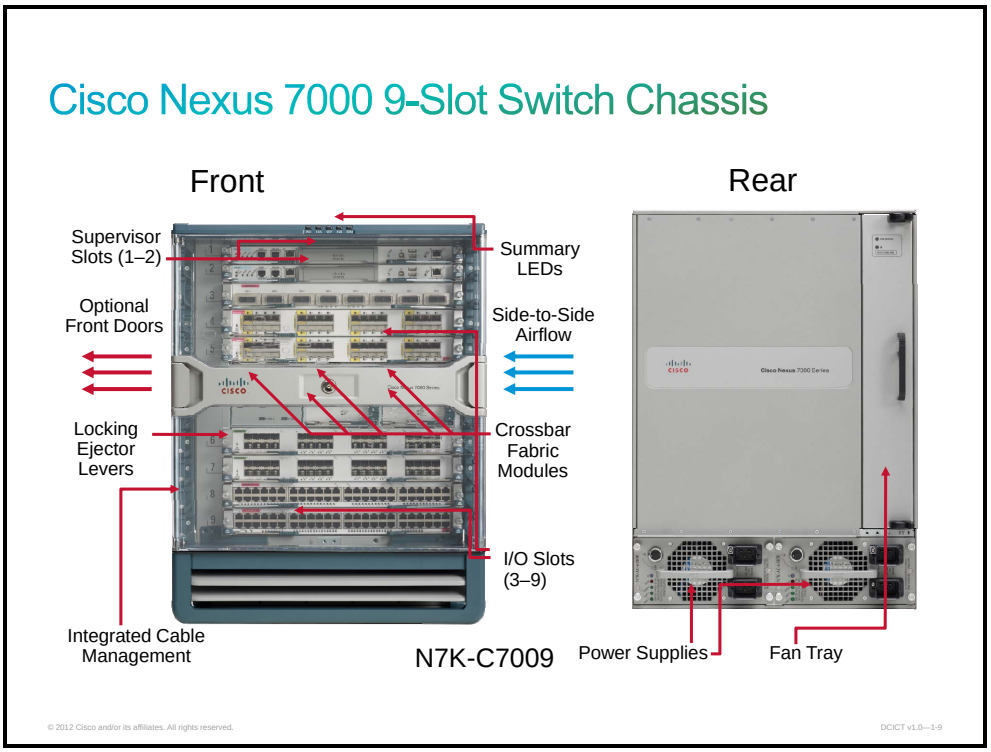
The Cisco Nexus 7000 Series Switches offer a modular data center-class product that is designed for highly scalable 10 Gigabit Ethernet networks with a fabric architecture that scales beyond 15 Tb/s. The Cisco Nexus 7000 Series provides integrated resilience that is combined with features that are optimized specifically for the data center for availability, reliability, scalability, and ease of management.

The Cisco Nexus 7000 Series Switches run the Cisco NX-OS Software to deliver a rich set of features with nonstop operation.

- This series features front-to-back airflow with 10 front-accessed vertical module slots and an integrated cable management system facilitates installation, operation, and cooling in both new and existing facilities.
- There are 9 or 18 front-accessed module slots with side-to-side airflow in a compact horizontal form factor with purpose-built integrated cable management easing operation and reducing complexity.
- Designed for reliability and maximum availability, all interface and supervisor modules are accessible from the front. Redundant power supplies, fan trays, and fabric modules are accessible completely from the rear to ensure that cabling is not disrupted during maintenance.

https://t.me/learningnets

- The system uses dual dedicated supervisor modules and fully distributed fabric architecture. There are five rear-mounted fabric modules in the 10- and 18-slot models and five front-mounted fabric modules in the 9-slot model. Combined with the chassis midplane they can deliver up to 230 Gb/s per slot for 4.1 Tb/s of forwarding capacity in the 10-slot form factor, and 7.8 Tb/s in the 18-slot form factor using the Cisco Nexus 7000 Series Fabric-1 Modules. The 9-slot form factor requires Cisco Nexus 7000 Series Fabric-2 Modules. Migrating to the Fabric-2 Module increases the bandwidth per slot to 550 Gb/s. This migration to the Fabric-2 Module increases the forwarding capacity on the 10-slot form factor to 9.9-Tb/s. On the 18-slot form factor, forwarding capacity is increased to 18.7-Tb/s, and 8.8-Tb/s for the 9-slot form factor.
- The midplane design supports flexible technology upgrades as your needs change and provides ongoing investment protection.

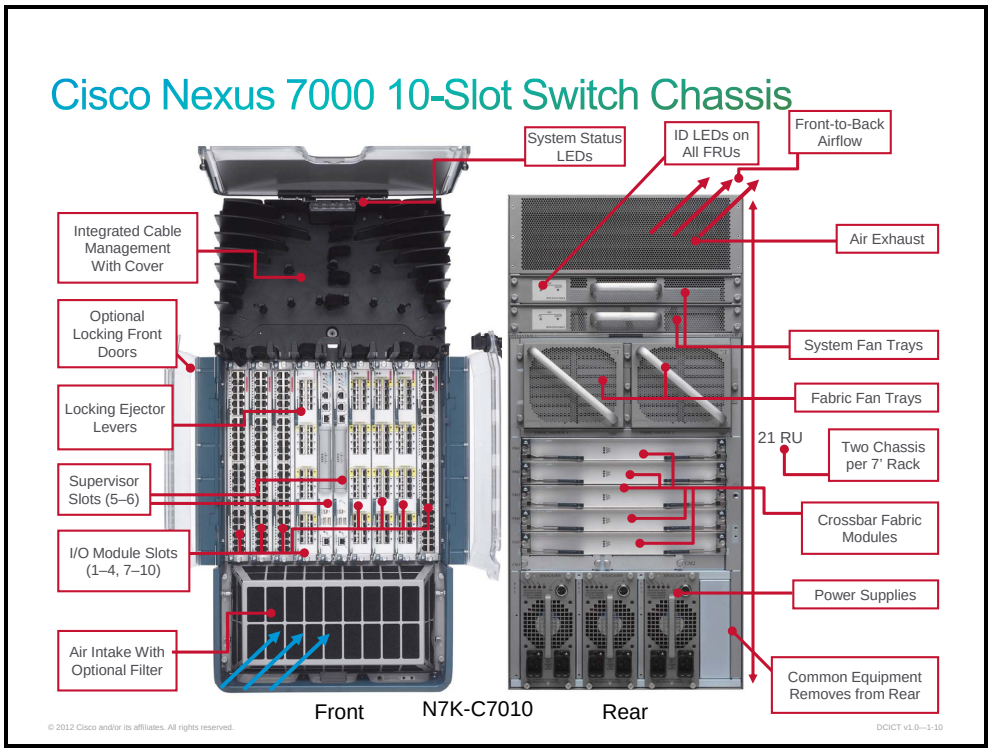


Cisco Nexus 7000 Series 9-Slot Switch Chassis

- The Cisco Nexus 7000 Series 9-Slot Switch chassis, with up to seven I/O module slots, supports up to 336 1 and 10 Gigabit Ethernet ports.
- Airflow is side-to-side.
- The integrated cable management system is designed to support the cabling requirements of a fully configured system to either or both sides of the switch, allowing maximum flexibility. All system components can easily be removed with the cabling in place, providing ease of use for maintenance tasks with minimal disruption.
- A series of LEDs at the top of the chassis provides a clear summary of the status of the major system components. The LEDs alert operators to the need to conduct further investigation. These LEDs report the power supply, fan, fabric, supervisor, and I/O module status.

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

- The purpose-built optional front module door provides protection from accidental interference with both the cabling and modules that are installed in the system. The transparent front door allows easy observation of cabling and module indicators and status lights without any need to open the doors. The door supports a dual-opening capability for flexible operation and cable installation while fitted. The door can be completely removed for both initial cabling and day-to-day management of the system.
- Independent variable-speed system and fabric fans provide efficient cooling capacity to the entire system. Fan tray redundancy features help ensure reliability of the system and support for hot-swapping of fan trays.
- The crossbar fabric modules are located in the front of the chassis, with support for two supervisors.

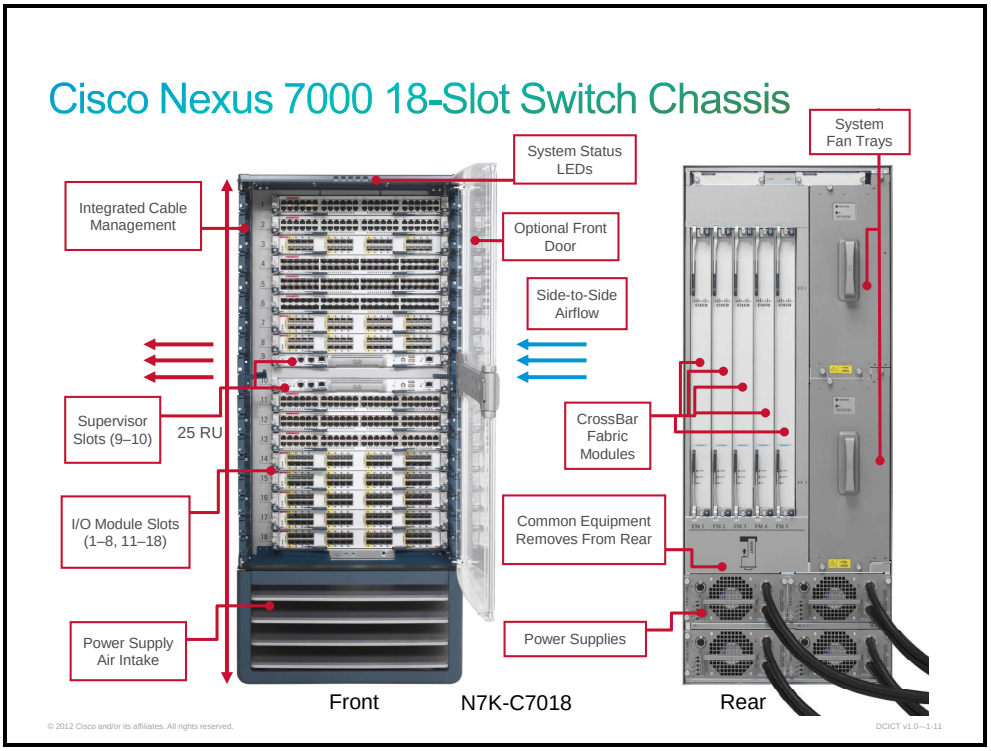


Cisco Nexus 7000 Series 10-Slot Switch Chassis

- The Cisco Nexus 7000 Series 10-Slot Switch chassis with up to eight I/O module slots supports up to 384 1 and 10 Gigabit Ethernet ports, meeting the demands of large deployments.
- Front-to-back airflow helps ensure that use of the Cisco Nexus 7000 Series 10-Slot Switch chassis addresses the requirement for hot-aisle and cold-aisle deployments without additional complexity.
- The system uses dual system and fabric fan trays for cooling. Each fan tray is redundant and composed of independent variable-speed fans that automatically adjust to the ambient temperature. This adjustment helps reduce power consumption in well-managed facilities while providing optimum operation of the switch. The system design increases cooling efficiency and provides redundancy capabilities, allowing hot-swapping without affecting the system. If either a single fan or a complete fan tray fails, the system continues to operate without a significant degradation in cooling capacity.

https://t.me/learningnets

- The integrated cable management system is designed for fully configured systems. The system allows cabling either to a single side or to both sides for maximum flexibility without obstructing any important components. This flexibility eases maintenance even when the system is fully cabled.
- The system supports an optional air filter to help ensure clean airflow through the system. The addition of the air filter satisfies Network Equipment Building Standards (NEBS) requirements.
- A series of LEDs at the top of the chassis provides a clear summary of the status of the major system components. The LEDs alert operators to the need to conduct further investigation. These LEDs report the power supply, fan, fabric, supervisor, and I/O module status.
- The cable management cover and optional front module doors provide protection from accidental interference with both the cabling and modules that are installed in the system. The transparent front door allows observation of cabling and module indicator and status lights.



Cisco Nexus 7000 Series 18-Slot Switch Chassis

- The Cisco Nexus 7000 Series 18-Slot Switch chassis with up to 16 I/O module slots supports up to 768 1 and 10 Gigabit Ethernet ports, meeting the demands of the largest deployments.
- Side-to-side airflow increases the system density within a 25 rack unit (25-RU) footprint, optimizing the use of rack space. The optimized density provides more than 16 RU of free space in a standard 42-RU rack for cable management and patching systems.
- The integrated cable management system is designed to support the cabling requirements of a fully configured system to either or both sides of the switch, allowing maximum flexibility. All system components can easily be removed with the cabling in place, providing ease of maintenance tasks with minimal disruption.

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

- A series of LEDs at the top of the chassis provides a clear summary of the status of the major system components. The LEDs alert operators to the need to conduct further investigation. These LEDs report the power supply, fan, fabric, supervisor, and I/O module status.
- The purpose-built optional front module door provides protection from accidental interference with both the cabling and modules that are installed in the system. The transparent front door allows easy observation of cabling and module indicators and status lights without any need to open the doors. The door supports a dual-opening capability for flexible operation and cable installation while fitted. The door can be completely removed for both initial cabling and day-to-day management of the system.
- Independent variable-speed system and fabric fans provide efficient cooling capacity to the entire system. Fan tray redundancy features help ensure reliability of the system and support for hot-swapping of fan trays.

Cisco Nexus 7000 Series Chassis
High Availability

High-availability software architecture

- ISSU
- VDCs
- Process modularity
- Process survivability
- Control plane and data plane separation
- Online diagnostics
- DOS-resilient
- RBAC
- Configuration verification and rollback
- Cisco IOS EEM
- Call Home



Highly available hardware architecture:

- Dual redundant supervisors
- Dual redundant connectivity management processors
- Dual redundant Ethernet OOB
- Dual redundant active central arbiters
- Fabric modules not dependent on centralized clock
- Majority voting on all critical signals
- Dual independent system management busses
- Multiple redundant fan modules

DOS = denial of service
RBAC = role-based access control
Cisco IOS EEM = Cisco IOS Embedded Event Manager

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

The Cisco Nexus 7000 Series Switches are modular in design with emphasis on redundant critical components throughout the subsystems. This modular approach has been applied across the physical, environmental, power, and system software aspects of the chassis architecture.

- **Supervisor module redundancy:** Active and standby operation with state and configuration synchronization between the two supervisors. Provides seamless and Stateful Switchover (SSO) in the event of a supervisor module failure.
- **Switch fabric redundancy:** A single chassis can be configured with one or more fabric modules up to a maximum of five modules, providing capacity as well as redundancy. Failure of a switch fabric module triggers an automatic reallocation and balancing of traffic across the remaining active switch fabric modules.
- **Cooling subsystem:** Two redundant system fan trays for I/O module cooling and two additional redundant fan trays for switch fabric module cooling. All fan trays are hot-swappable.

https://t.me/learningnets

- **Power subsystem availability features:** Three internally redundant power supplies. Each is composed of two internalized, isolated power units providing two power paths per modular power supply, and six paths total per chassis when fully populated.
- **Internal EOBC:** A switched Ethernet out-of-band channel (EOBC) is used for management and control of traffic between the supervisors and I/O modules. EOBC is also used between the supervisors and fabric modules. On the supervisor modules, there is an onboard 24-port Ethernet switch on a chip. It provides one 1-Gb/s Ethernet link from each supervisor to each I/O module, and from each supervisor to each switch fabric module (up to five), and between the two supervisors. Two additional redundant 1-Gb/s Ethernet links are used on each supervisor to connect to the local CPU within the supervisor. This configuration provides a highly redundant switched Ethernet-based fabric for control and management of traffic.
- **Cisco NX-OS Software redundancy options and features:** The Cisco NX-OS Software compartmentalizes components for redundancy, fault isolation, and resource efficiency. Functional feature components operate as independent processes referred to as services, with availability features implemented into each service. Most system services are capable of performing stateful restarts that are transparent to other services within the platform and neighboring devices within the network.
- **Service modularity:** Services within the Cisco NX-OS Software are designed as nonkernel space processes. They perform a function or set of functions for a subsystem or feature set, with each service instance running as a separate independent protected process. This modular architecture permits the system to provide a high level of protection and fault tolerance for all services running on the platform.
- **Modular software upgrades:** These upgrades address specific issues while minimizing the impact to other critical services and the system overall.
- **Rapid restart:** The services in the Cisco NX-OS Software can be restarted automatically by the system manager in the event of critical fault detection. Notification of restart and restart both happen in milliseconds.
- **Cisco NX-OS In-Service Software Upgrade (ISSU):** The modular software architecture of the Cisco NX-OS Software supports plug-in-based services and features. This architecture makes it possible to perform complete image upgrades nondisruptively without impacting the data-forwarding plane.
- **Virtual device contexts (VDCs):** The Cisco NX-OS Software implements a logical virtualization at the device level. Logical virtualization allows multiple instances of the device to operate on the same physical switch simultaneously. These logical operating environments are known as VDCs.

Note Currently, the clock modules are not used, and there is no systemwide clock.

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

Cisco Nexus 7000 Series Supervisor Module

A Cisco Nexus 7000 Series switch has two slots that are available for supervisor modules. High availability is provided through populating both supervisor slots. This topic describes the capabilities of Cisco Nexus 7000 Series Supervisor Module.

Cisco Nexus 7000 Series
Supervisor Engine 1

- Interfaces
  - Supervisor management port: 10/100/1000-Mb/s Ethernet port, support for inline encryption through MAC security (IEEE 802.1AE)
  - CMP management port: 10/100/1000-Mb/s Ethernet port
  - Console serial port: RJ-45 connector
  - Auxiliary serial port: RJ-45 connector
  - Three USB ports: two host and one device port for peripheral devices
- Memory
  - DRAM: 8 GB
  - Flash memory: 2 GB
  - NVRAM: 2-MB battery backup



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

The Cisco Nexus 7000 Series Supervisor Module is designed to deliver scalable control plane and management functions for the Cisco Nexus 7000 Series chassis. It is based on a dual-core processor that scales the control plane by harnessing the flexibility and power of the dual cores. The supervisors control the Layer 2 and Layer 3 services, redundancy capabilities, configuration management, status monitoring, power and environmental management, and more. Supervisors provide centralized arbitration to the system fabric for all line cards. The fully distributed forwarding architecture allows the supervisor to support transparent upgrades to higher forwarding capacity-capable I/O and fabric modules. The supervisor incorporates an innovative dedicated connectivity management processor (CMP) to support remote management and troubleshooting of the complete system. Two supervisors are required for a fully redundant system. One supervisor module runs as the active device while the other is in hot-standby mode. This redundancy provides exceptional high-availability features in data center-class products.

To deliver a comprehensive set of features, the Cisco Nexus 7000 Series Supervisor Module offers the following:

- Continuous system operation
  - Active and standby supervisor
  - Segmented and redundant out-of-band (OOB) provisioning and management paths
  - Virtualization of the management plane
  - Integrated diagnostics and protocol decoding with an embedded control plane packet analyzer

https://t.me/learningnets

- Upgradable architecture
  - Fully decoupled control plane and data plane with no hardware forwarding on the module
  - Distributed forwarding architecture, allowing independent upgrades of the supervisor and fabric
  - Cisco Unified Fabric-ready
  - Transparent upgrade capacity and capability, which is designed to support 40 and 100 Gigabit Ethernet
- Superior operational efficiency
  - System locator and beacon LEDs for simplified operations
  - Dedicated OOB management processor for “lights out” management

Supervisor CMP

The CMP provides a complete OOB management and monitoring capability independent from the primary operating system. The CMP enables “lights out” remote monitoring and management of the supervisor module, all modules, and the Cisco Nexus 7000 Series system without the need for separate terminal servers with the associated additional complexity and cost. The CMP delivers the remote control through its own dedicated processor, memory, and bootflash memory and a separate Ethernet management port. The CMP can reset all system components, including power supplies. It can also reset the host supervisor module to which it is attached, allowing a complete system restart.

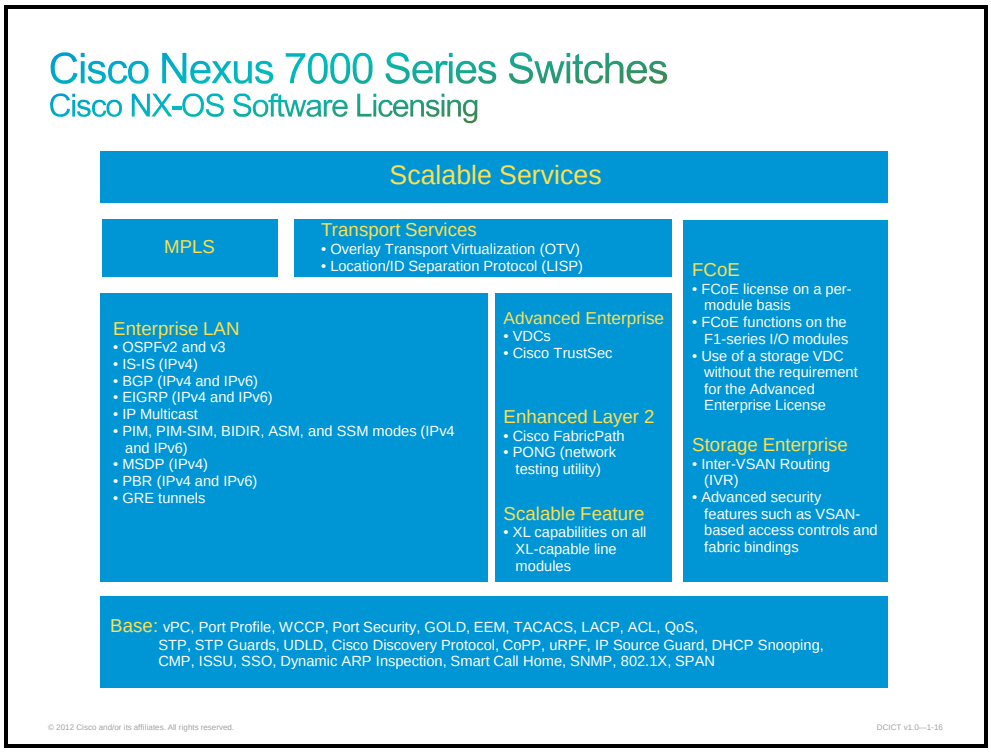
Advanced Diagnostics and Troubleshooting Tools

Management of large-scale data center networks requires proactive management tools to verify connectivity and mechanisms for capturing and analyzing traffic. The Cisco Nexus 7000 Series Supervisor Module incorporates highly advanced analysis and debugging capabilities. The power-on self-test (POST) and Cisco Generic Online Diagnostics (Cisco GOLD) provide proactive health monitoring both at startup and during system operation. The supervisor module uniquely provides a built-in packet capture and protocol decoding tool that allows analysis of control plane traffic. This analysis improves network planning, provides faster operation response times to events, and reduces operating costs.

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

Cisco Nexus 7000 Series Licensing Options

The Cisco Nexus switches use a licensing option for additional features. This topic describes the licensing options for the Cisco Nexus 7000 Series Switches.

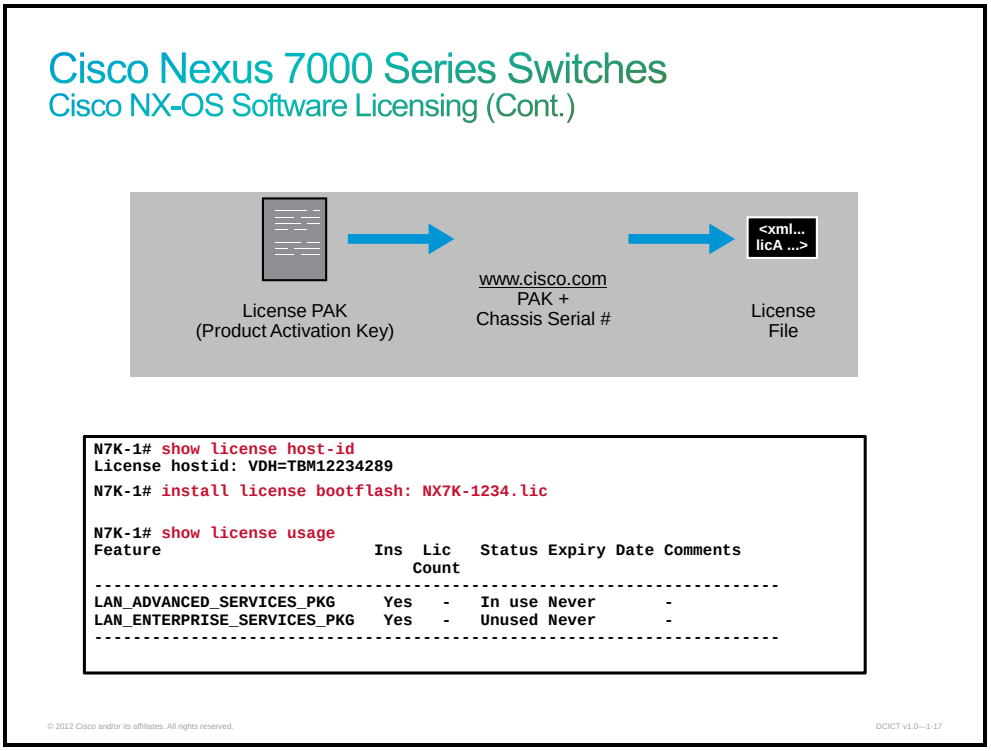


The Cisco NX-OS Software for the Cisco Nexus 7000 Series Switches is available in incremental license levels.

- **Base:** A rich feature set is provided with the base software, which is bundled with the hardware at no extra cost.
- **Enterprise LAN:** The Enterprise LAN License enables incremental functions that are applicable to many enterprise deployments, such as dynamic routing protocols (Open Shortest Path First [OSPF], Enhanced Interior Gateway Routing Protocol [EIGRP]; Intermediate System-to Intermediate System [IS-IS], Border Gateway Protocol [BGP]), IP multicast; Protocol Independent Multicast (PIM), PIM sparse mode (PIM-SM), Bidirectional PIM (BIDIR-PIM), Advanced Services Module (ASM), and Source Specific Multicast (SSM) modes; Multicast Source Discovery Protocol (MSDP), policy-based routing (PBR), and Generic Routing Encapsulation (GRE).
- **Advanced LAN Enterprise:** The Advanced LAN Enterprise License enables next-generation functions such as VDCs and the Cisco TrustSec solution.
- **MPLS:** The MPLS License enables Multiprotocol Label Switching (MPLS).
- **Transport Services:** The Transport Services License enables Overlay Transport Virtualization (OTV), an IP-based data center interconnect technology, and Cisco Locator/ID Separation Protocol (Cisco LISP). LISP is an evolutionary routing architecture that is designed for Internet scale and global reach across organizations.
- **Enhanced Layer 2:** The Enhanced Layer 2 License enables Cisco FabricPath, the latest Cisco technology to massively scale Layer 2 data centers.

https://t.me/learningnets

- **Scalable Services:** The Scalable Services License is applied on a per-chassis basis and enables XL capabilities on the line cards. A single license per system enables all XL-capable I/O modules to operate in XL mode. After the single system license is added to a system, all modules that are XL-capable are enabled with no additional licensing.
- **SAN Enterprise:** This license enables Inter-VSAN Routing (IVR), and advanced security features such as VSAN-based access controls, and fabric bindings for open systems.
- **FCoE License for the 32-port 10 Gigabit Ethernet:** This license enables director-class multihop FCoE implementation in a highly available modular switching platform for access and core of converged network fabric. FCoE is supported on the Cisco Nexus 7000 F1-Series I/O modules. This license also enables the use of a storage VDC for the FCoE traffic within the Cisco Nexus 7000 Series, with requiring the enablement of the Advanced package.
- **DCNM for Cisco Nexus 7000**
- **DCNM for SAN Advanced**



Cisco NX-OS Software uses feature-based licenses that are enforced on the switch. They are tied to the chassis serial number that is stored in dual redundant NVRAM modules on the backplane.

Licenses are issued in the form of a digitally signed text file that is installed in the bootflash, and run using the **install license** command.

Associated with a license is a grace period, during which a feature can be run without having a license installed. The grace period permits the feature to be trial-tested before committing to using that feature and purchasing the relevant license. Periodic syslog, Call Home, and Simple Network Management Protocol (SNMP) traps issue warnings when the grace period is nearing expiration. When the grace period ends, the associated features become unavailable for use until a license is installed to activate them. The grace period is 120 days.

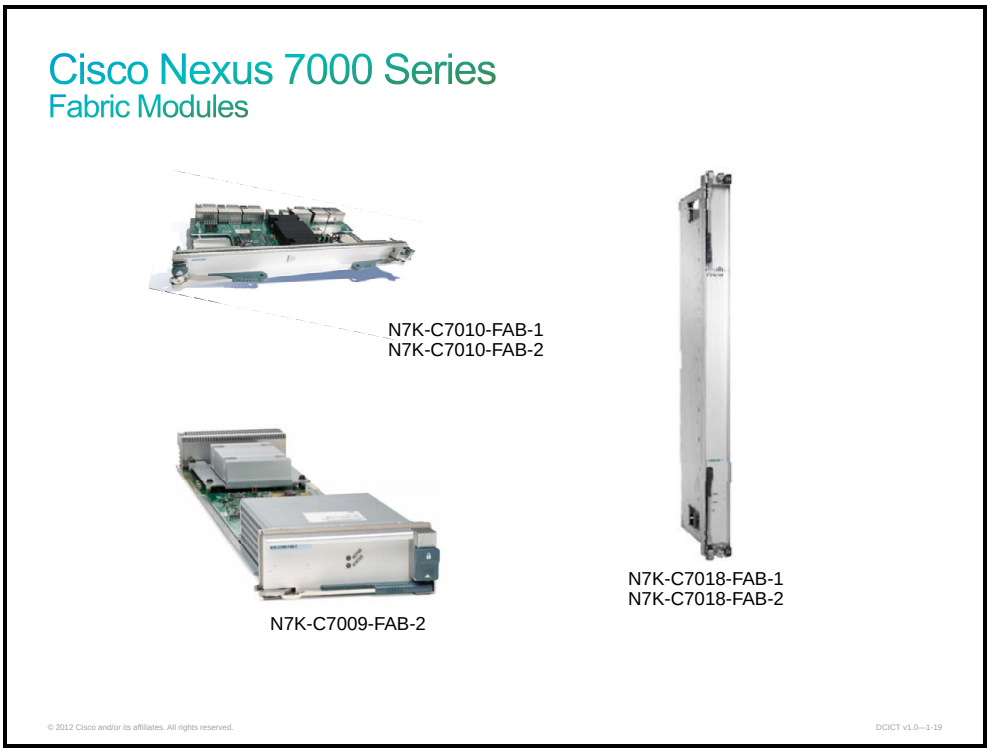
© 2012 Cisco Systems, Inc.
https://t.me/learningnets

In addition, there are time-bound licenses that are currently used in an emergency. An example would be when a grace period has expired and you need additional time to purchase the license but do not want to lose the use of a feature. The expiration date on a time-bound license is absolute and expires at midnight Coordinated Universal Time (UTC) on the set date. Periodic syslog, Call Home, and SNMP trap warnings are issued when time-bound licenses near their expiration date. When the time-bound expiration date is reached, the relevant features are no longer available unless a license is installed or there is additional time in the grace period.

https://t.me/learningnets

Cisco Nexus 7000 Series Fabric Modules

The Cisco Nexus 7000 Series uses switch fabric modules either located in the rear of the chassis or the front, as on the Cisco Nexus 9-Slot Switch chassis. This topic describes the capacities of the Cisco Nexus 7000 Series Fabric Modules.



The Cisco Nexus 7000 Series Fabric Modules for the Cisco Nexus 7000 Series chassis are separate fabric modules that provide parallel fabric channels to each I/O and supervisor module slot. Up to five simultaneously active fabric modules work together delivering up to 230 Gb/s per slot (Fabric-1 Modules) or up to 550 Gb/s per slot (Fabric-2 Modules). The fabric module provides the central switching element for the fully distributed forwarding on the I/O modules.

Switch fabric scalability is made possible through the support of one to five concurrently active fabric modules for increased performance as your needs grow. All fabric modules are connected to all module slots. The addition of each fabric module increases the bandwidth to all module slots up to the system limit of five modules. The architecture supports lossless fabric failover, with the remaining fabric modules load balancing the bandwidth to all the I/O module slots, helping ensure graceful degradation.

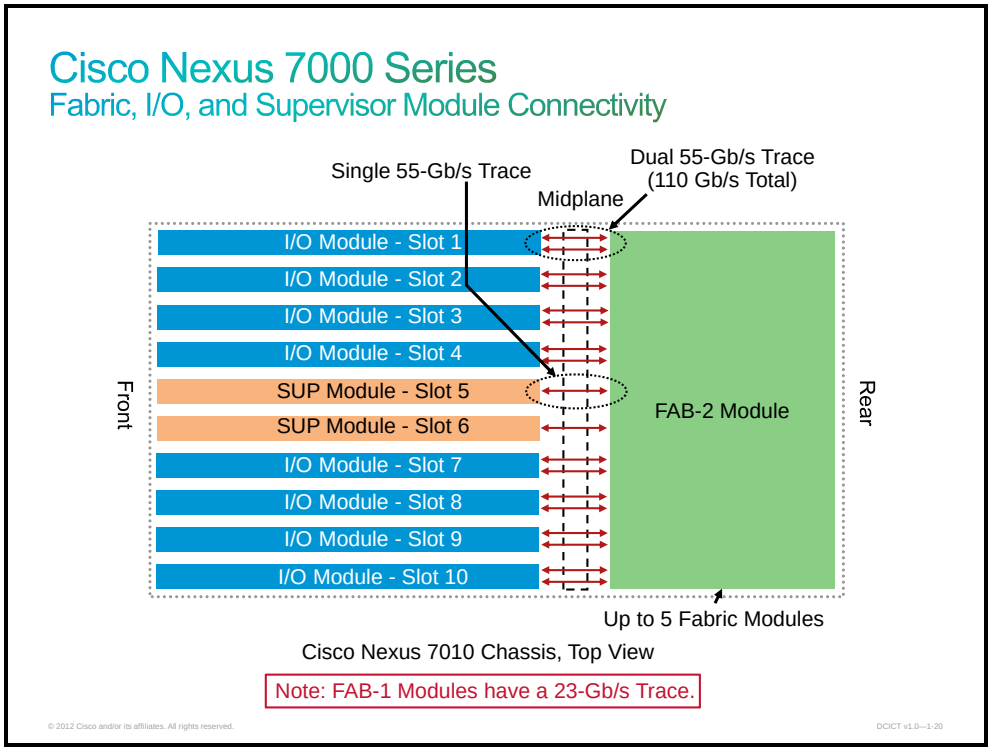
The combination of a Cisco Nexus 7000 Fabric Module and the supervisor and I/O modules supports virtual output queuing (VOQ) and credit-based arbitration to the crossbar switch to increase performance of the distributed forwarding system. VOQ and credit-based arbitration facilitate fair sharing of resources when a speed mismatch or contention for an uplink interface exists. The fabric architecture also enables support for lossless Ethernet and unified I/O capabilities.

The following should be considered for Cisco Nexus 7000 Fabric-2 Modules:

- **Cisco Nexus 7000 9-Slot Switch Chassis**:
  - Cisco NX-OS Software Release 5.2 or higher
  - Does not support Cisco Nexus 7000 Fabric-1 Modules
- **Cisco Nexus 7000 8- and 10-Slot Switch Chassis**:
  - Cisco NX-OS Software Release 6.0 or higher

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



This figure describes the connectivity between the Cisco Nexus 7000 Series I/O and Supervisor Modules and the Fabric Modules. Each I/O module has two 55-Gb/s traces to each fabric module. Therefore, a fully loaded Cisco Nexus 7000 Series chassis provides 550 Gb/s of switching capacity per I/O slot.

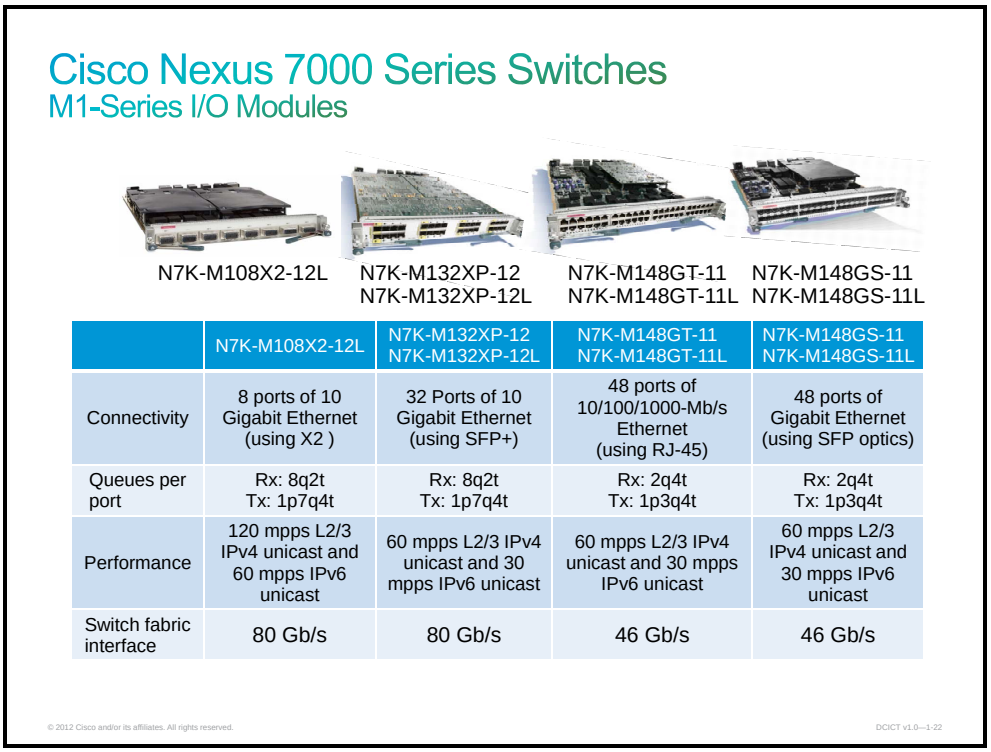
In addition, each supervisor module has a single 55-Gb/s trace to each fabric module. This means that a fully loaded Cisco Nexus 7000 10-Slot Switch chassis provides 275 Gb/s of switching capacity to each supervisor slot.

Note	Cisco Nexus 7000 Fabric-1 Modules provide two 23-Gb/s traces to each fabric module, providing 230 Gb/s of switching capacity per I/O slot for a fully loaded chassis. Each supervisor module has a single 23-Gb/s trace to each fabric module.

https://t.me/learningnets

Cisco Nexus 7000 Series I/O Modules

The Cisco Nexus 7000 Series is a module series of switches supporting various I/O modules. These I/O modules can be feature- or performance-rich. This topic describes the capabilities of the Cisco Nexus 7000 Series I/O modules.



	N7K-M108X2-12L	N7K-M132XP-12 N7K-M132XP-12L	N7K-M148GT-11 N7K-M148GT-11L	N7K-M148GS-11 N7K-M148GS-11L
Connectivity	8 ports of 10 Gigabit Ethernet (using X2)	32 Ports of 10 Gigabit Ethernet (using SFP+)	48 ports of 10/100/1000-Mb/s Ethernet (using RJ-45)	48 ports of Gigabit Ethernet (using SFP optics)
Queues per port	Rx: 8q2t Tx: 1p7q4t	Rx: 8q2t Tx: 1p7q4t	Rx: 2q4t Tx: 1p3q4t	Rx: 2q4t Tx: 1p3q4t
Performance	120 mpps L2/3 IPv4 unicast and 60 mpps IPv6 unicast	60 mpps L2/3 IPv4 unicast and 30 mpps IPv6 unicast	60 mpps L2/3 IPv4 unicast and 30 mpps IPv6 unicast	60 mpps L2/3 IPv4 unicast and 30 mpps IPv6 unicast
Switch fabric interface	80 Gb/s	80 Gb/s	46 Gb/s	46 Gb/s

M1-Series 8-Port 10 Gigabit Ethernet Module

The Cisco Nexus 7000 M1-Series 8-port 10 Gigabit Ethernet Module is designed for performance-driven, mission-critical Ethernet networks, with 80 Gb/s (full-duplex 160 Gb/s) of bandwidth to the fabric. The module provides the same features as the N7K-M132XP-12 but without oversubscription at the front panel.

M1-Series 32-Port 10 Gigabit Ethernet Module

The Cisco Nexus 7000 M1-Series 32-port 10 Gigabit Ethernet Module is designed for performance-driven, mission-critical Ethernet networks, with 80 Gb/s (full duplex 160 Gb/s) of bandwidth to the fabric. Populating the 10-slot chassis with this module provides up to 256 ports of 10 Gigabit Ethernet in a single chassis. You may configure up to 512 ports of 10 Gigabit Ethernet in a single rack. The physical interfaces on the module support SFP+ optics, which support various distances and types of fiber cable.

All I/O modules contain an integrated forwarding engine, which is part of the Cisco Nexus 7000 M Series of forwarding engines. It is the first generation of the M Series and is referred to as the M1 forwarding engine.

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

M1-Series 48-Port Gigabit Ethernet Module

The Cisco Nexus 7000 M1-Series 48-port Gigabit Ethernet modules N7K-M148GT-11 and N7K-M148GT-11L provide 40 Gb/s of bandwidth (full-duplex 80 Gb/s) to the fabric and are designed for mission-critical Ethernet networks. Populating a 10-slot chassis with this module delivers up to 384 ports of 10-, 100-, or 1000-Mb/s Ethernet in a single chassis, or up to 768 ports of 10-, 100-, or 1000-Mb/s Ethernet in a single rack.

The 48-port Gigabit Ethernet small form-factor pluggable (SFP) I/O modules N7K-M148GS-11 and N7K-M148GS-11L are designed for mission-critical Ethernet networks that require high performance. The 48-port Gigabit Ethernet SFP modules are supported by both the 18-slot and the 10-slot chassis.

The 48-port Gigabit Ethernet modules contain an integrated forwarding engine as do all I/O modules. This one is part of the Cisco Nexus 7000 M Series of forwarding engines. This series is the first generation of the M series and is referred to as the M1 forwarding engine.

Cisco Nexus 7000 Series Switches
M2-Series I/O Modules



N7K-M206FQ-23L N7K-M202CF-22L

	N7K-M206FQ-23L	N7K-M202CF-22L
Connectivity	6 ports of 40 Gigabit Ethernet (using QSFP+)	2 Ports of 100 Gigabit Ethernet (using CFP)
Queues per port	Rx: 8q2t Tx: 1p7q4t	Rx: 8q2t Tx: 1p7q4t
Performance	120 mpps L2/3 IPv4 unicast and 60 mpps IPv6 unicast	120 mpps L2/3 IPv4 unicast and 60 mpps IPv6 unicast
Switch fabric interface	550 Gb/s (FAB-2) 230 Gb/s (FAB-1)	550 Gb/s (FAB-2) 230 Gb/s (FAB-1)

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

The Cisco Nexus 7000 M2-Series I/O modules are highly scalable, high-performance modules offering outstanding flexibility, and full-featured, nonblocking performance on each port. The Cisco Nexus 7000 M2-Series modules facilitate the deployment of high-density, high-bandwidth, scalable network architectures, especially in large network cores and in service provider and Internet peering environments.

The first two Cisco Nexus 7000 M2-Series I/O modules are the Cisco Nexus 7000 M2-Series 6-Port 40 Gigabit Ethernet Module and the Cisco Nexus 7000 M2-Series 2-port 100 Gigabit Ethernet module. These two modules will be discussed in the following two figures.

Both modules can use either the Cisco Fabric-1 Module or Fabric-2 Module.

https://t.me/learningnets

Cisco Nexus 7000 Series Switches
F-Series I/O Modules



N7K-F132XP-15



N7K-F248XP-25

	N7K-F132XP-15	N7K-F248XP-25
Connectivity	32 ports of Gigabit Ethernet and 10 Gigabit Ethernet (using SFP or SFP+)	48 ports of Gigabit Ethernet and 10 Gigabit Ethernet (using SFP or SFP+)
Queues Per Port	Rx: 4q4t and 2q4t Tx: 1p3q1t, 2p2q1t, 3p1q1t, 2p6q1t, 3p5q1t, and 1p7q1t	Rx: 4q4t and 2q4t Tx: 1p3q1t, 2p2q1t, 3p1q1t
Performance	480 mpps Layer 2 forwarding capacity	720 mpps Layer 2 and Layer 3 forwarding capacity for both IPv4 and IPv6
Switch Fabric Interface	320 Gb/s Local 230 Gb/s Fabric	550 Gb/s Fabric

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

F1-Series 32-Port 1 and 10 Gigabit Ethernet Module

The Cisco Nexus 7000 F1-Series 32-port 1 and 10 Gigabit Ethernet Module offers outstanding flexibility and performance, with extensive virtualization and multipath capabilities. The module enables the deployment of high-density, low-latency, scalable data center architectures. The Cisco Nexus 7000 F1-Series 32-port 1 and 10 Gigabit Ethernet module is the first of the Cisco Nexus 7000 F-Series Ethernet modules. It extends the capabilities of the Cisco Nexus 7000 Series in the data center. Powered by the F1 Forwarding Engine, the module delivers 480 millions of packets per second (mpps) of distributed Layer 2 Forwarding. The module can also deliver up to 320 Gb/s of data throughput through a custom dedicated forwarding ASIC, enabling the creation of extensible data centers.

A Cisco Nexus 7000 18-slot Switch that is fully populated with Cisco Nexus 32-port 1 and 10 Gigabit Ethernet modules can deliver up to 10.2 Tb/s of switching performance, with a typical power consumption of less than 10 W per port. The F1-Series delivers integrated hardware support for FCoE and IEEE Data Center Bridging (DCB) features. It greatly simplifies the network infrastructure. It also reduces costs by enabling the deployment of unified data center fabrics to consolidate data center traffic onto a single, general-purpose, high-performance, highly available network.

F2-Series 48-Port 1 and 10 Gigabit Ethernet Module

The Cisco Nexus 7000 F2-Series 48-Port 1 and 10 Gigabit Ethernet Module offers outstanding flexibility and wire-rate performance on each port. The module enables the deployment of high-density, low-latency, scalable data center architectures.

The Cisco Nexus 7000 F2-Series module is a low-latency, high-performance, high-density 10 Gigabit Ethernet module that is designed for mission-critical data center networks. Up to 768 wire-rate 10 Gigabit Ethernet ports are supported in a single system by using the Cisco Nexus 7000 18-Slot Switch chassis, providing the highest-density of wire-rate 10 Gigabit Ethernet ports on the market. Populating the Cisco Nexus 7000 10- and 9-Slot Switch chassis with this module delivers, respectively, up to 384 and 336 ports of 10 Gigabit Ethernet in a single chassis.

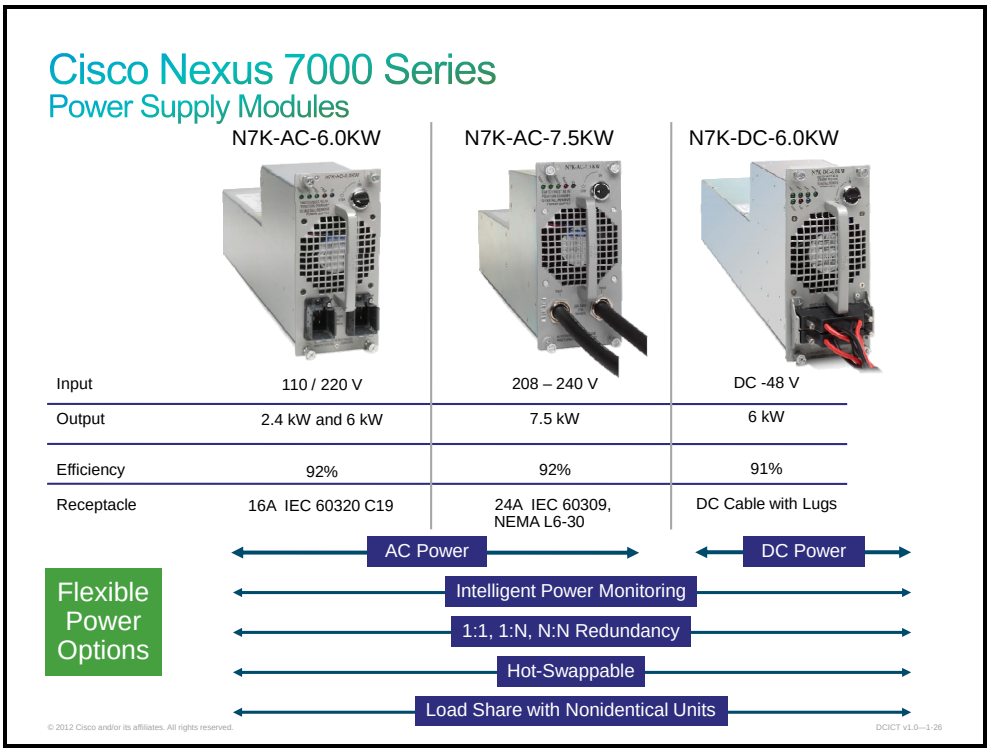
https://t.me/learningnets

The Cisco Nexus 7000 F2-Series Module can also be used with the Cisco Nexus 2000 Series Fabric Extenders.

https://t.me/learningnets

Cisco Nexus 7000 Series Power Supply Options

There are four redundancy modes that are available for power supplies on the Cisco Nexus 7000 Series Switches. This topic describes capabilities of the Cisco Nexus 7000 Series power supply modules and the redundancy modes that are available.



Cisco Nexus 7000 6.0-kW AC Power Supply Module

The 6-kilowatt (6-kW) AC power supply module is the first of the Cisco Nexus 7000 Series power supply modules. Each 10-slot chassis can hold up to three load-sharing, fault-tolerant, hot-swappable power supply modules. The power supply modules are installed in the rear of the chassis for easy installation and removal without obstruction by the cables at the front.

The power supply module is a dual 20-ampere (20-A) AC input unit providing the following:

- **Single Input:** 220 volt (220 V), 3000 W output, 110 V, 1200 W output
- **Dual Input:** 220 V, 6000 W output, 110 V, 2400 W output
- **Dual Input:** 110 and 220 V, 4200 W output

The power supply module has four user-configurable power-redundancy modes.

Key features of the power supply modules include the following:

- Multiple inputs providing redundancy if one input fails
- Universal input providing flexibility
- Compatibility with future Cisco Nexus 7000 Series chassis
- Hot-swappable, therefore no downtime when replacing power supply modules
- Temperature sensor and instrumentation that shut down the power supply if the temperature exceeds the thresholds, preventing damage due to overheating

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

- Internal fault monitoring so that if a short circuit and component failure are detected, the power supply unit can be shut down automatically
- Intelligent remote management so that users can remotely power-cycle one or all power supply modules using the supervisor command-line interface
- Real-time power draw showing real-time actual power consumption (not available in the initial software release)
- Variable fan speed, allowing reduction in fan speed for lower power usage in well-controlled environments

Cisco Nexus 7000 7.5-kW AC Power Supply Module

The Cisco Nexus7000 7.5-kW AC Dual 30 A power supply module that is shown in the figure delivers fault-tolerance, high-efficiency, load-sharing, and hot-swap features to the Cisco Nexus 7000 Series. Each Cisco Nexus 7000 Series chassis can accommodate multiple power supply modules, providing both chassis-level and facility power fault tolerance.

Cisco Nexus 7000 6.0-kW DC Power Supply Module

The Cisco Nexus 7000 Series Switch 6.0-kW DC power supply is designed for DC environments. It is a variable-output, high-capacity power supply module that is scalable from 3000 to 6000 W, delivering fault-tolerant load-sharing capability. The power supply modules are hot-swappable.

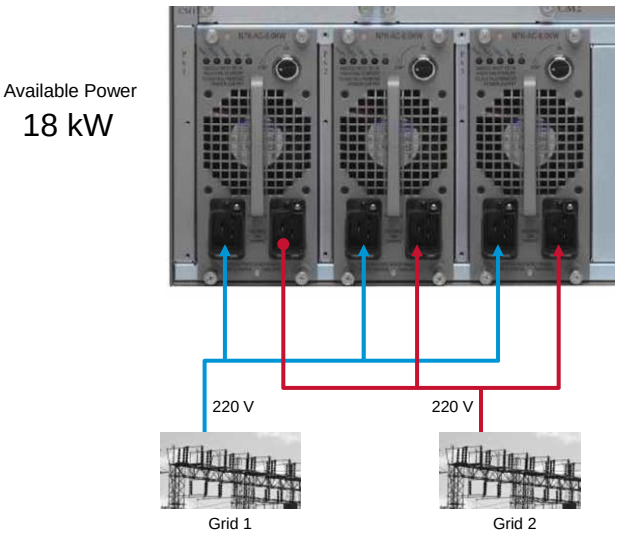
DC power connections are made using a hot-swappable DC power cable that enables quick and easy installation of the power supply modules without the need to disturb the DC terminal blocks. The DC cable supports both direct connection to DC power sources and connection to an intermediate power interface unit. This cable is used in situations where connections to the source are beyond the cable length.

The Cisco Nexus 7000 Series DC Power Interface Unit (PIU) is an optional element that is provided for environments in which the Cisco Nexus 7000 Series DC cable needs to connect to existing DC power cabling. It provides 16 two-pole terminal connections. The PIU supports one or two Cisco Nexus 7000 6.0-kW DC power supply modules, with each module using two DC power cables for a total of four connections to the PIU.

The Cisco Nexus 7000 9-Slot Switch chassis supports up to three power supply modules in a single chassis, the Cisco Nexus 7000 10-Slot Switch chassis supports up to three power supply modules, and the Cisco Nexus 7000 18-Slot Switch chassis supports up to four power supply modules.

https://t.me/learningnets

Cisco Nexus 7000 10-Slot Switch
Power Redundancy



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

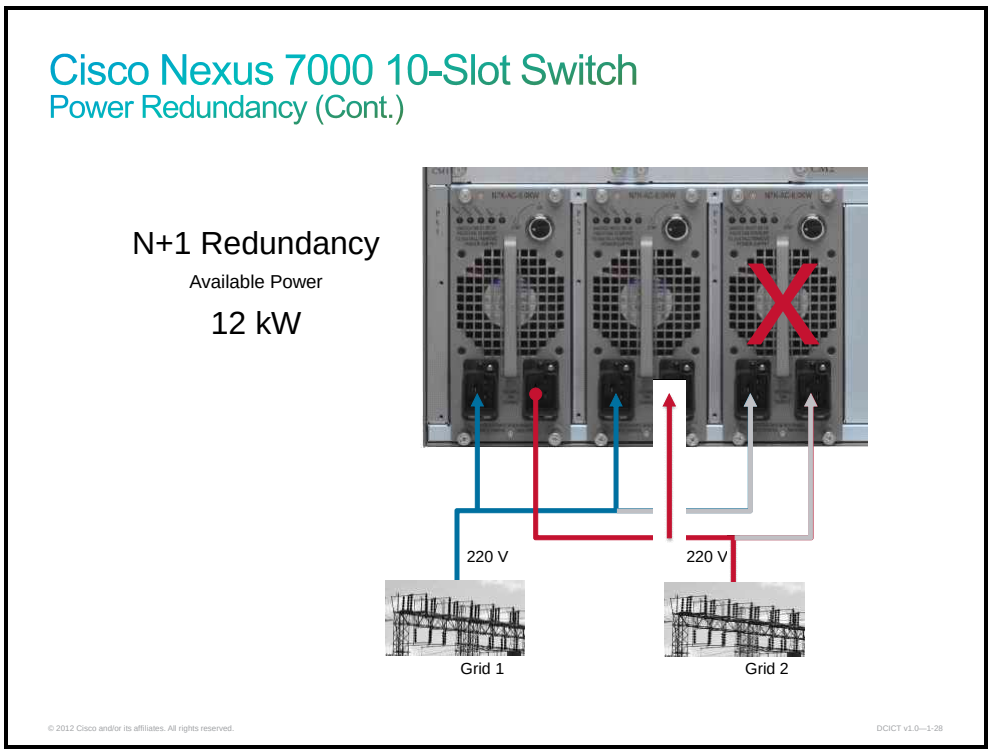
The power redundancy mode dictates how the system budgets its power. There are four user-configurable power-redundancy modes:

- Combined
- Power supply redundancy (N+1)
- Input source redundancy (grid redundancy)
- Power supply and input source redundancy (complete redundancy)

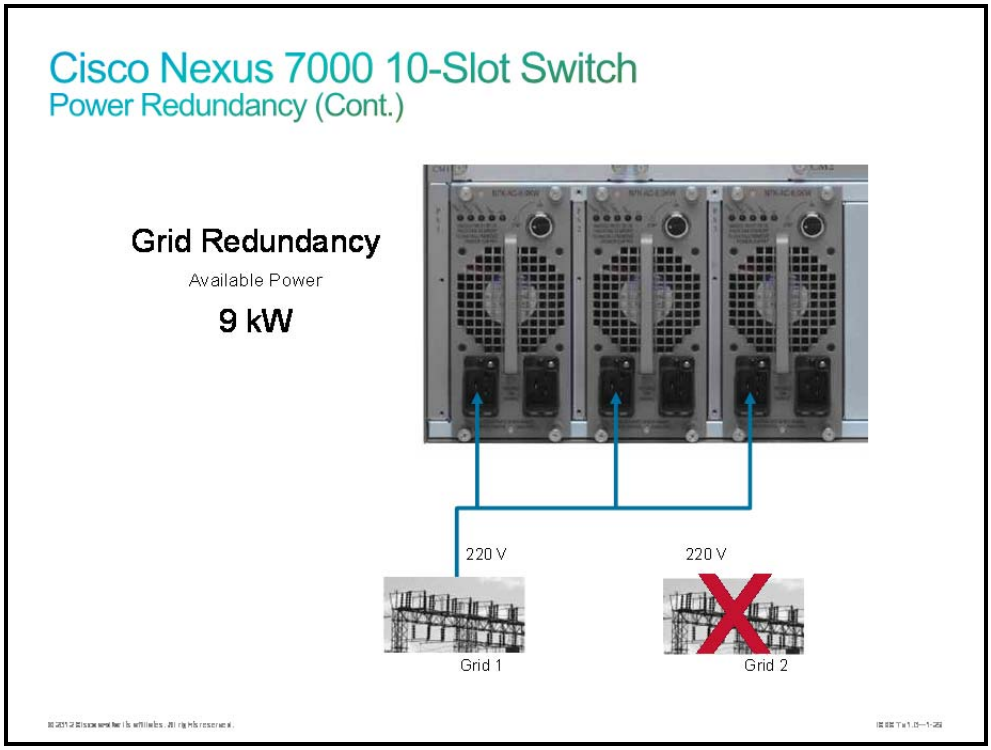
Combined mode has no redundancy, with the power that is available to the system being the sum of the power outputs of all the power supply modules in the chassis.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

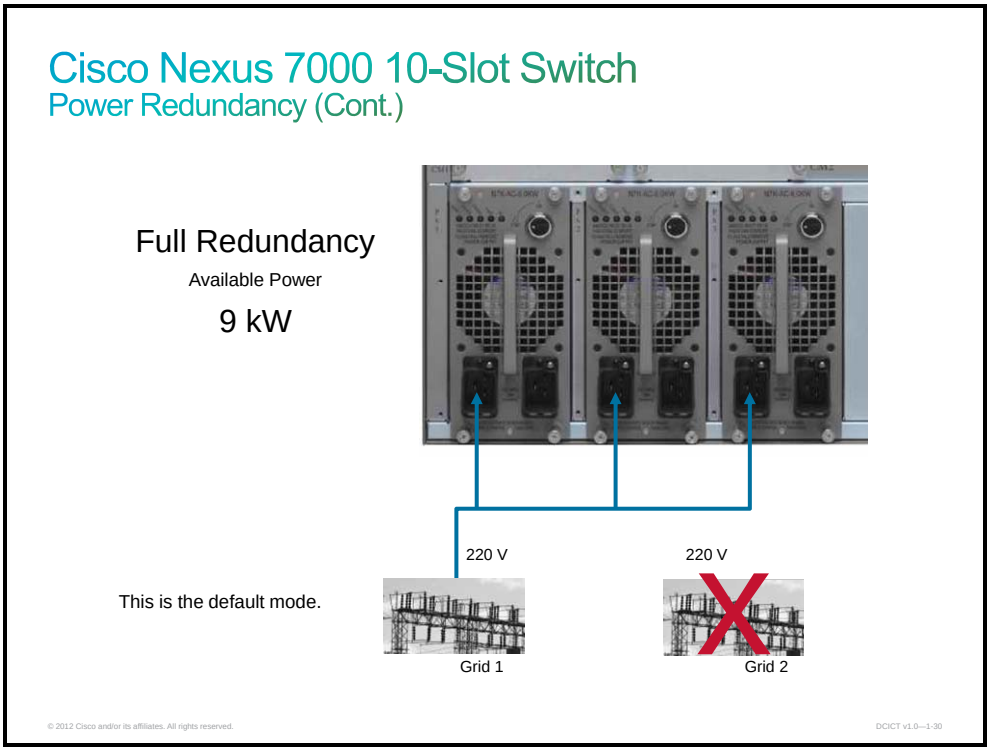


N+1 redundancy is a feature that guards against failure of one of the power supply modules. The power that is available to the system is the sum of the two least-rated power supply modules.



Grid redundancy guards against failure of one input circuit (grid). For grid redundancy, each input on the power supply is connected to an independent AC feed, and the power that is available to the system is the minimum power from either of the input sources (grids).

https://t.me/learningnets



Complete redundancy is the system default redundancy mode. This mode guards against failure of either one power supply or one AC grid, and the power that is available is always the minimum of input source and power supply redundancy.

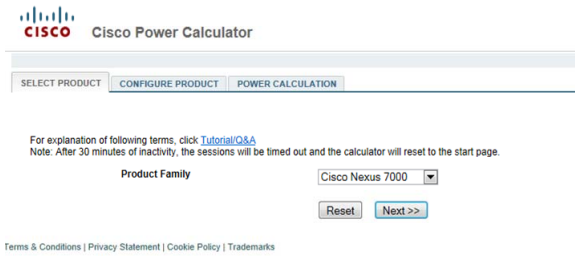
© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Power Calculator

Click the link to launch the Cisco Power Calculator tool.





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

The Cisco Power Calculator enables you to calculate the power supply requirements for a specific configuration. The results will show output current, output power, and system heat dissipation. The Cisco Power Calculator supports the Cisco Nexus 7000 Series Switches, Cisco Catalyst 6500 Series Switches, Cisco Catalyst 4500-E Series chassis and 4500 Series Switches, Cisco Catalyst 3750-E and 3750 Series Switches, Cisco Catalyst 3560-E and 3560 Series Switches, Cisco Catalyst 2960 Series Switches, Cisco Catalyst 2975 Series Switches, Cisco Catalyst Express 500 Series Switches, and the Cisco 7600 Series Routers.

Note	The calculator is a starting point in planning your power requirements; it does not provide a final power recommendation.

The Cisco Power Calculator will guide you through a series of selections for configurable products. If you need to change a previous selection, there is a Back button.

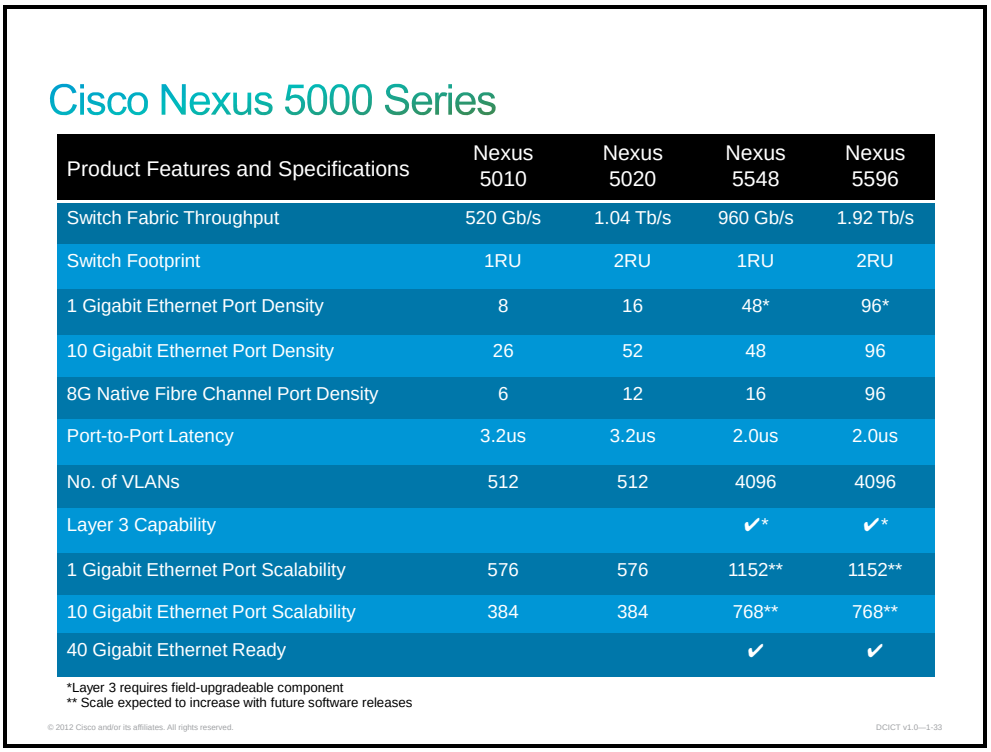
To launch the Cisco Power Calculator, go to the following URL and click the **Launch Cisco Power Calculator** link: http://tools.cisco.com/cpc/.

Note	You will need a Cisco.com account to launch the tool.

https://t.me/learningnets

Cisco Nexus 5000 Series Chassis Options

The Cisco Nexus 5000 Series Switches includes the Cisco Nexus 5000 Platform switch chassis and the Cisco Nexus 5500 Platform switch chassis. This topic compares the chassis for the Cisco Nexus 5000 and 5500 Platform switches.



The table in the figure describes the differences between the Cisco Nexus 5000 and 5500 Platform switches. The port counts are based on 24 Cisco Nexus 2000 Series Fabric Extenders per Cisco Nexus 5500 switch.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 5010 and 5020 Switches Features

The Cisco Nexus 5010 and 5020 switches were the first switches in the Cisco Nexus 5000 Series. These switches provide Layer 2 connectivity. This topic describes the features of the Cisco Nexus 5010 and 5020 switches.

Cisco Nexus 5010 Switch

- 28-port Layer 2 switch
  - OOB 10/100/1000-Mb/s management port
- 20 wire-speed 10 Gb/s, DCB- and FCoE-capable, SFP+ ports
  - 520-Gb/s throughput
- Single GEM slot



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

The Cisco Nexus 5000 Series Switches are a family of line-rate, low-latency, cost-effective 10 Gigabit Ethernet switches that are designed for access-layer applications.

The following are key features of the Cisco Nexus 5000 Series Switches:

- Wire-speed 10 Gigabit Ethernet ports supporting DCB and FCoE
- One expansion slot supporting any of the Generic Expansion Modules (GEM) for the Cisco Nexus 5000 Series Switches
- Redundant power entry connections
- Support for 1 Gigabit Ethernet on the first eight ports if required
- Low-latency switch

https://t.me/learningnets

Cisco Nexus 5020 Switch

- 56-port Layer 2 switch
  - OOB 10/100/1000-Mb/s management port
- 40 wire-speed 10 Gb/s, DCB- and FCoE-capable, SFP+ ports
  - 1.04-Tb/s throughput
- 2 GEM slots



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

The Cisco Nexus 5020 Switch is a 2-RU, 56-port, Layer 2 switch that provides an Ethernet-based unified fabric. It delivers 1.04 Tb/s of nonblocking switching capacity with 40 fixed wire-speed 10 Gigabit Ethernet ports that accept modules and cables meeting the SFP+ form factor. All of the 10 Gigabit Ethernet ports support DCB and FCoE.

The following are key features of the Cisco Nexus 5020 switch:

- Wire-speed 10 Gigabit Ethernet ports supporting DCB and FCoE
- Two expansion slots supporting any of the Generic Expansion Modules (GEM) for the Cisco Nexus 5000 switch
- Redundant power entry connections
- Support for 1 Gigabit Ethernet on the first 16 ports if required
- Low-latency switch

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 5000 Series Switches Front Panel

- Front-to-back airflow
- Power supplies and fans serviced from front
- N+1 redundancy for all front panel components



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

For Cisco Nexus 5000 Platform Switches (5010 and 5020), cooling is front-to-back, supporting hot- and cold-aisle configurations that help increase cooling efficiency. All serviceable components are accessible from the front panel, allowing the switch to be serviced while in operation and without disturbing the network cabling.

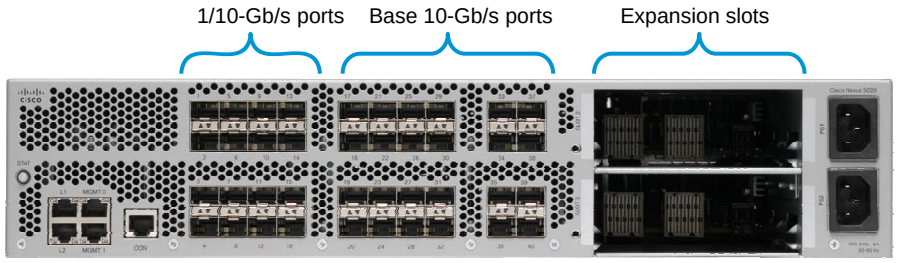
The following are key features of the front panel:

- Two N+1 redundant, hot-pluggable power supplies
- Five N+1 redundant hot-pluggable fan modules
- LED status indicators
- Slots for two power supplies

https://t.me/learningnets

Cisco Nexus Series 5000 Switches Rear Panel

- Interfaces are on the back, aligned in the rack with server ports
- Hot-swappable expansion modules
- 10/100/1000-Mb/s OOB management Ethernet port



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

The network and management interfaces are located on the rear panel of the Cisco Nexus 5000 Platform Switches (5010 and 5020), which are aligned in the rack with server ports. When the unit is rack-mounted, the interface connections on the front panel align with the server connections in the rack to allow easy cabling runs from server to switch.

All the Cisco Nexus 5000 Series Switches have a bank of four management ports. These ports include two internal cross-connect ports that are currently unused, the 10-, 100-, 1000-Mb/s OOB management Ethernet port and the console port.

The Cisco Nexus 5010 Switch has 20 fixed 10-Gigabit Ethernet ports for server or network connectivity. The first bank of ports is 1-Gigabit Ethernet-capable, while the remaining Ethernet ports are 10 Gigabit Ethernet. There is one slot for a hot-swappable, optional expansion module (GEM).

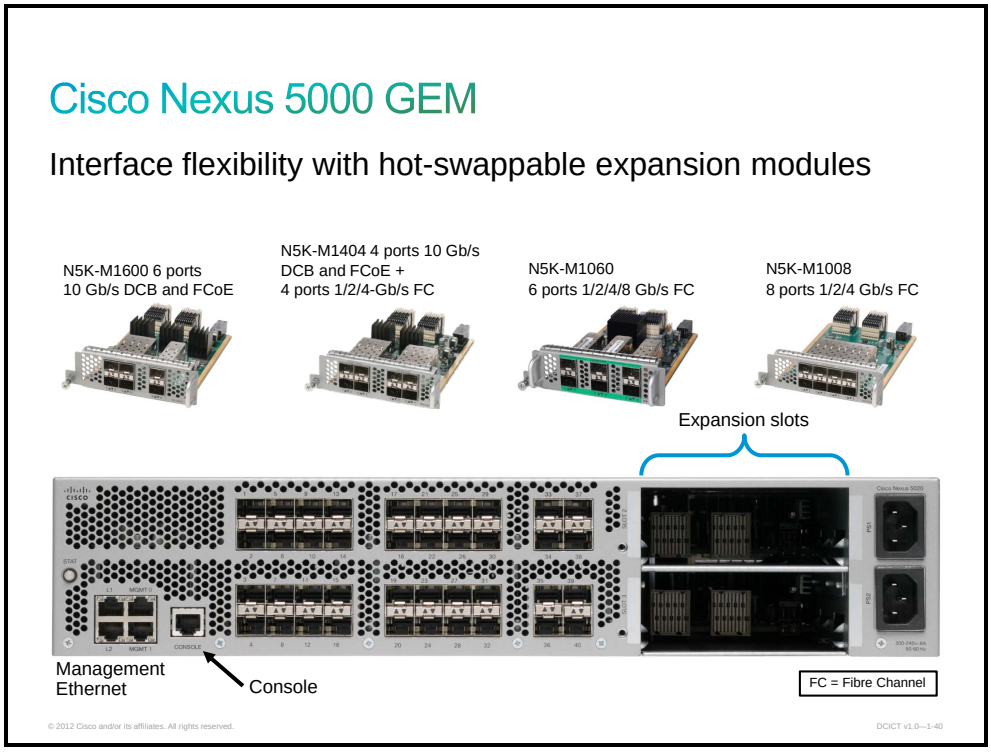
The Cisco Nexus 5020 Switch has 40 fixed 10-Gigabit Ethernet ports for server or network connectivity. The first bank of ports is 1-Gigabit Ethernet-capable, while the remaining Ethernet ports are 10 Gigabit Ethernet. There are two slots for hot-swappable, optional expansion modules (GEM).

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 5010 and 5020 Expansion Modules

Additional Ethernet or Fibre Channel ports can be added to the Cisco Nexus 5010 and 5020 Switches by using expansion modules. The Cisco Nexus 5010 switch has one expansion slot. The Cisco Nexus 5020 switch has two. This topic describes the expansion modules that are available for the Cisco Nexus 5010 and 5020 Switches.



Expansion modules allow Cisco Nexus 5000 Series Switches to be configured as cost-effective 10 Gigabit Ethernet switches and as I/O consolidation platforms with native Fibre Channel connectivity. There are currently four expansion module options that can be used to increase the number of 10 Gigabit Ethernet, FCoE-capable ports, connect to Fibre Channel SANs, or do both.

The Cisco Nexus 5010 Switch supports a single module, while the Cisco Nexus 5020 Switch supports any combination of the following two modules. Two expansion module slots can be configured to support up to 12 additional 10 Gigabit Ethernet, FCoE-capable ports. The modules may also be configured with up to 16 Fibre Channel ports, or a combination of both in a nonblocking, nonoversubscribed fashion. All modules are hot-swappable.

- **Eight-port 1/2/4-Gigabit Fibre Channel:** A Fibre Channel module that provides eight ports of 1-, 2-, or 4-Gb/s Fibre Channel through small form-factor pluggable (SFP) ports for transparent connectivity with existing Fibre Channel networks. This module is ideal in environments where storage I/O consolidation is the main focus.
- **Six-port 1/2/4/8-Gigabit Fibre Channel:** A Fibre Channel module that provides six ports of 1-, 2-, 4-, or 8-Gb/s Fibre Channel through SFP ports for higher speed or longer distances over Fibre Channel. Requires Cisco NX-OS Software Release 4.1(3)N2 or later release.
- **Four-port 10 Gigabit Ethernet (DCB and FCoE) and 4-port 1/2/4-Gigabit Fibre Channel:** A combination Fibre Channel and Ethernet module provides four 10 Gigabit Ethernet/, FCoE-capable ports through SFP+ ports and four ports of 1-, 2-, or 4-Gb/s native Fibre Channel connectivity through SFP ports.

https://t.me/learningnets

- **Six-port 10 Gigabit Ethernet (DCB and FCoE):** A 10 Gigabit Ethernet module provides an additional six 10 Gigabit Ethernet, FCoE-capable SFP+ ports per module, helping the switch support even denser server configurations.

Note	When you calculate port requirements for a design, note that the maximum number of 10-Gigabit Ethernet ports that are available on a Cisco Nexus 5010 Switch is 26 and 52 on a Cisco Nexus 5020 Switch.

Expansion modules are hot-swappable and contain no forwarding logic.

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

Cisco Nexus 5500 Platform Switches Features

The Cisco Nexus 5500 Platform switches were the second generation of Cisco Nexus 5000 Series Switches to be released. This topic describes the features of the Cisco Nexus 5548P, 5548UP, and 5596UP Switches.

Cisco Nexus 5548 Switches

- 5548P and 5548UP
- 48-port Layer 3-capable switch
  - Out-of-band 10/100/1000-Mb/s management port
- 32 wire-speed 1/10 Gb/s, DCB- and FCoE-capable ports
  - 960-Gb/s throughput
- 1 GEM2 expansion slot
- Layer 3 daughter card



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

The Cisco Nexus 5500 Platform switches are the second generation in the Cisco Nexus 5000 Series Switches, a series of line-rate, low-latency, cost-effective 10 Gigabit Ethernet switches. The initial release of the Cisco Nexus 5500 Platform switches was the Cisco Nexus 5548P Switch, which was followed by the release of the Cisco Nexus 5548UP Switch, supporting unified ports.

The Cisco Nexus 5500 Platform switches are well suited for enterprise-class data center server access-layer deployments and smaller-scale, midmarket data center aggregation deployments across a diverse set of physical, virtual, storage access, and unified data center environments. The Cisco Nexus 5500 Platform switches have the hardware capability to support Cisco FabricPath and IETF Transparent Interconnection of Lots of Links (TRILL) to build scalable and highly available Layer 2 networks.

The Cisco Nexus 5548 Switch is a 1-RU, 48-port Layer 3-capable switch that provides an Ethernet-based unified fabric. It delivers 960 Gb/s of nonblocking switching capacity with 32 fixed wire-speed 1 and 10 Gigabit Ethernet ports. All of the 10 Gigabit Ethernet ports support DCB and FCoE.

The switch has a single serial console port and a single OOB 10-, 100-, and 1000-Mb/s Ethernet management port.

There is a single expansion slot that supports any Generic Expansion M

odule 2 (GEM2) series module.

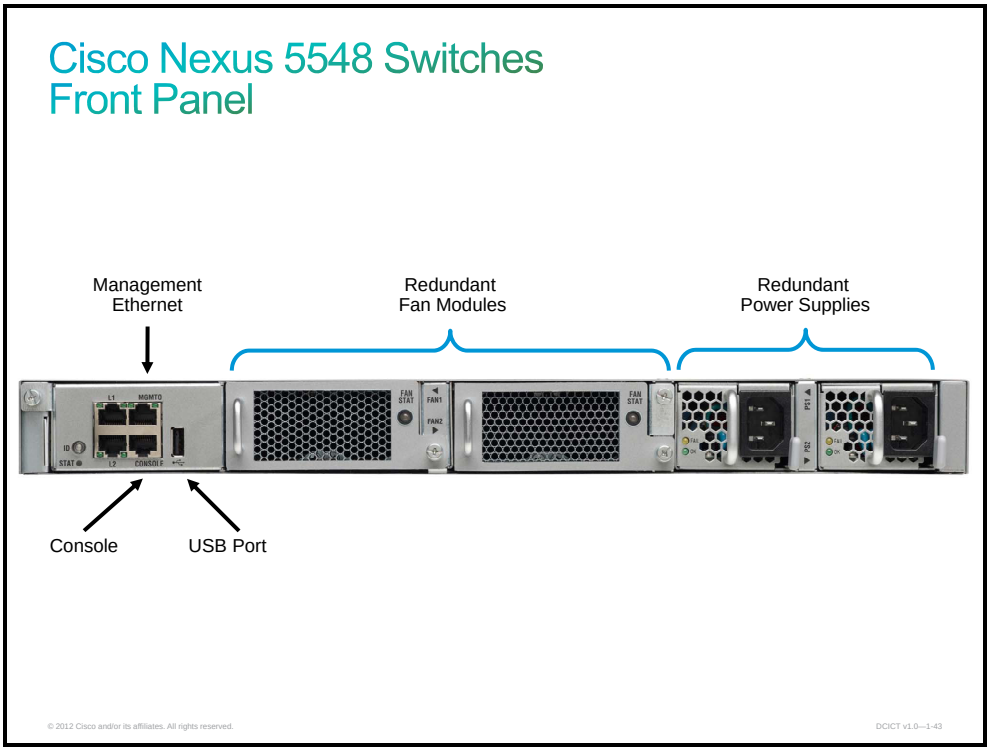
The Cisco Nexus 5500 Platform switches can be used as a Layer 3 switch through the addition of a Layer 3 daughter card routing module, enabling the deployment of Layer 3 services at the access layer.

https://t.me/learningnets

The Cisco Nexus 5500 Platform switches support Cisco FabricPath. Cisco FabricPath is a set of multipath Ethernet technologies that combine the reliability and scalability benefits of Layer 3 routing with the flexibility of Layer 2 networks, enabling IT to build massively scalable data centers. Cisco FabricPath offers a topology-based Layer 2 routing mechanism that provides an Equal-Cost Multipath (ECMP) forwarding model. Cisco FabricPath implements an enhancement that solves the MAC address table scalability problem that is characteristic of switched Layer 2 networks. Furthermore, Cisco FabricPath supports enhanced virtual port channel (vPC+), a technology that is similar to vPC that allows redundant interconnection of the existing Ethernet infrastructure to Cisco FabricPath without using Spanning Tree Protocol (STP).

The figure shows a Cisco Nexus 5548P Switch and a Cisco Nexus 5548UP Switch, each with a GEM2 expansion module installed.

The Cisco Nexus 5500 Platform switches start with the Cisco NX-OS Software Release 5.0(2)N1(1).



The front panel of the Cisco Nexus 5548 switches contains the management interfaces. There is a bank of four management ports. The ports include two internal cross-connect ports that are currently unused, the 10-, 100-, and 1000-Mb/s OOB management Ethernet port, and the console port. In addition, a fully functional USB port is present, useful for transferring files to or from the device bootflash, for backup, or other purposes.

Similar to the Cisco Nexus 5000 Platform switches, cooling is front-to-back, supporting hot- and cold-aisle configurations that help increase cooling efficiency. All serviceable components are accessible from the front panel, allowing the switch to be serviced while in operation and without disturbing the network cabling.

The Cisco Nexus 5548 Switch (P and UP) front panel includes two, N+1 redundant, hot-pluggable power supply modules and two, N+1 redundant, hot-pluggable fan modules for highly reliable front-to-back cooling.

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

Two power supplies can be used for redundancy, but the switch is fully functional with one power supply. The power supply has two LEDs, one for power status and one to indicate a failure condition.

Note	It is not recommended that you leave a power supply slot empty. If you remove a power supply, replace it with another one. If you do not have a replacement power supply, leave the nonfunctioning one in place until you can replace it.

The Cisco Nexus 5548 Switch (P and UP) requires two fan modules. Each fan module has four fans. If more than one fan fails in one of these modules, you must replace the module.

The fan module LED indicates the fan tray health. Green indicates normal operation, while amber indicates a fan failure.

Cisco Nexus 5596UP Switch

- 96 unified ports capable of 1 Gigabit Ethernet, 10 Gigabit Ethernet, FCoE, and Fibre Channel
  - OOB 10/100/1000-Mb/s management port
- 48 wire-speed 1/10 Gb/s, DCB- and FCoE-capable ports
  - 1.92-Tb/s throughput
- 3 GEM2 expansion slots
  - Layer 3 hot-swappable module



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

The Cisco Nexus 5596UP Switch is a 2-RU, 96-port Layer 3-capable switch that provides an Ethernet-based unified fabric. It delivers 1.92 Tb/s of nonblocking switching capacity with 48 fixed wire-speed 1 and 10 Gigabit Ethernet ports. All of the 10 Gigabit Ethernet ports are unified ports. In the future, 40 Gigabit Ethernet uplinks will be supported.

The switch has a single serial console port and a single OOB 10-, 100-, and 1000-Mb/s Ethernet management port.

There are three expansion slots that support any GEM2.

The Cisco Nexus 5500 Platform switches can be used as a Layer 3 switch through the addition of a Layer 3 hot-swappable GEM2 routing module. This addition enables the deployment of Layer 3 services at the access layer.

The Cisco Nexus 5500 Platform switches start with the Cisco NX-OS Software Release 5.0(2)N1(1).

https://t.me/learningnets

The Cisco Nexus 5500 Platform switches include the following features:

- 1 and 10 Gigabit Ethernet classical Ethernet, DCB-, and FCoE-capable switch
- 1-. 2-, 4-, and 8-Gigabit Fibre Channel switch, T11 FCoE
- 4096 VLANs (some are reserved by Cisco NX-OS in software) (507 on the Cisco Nexus 5010 or 5020)
- 32,000 MAC addresses (16,000 on Cisco Nexus 5010 or 5020)
- As many port channels as the number of ports permit
- Support for Layer 3 switching (with future daughter card or GEM module) connected
- 96 ports of line rate 10 Gigabit Ethernet (5596)
- Unicast latency: approximately 2 ms
- Quality of service (QoS) and multicast enhancements (differentiated services code point [DSCP] marking, more multicast queues, and so on)
- Switched Port Analyzer (SPAN) enhancements
- Hardware support for IEEE 1588 (Precision Time Protocol [PTP], microsecond accuracy, and time stamp)
- Flexible port configurations
- Unified ports on all ports of the Cisco Nexus 5596UP and 5548UP Switches and via an expansion module for the Cisco Nexus 5548P Switch
- Layer 2 and Layer 3 support
- Supports up to 16 fabric extenders (FEXs)
- 40 Gigabit Ethernet uplinks (future)
- 48 Cisco Nexus 5000 Series Switches port channels plus 384 Cisco Nexus 2000 Series Fabric Extenders port channels and 768 vPCs
- Support for Cisco FabricPath

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

Cisco Nexus 5500 Platform Switches Expansion Modules

Additional Ethernet or Fibre Channel ports can be added to the Cisco Nexus 5500 Platform switches by using expansion modules. The Cisco Nexus 5548 Switches have one expansion slot. The Cisco Nexus 5596 switch has three. This topic describes the expansion modules that are available for Cisco Nexus 5548P, 5548UP, and 5596UP Switches.

Cisco Nexus 5500 Platform Switches Ethernet and Fibre Channel GEM2

- N55-M8P8FP8 GEM2
- 8 SFP+ 1 and 10 Gigabit Ethernet ports
- 8 SFP ports 1/2/4/8 Gigabit Fibre Channel



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

The N55-M8P8FP GEM provides eight 1 or 10 Gigabit Ethernet and FCoE ports using the SFP+ interface. The module also provides eight ports of 8-, 4-, 2-, or 1-Gb/s native Fibre Channel connectivity using the SFP interface.

https://t.me/learningnets

Cisco Nexus 5500 Platform Switches Ethernet GEM2

- N55-M16P GEM2
- 16 SFP+ Ethernet ports, hardware capable of 1 and 10 Gigabit Ethernet



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

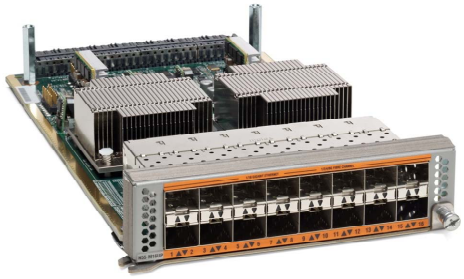
The N55-M16P GEM provides 16 1 and 10Gb/s ports using the SFP+ transceiver.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 5500 Platform Switches Unified Ports GEM2

- Unified ports can be configured as either Ethernet or native Fibre Channel ports
- N55-M16UP GEM2 provides 16 unified ports
  - Ethernet operations at 1 and 10 Gigabit Ethernet
  - Fibre Channel operations at 8/4/2/1 Gb/s



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

The N55-M16UP GEM provides 16 unified ports. A unified port is a single port that can be configured as either an Ethernet or a native Fibre Channel port. In Ethernet operation, it functions as a 1 or 10-Gb/s port, or in Fibre Channel operations as 8/4/2/1-Gb/s.

The N55-M16UP has the following characteristics:

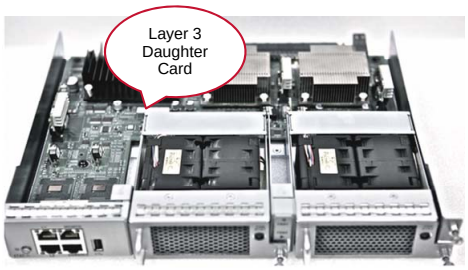
- 16 unified ports
- Ports can be configured as either Ethernet or native Fibre Channel ports
- Ethernet operation at 1 and 10 Gigabit Ethernet
- Fibre Channel operation at 8/4/2/1 G
- Uses existing Ethernet SFP+ and Cisco 8/4/2-Gb/s and 4/2/1–Gb/s Fibre Channel optics

The unified port expansion module may be installed in any of the Cisco Nexus 5500 Platform Switches chassis.

https://t.me/learningnets

Cisco Nexus 5548 Switch
Layer 3 Daughter Card

- Layer 3 daughter card for Cisco Nexus 5548 Switch
- N55-D160L3
  - Field replaceable module
  - In-rack upgradeability
  - No unmounting required

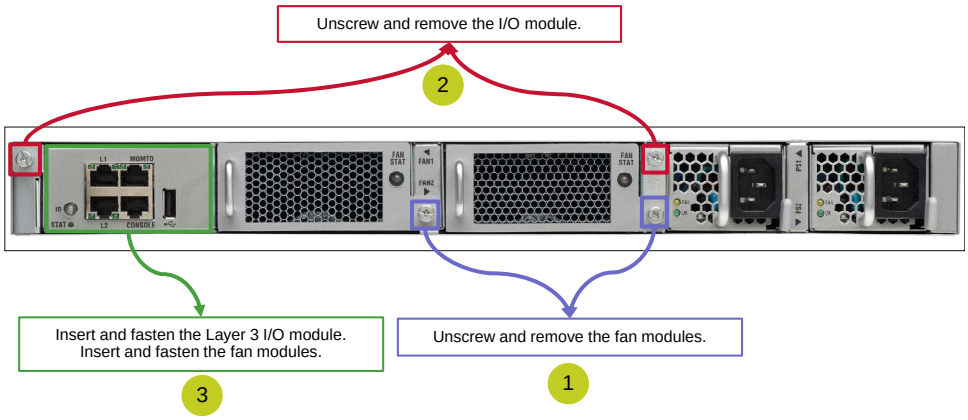


© 2012 Cisco and/or its affiliates. All rights reserved.

The N55-D160L3 is the Layer 3 daughter card for the Cisco Nexus 5548Switch. Layer 3 support is enabled via a field-upgradable routing card that may be installed while the switch remains mounted in the rack.

Cisco Nexus 5548 Switch
Layer 3 Daughter Card Installation

It is recommended to power down the switch before installing the Layer 3 daughter card.



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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

The N55-D160L3 is not a GEM. It is a field upgradable card that is a component of the management complex of the switch. The following steps are required to upgrade to Layer 3 services.

Step 1 Power down the switch.

Step 2 Unscrew and remove the fan modules.

Step 3 Unscrew and remove the I/O module (the management port complex).

Step 4 Insert and fasten the screws on the Layer 3 I/O daughter card.

Step 5 Replace and fasten the screws on the fan modules.

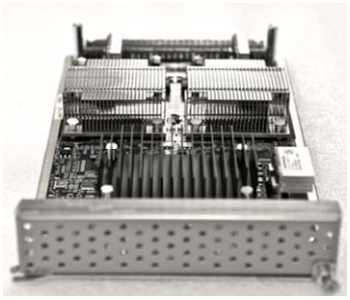
Step 6 Power on the switch.

Note If two fan modules are removed, a major alarm will be generated. The system starts a 120-second shutdown timer. If the module is reinserted within 120 seconds, the major alarm will be cleared and the shutdown timer will be stopped.

It is recommended to power down the switch before installing the Layer 3 module. This procedure affects service.

Cisco Nexus 5596 Switch Layer 3 Expansion Module

- Layer 3 hot-swappable module for Cisco Nexus 5596UP
- N55-M160L3
  - Field replaceable module
  - Provides 160 Gb/s of Layer 3 services



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

The Layer 3 GEM2 (N55-M160L3) provides 160 Gb/s of Layer 3 services to the Cisco Nexus 5596UP Switch. This expansion module is a field-replaceable unit (FRU).

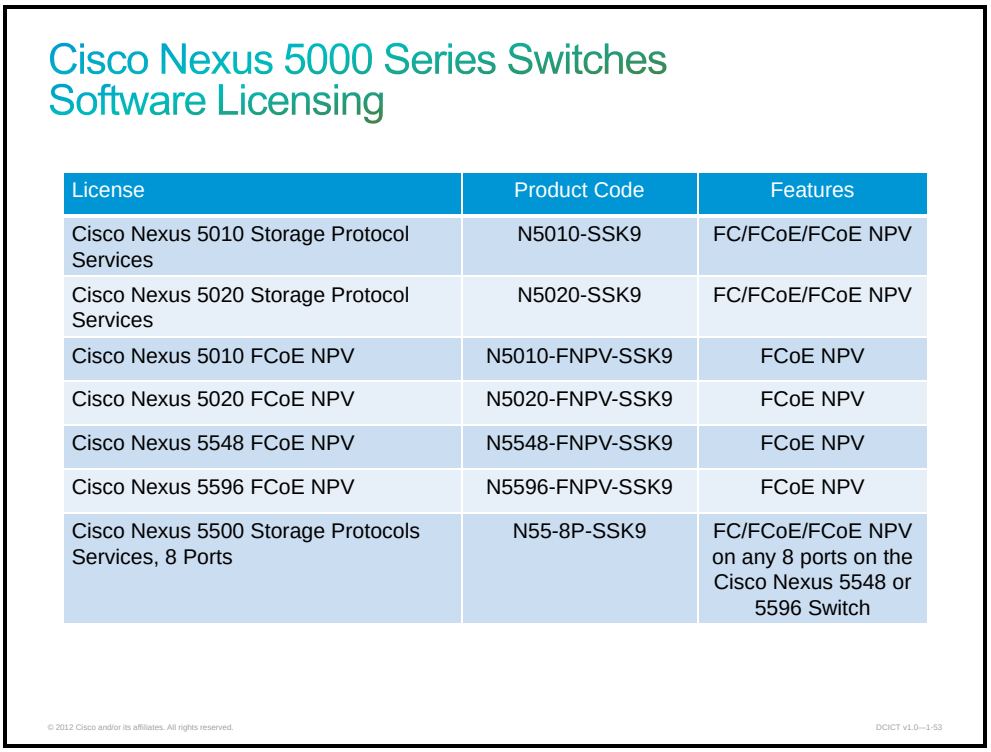
Each Layer 3 expansion module provides 160 Gb/s of Layer 3 services to the chassis, 10 Gb/s across any 16 ports that are as configured by the administrator. The Cisco Nexus 5596UP chassis has three available expansion slots; however, only one Layer 3 module is supported at the initial release.

Each Layer 3 module that is installed reduces the overall system scalability, primarily regarding the number of FEXs that are supported per switch. Up to 24 FEXs are supported per Cisco Nexus 5548P, 5548UP, 5596UP Switches, but this number is reduced to eight FEXs in Layer 3 configurations. This configuration correspondingly reduces the available server ports.

https://t.me/learningnets

Cisco Nexus 5000 Series Software Licensing

For ease of adding additional features, Cisco uses a software licensing model on the Cisco Nexus 5000 Series Switches. This topic describes licensing for the Cisco Nexus 5000 Series Switches.



Licensing for the Cisco Nexus 5000 Platform switches is tied to the physical chassis serial number. The Cisco Nexus 5500 Platform licensing is port-based. Both platforms allow a grace period or trial licenses to be used, sometimes referred to as honor-based licensing.

The following terminology is often used when describing Cisco NX-OS Software licensing:

- **Licensed feature:** Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- **Licensed application:** A software feature that requires a license to be used.
- **License enforcement:** A mechanism that prevents a feature from being used without first obtaining a license.
- **Node-locked license:** A license that can only be used on a particular switch using the unique host ID of the switch.
- **Host ID:** A unique chassis serial number that is specific to each switch.
- **Proof of purchase:** A document entitling its rightful owner to use licensed features on one switch as described in that document. The proof of purchase document is also known as the claim certificate.
- **Product Authorization Key (PAK):** The PAK allows you to obtain a license key from one of the sites that is listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through email.

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

- **License key file:** A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
- **Missing license:** If the bootflash has been corrupted or a supervisor module has been replaced after you have installed a license, that license shows as “missing.” The feature still works, but the license count is inaccurate. You should reinstall the license as soon as possible.
- **Incremental license:** An additional licensed feature that was not in the initial license file. License keys are incremental. If you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- **Evaluation license:** A temporary license. Evaluation licenses are time-bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- **Permanent license:** A license that is not time-bound is called a permanent license.
- **Grace period:** The amount of time the features in a license package can continue functioning without a license.

You can either obtain a factory-installed license (only applies to new switch orders) or perform a manual license installation of the license (applies to existing switches in your network).

The Base Services Package (N5000-AS) is included with the switch hardware at no additional charge. It includes all available Ethernet and system features, except features that are explicitly listed in the Basic Storage Services Package. The Basic Storage Services Package provides all Fibre Channel functionality. That functionality includes FCoE, DCB, Fibre Channel SAN services, Cisco N_Port Virtualizer (Cisco NPV), Fibre Channel Port Security, and Fabric Binding.

Layer 3 support on the Cisco Nexus 5500 Platform switches requires additional field upgradable hardware and a software license.

The Cisco Nexus 5000 Series Switches licensing is tied to the physical chassis serial number. The Storage Protocol Services Package license is installed once per chassis on the Cisco Nexus 5000 Series Switches.

The licenses that are supported on the Cisco Nexus 5000 Series Switches are described in the table in this figure and on the subsequent two figures.

https://t.me/learningnets

Cisco Nexus 5000 Series Switches
Software Licensing (Cont.)

License	Product Code	Features
Cisco Nexus 5500 Storage Protocols Services, 48 Ports	N55-48P-SSK9	FC/FCoE/FCoE NPV on 48 ports on the Cisco Nexus 5500 Platform switches
Cisco Nexus 5500 Layer 3 Base	N55-BAS1K9	Static routing, RIPv2, OSPF vs. EIGRP stub, HSRP, VRRP, IGMP v2 and v3, PIMv2 PIM-SM, routed ACL, uRPF, OSPF scalability limited to 256 dynamically learned routes
Cisco Nexus 5500 Layer 3 Enterprise	N55-LAN1K9	Full EIGRP, OSPF with scalability up to 8000 routes, BGP, VRF-lite (IP VPN), maximum routes supported by the Layer 3 hardware is 8000 entries

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

This table is a continuation of the available licenses for the Cisco Nexus 5000 Series Switches.

Cisco Nexus 5000 Series Switches
Software Licensing (Cont.)

License	Product Code	Features
Cisco Nexus 5500 VM-FEX	N55-VMFEXK9	VM-FEX support for the Cisco Nexus 5500 Platform switches
Cisco Nexus 5548 Enhanced Layer 2	N5548-EL2-SSK9	Cisco FabricPath on the Cisco Nexus 5548 Switch
Cisco Nexus 5596 Enhanced Layer 2	N5596-EL2-SSK9	Cisco FabricPath on the Cisco Nexus 5596 Switch
DCNM SAN	DCNM-SAN-N5K-K9	DCNM SAN Advanced Edition for the Cisco Nexus 5000 Platform switches
DCNM LAN	DCNM-L-NXACCK9	DCNM LAN Advanced Edition for the Cisco Nexus 3000 and 5000 Platform switches

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

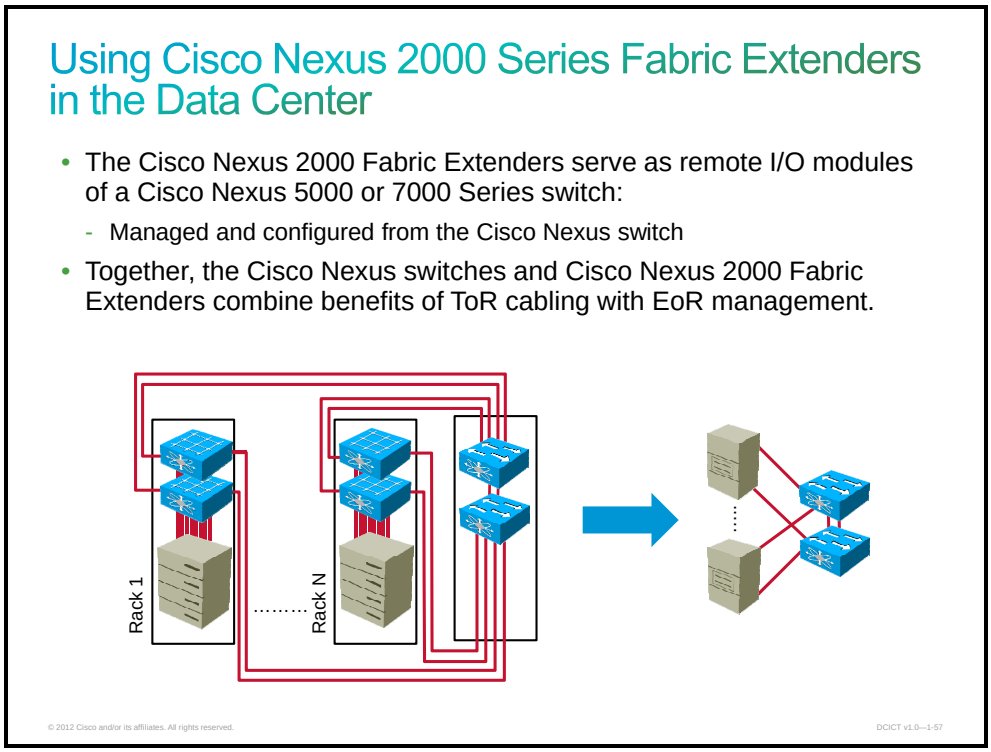
This table is a continuation of the available licenses for the Cisco Nexus 5000 Series Switches.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders Function in the Cisco Data Center

The Cisco Nexus 2000 Series Fabric Extenders are a remote I/O module for a Cisco Nexus 7000 or 5000 Series switch. By using the Cisco Nexus 2000 Series Fabric Extenders, many additional Ethernet ports can be added to the switches to create a virtualized switch chassis. This topic describes how the Cisco Nexus 2000 Series Fabric Extenders act as remote I/O modules in the data center.



Cisco Nexus 2000 Series Fabric Extenders can be deployed together with Cisco Nexus 5000 or Cisco Nexus 7000 Series Switches to create a data center network that combines the advantages of a top-of-rack (ToR) design with the advantages of an end-of-row (EoR) design.

Dual redundant Cisco Nexus 2000 Series Fabric Extenders are placed at the top of each rack. The uplink ports on the Cisco Nexus 2000 Series Fabric Extenders are connected to a Cisco Nexus 5000 or Cisco Nexus 7000 Series switch that is installed in the EoR position. From a cabling standpoint, this design is a ToR design. The cabling between the servers and the Cisco Nexus 2000 Series Fabric Extenders is contained within the rack. Only a limited number of cables need to be run between the racks to support the 10 Gigabit Ethernet connections between the Cisco Nexus 2000 Series Fabric Extenders and the Cisco Nexus switches in the EoR position.

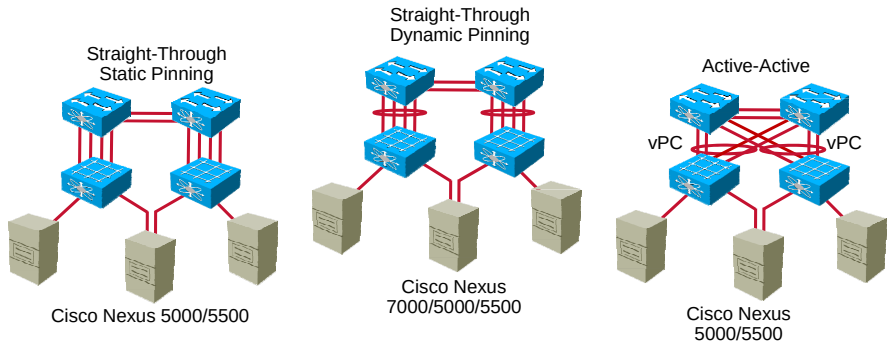
From a network deployment standpoint, however, this design is an EoR design. The fabric extenders (FEXs) act as remote I/O modules for the Cisco Nexus switches, which means that the ports on the Cisco Nexus 2000 Series Fabric Extenders act as ports on the associated switch. In the logical network topology, the FEXs disappear from the picture, and all servers appear as directly connected to the Cisco Nexus switch. From a network operations perspective, this design has the simplicity that is normally associated with EoR designs. All the configuration tasks for this type of data center design are performed on the EoR switches. There are no configuration or software maintenance tasks that are associated with the FEXs

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders Deployment Models for Native Ethernet Links

The Cisco Nexus 2000 Series Fabric Extenders can be deployed using three different models:

- Straight-through FEX using static pinning
- Straight-through FEX using dynamic pinning
- Active-active FEX using vPC



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

There are three deployment models that are used to deploy Cisco Nexus 2000 Series Fabric Extenders together with the Cisco Nexus 5000 and Cisco Nexus 7000 Series Switches:

- **Straight-through, using static pinning:** In the straight-through model each FEX is connected to a single Cisco Nexus switch. The single switch that the FEX is connected to exclusively manages the ports on that FEX. Static pinning means that each downlink server port on the FEX is statically pinned to one of the uplinks 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 the FEXs and the 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 that are based on the PortChannel hashing mechanism.
- **Active-active FEX using 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 that are based on vPC forwarding mechanisms.

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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders
Parent Switches

Parent Switch	FEX Supported	No. FEX	Optics/Transceivers Supported for FEX
Cisco Nexus 5010 Cisco Nexus 5020	All models	12	Passive CX-1 SFP+ (1/3/5m) Active CX-1 SFP+ (7/10m) SR SFP+ (MMF) – OM3 300m LR SFP+ (SMF) – 300m (FCoE) FET SFP+ (MMF) – OM2 20m, OM3 100m LRM SFP+ - Not Supported
Cisco Nexus 5548 Cisco Nexus 5596	All models	24	Passive CX-1 SFP+ (1/3/5m) Active CX-1 SFP+ (7/10m) SR SFP+ (MMF) – OM3 300m LR SFP+ (SMF) – 300m (FCoE) FET SFP+ (MMF) – OM2 20m, OM3 100m LRM SFP+ - Not Supported
N7K-M132XP-12 N7K-M132XP-12L N7K-F248XP-25	2224TP 2248TP 2232PP	32	Passive CX-1 SFP+ (1/3/5m) - on M132XP-12 Active CX-1 SFP+ (7m/10m) SR SFP+ (MMF): OM1 26m OM3 300m LR SFP+ (SMF): up to 10km FET SFP+ (MMF): OM2 25m OM3 100m LRM SFP+: Not Supported

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

The table shows the model and number of the Cisco Nexus 2000 Series Fabric Extenders that are supported by each parent switch model.

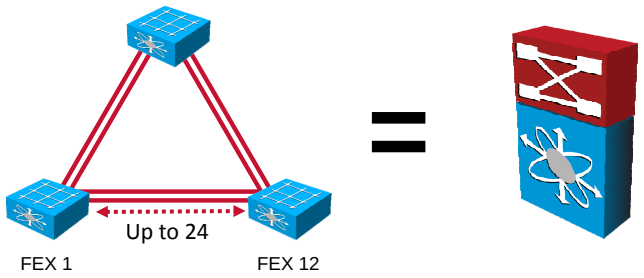
All Cisco products that support passive Twinax cables must support up to a 16 foot (5 m) distance. Beyond that length, Cisco supports only active Twinax cables.

Software support for active 23- and 33-foot (7- and 10-m) Twinax cables came with Cisco NX-OS Software Release 4.2(1).

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders Virtualized Switch Chassis

- The parent switch and the connected FEXs form a virtual switch chassis.
- Different models of FEXs can be connected to the same parent switch.



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

The figure explains the virtualized switch chassis from both the physical and logical view.

The Cisco Nexus 7000 or 5000 Series Switches and the Cisco Nexus 2000 Series Fabric Extenders that are connected to them combine to form a scalable virtual modular system, also called a virtualized switch chassis.

Different models of Cisco Nexus 2000 Series Fabric Extenders can be connected to the same parent switch. This connection type is similar to a physical modular chassis that may have physical line cards of different types, which are located in different slots.

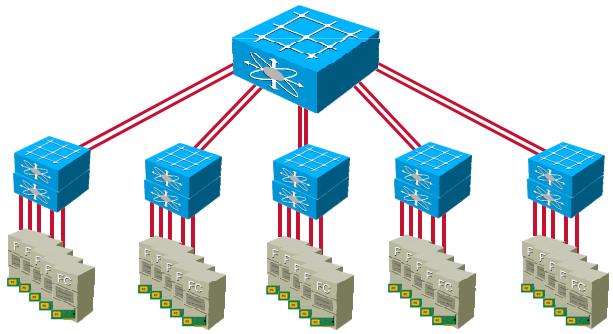
Physically, the parent switch is a separate device that is connected via uplink ports (fabric extensions) down to each FEX. Logically, the FEX is connected into the parent switch as a module. The fabric ports appear to the FEX the same way the universal feature card ASIC would connect to a GEM via the unified port control.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders Virtualized Switch Chassis (Cont.)

A single virtualized switch chassis with 1 or 2 expansion modules and up to 24 or 32 I/O modules



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

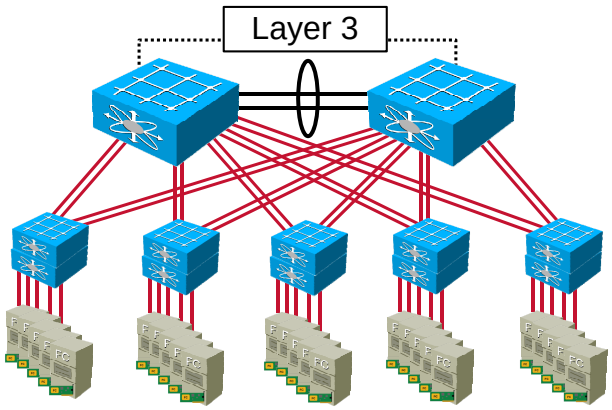
The scalable virtual modular system, also called a virtualized switch chassis, can contain up to 32 virtual I/O modules depending on the Cisco Nexus chassis that is being used.

Typically the FEX would be at the top of the rack for easy access to server connections and reduced cabling runs. The fabric ports would be connected back to the parent switch that is located at the middle or end of the row.

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders
Virtualized Switch Chassis High Availability

High availability can be achieved by linking two parent switches.



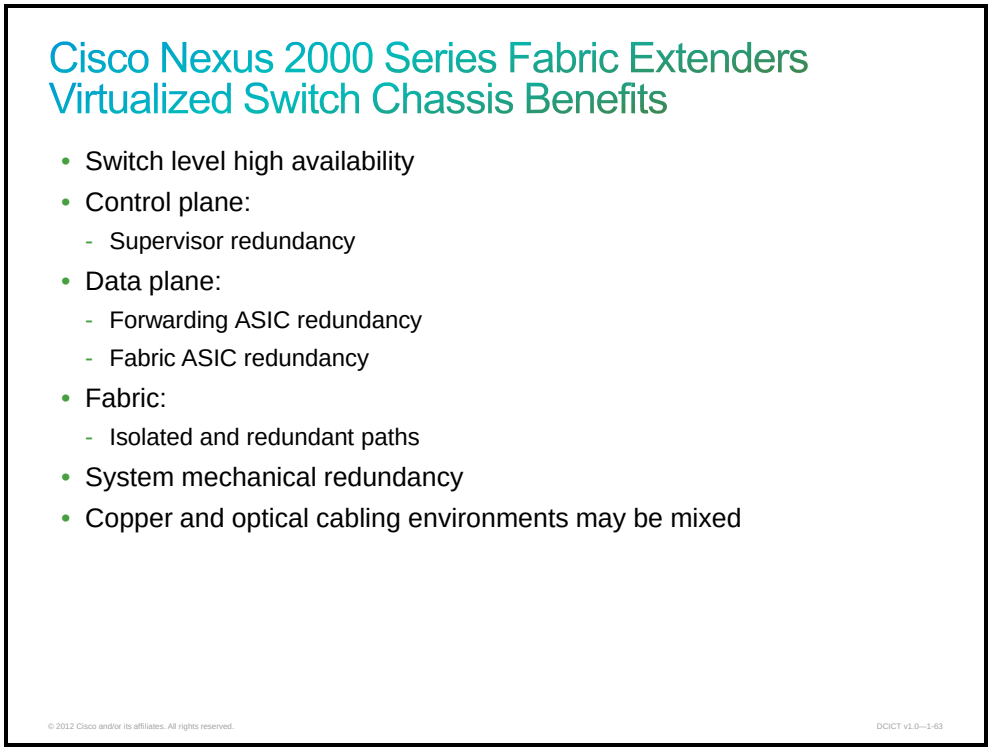
© 2012 Cisco and/or its affiliates. All rights reserved.

High availability can be achieved by linking two parent switches. This configuration extends the concept of the virtualized switch chassis by logically linking two parent switches and becomes a dual supervisor configuration with redundancy across all elements.

When two parent switches are linked, they logically combine to form one virtual switch. The number of FEXs that are supported by the single virtual switch is the same as if there was one single physical switch present.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



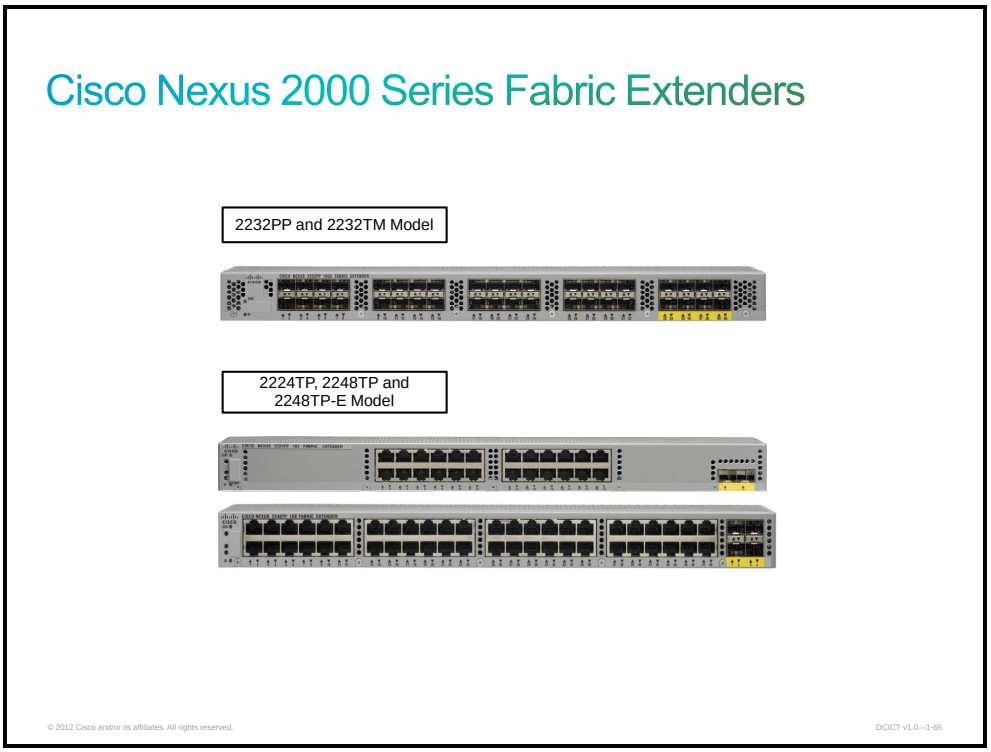
High switch level high availability addresses the following features:

- Control plane:
  - Supervisor redundancy
- Data plane:
  - Forwarding ASIC redundancy
  - Fabric ASIC redundancy
- Fabric:
  - Isolated and redundant paths
- System mechanical redundancy:
  - Power supply
  - Fan

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders Features

There are a number of Cisco Nexus 2000 Series Fabric Extenders that are available. This topic describes the features and range of the Cisco Nexus 2000 Series Fabric Extenders.



The Cisco Nexus 22XX (2248TP GE, 2248TP-E, 2232PP 10GE, 2232TM 10GE, and 2224TP GE) Fabric Extenders are stackable 1-RU chassis that are designed for rack mounting as a ToR solution. The second generation of FEXs uses a new ASIC, with additional features. The features include enhanced QoS with eight hardware queues, port channels, access control list (ACL) classification, and SPAN sessions on host ports. In addition, host ports can be configured as part of a vPC.

The fabric interfaces are fixed and reserved, marked in yellow on the chassis, as shown in the figure.

The following features should be noted for the Cisco Nexus 2232 Fabric Extender:

- The Cisco Nexus 2232PP 10GE Fabric Extender supports SFP and SFP+ connectivity on all ports.
- The Cisco Nexus 2232TM 10GE Fabric Extender supports RJ-45 connectivity with eight SFP+ uplinks. (There is no FCoE support on this model.)
- The Cisco Nexus 2232PP 10GE Fabric Extender can be connected to the Cisco Nexus 5000 or 7000 Series Switches.
- The Cisco Nexus 2232TM 10GEFabric Extender can be connected to the Cisco Nexus 5000 Series switch
- The Cisco Nexus 2232TM 10GE Fabric Extender requires Cisco NX-OS Release 5.0(2) N2(1) for connectivity to the Cisco Nexus 5000 Series Switches.

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

Additional features of the Cisco Nexus 2248TP-E Fabric Extender are listed here:

- Support for 48 100/1000BASE-T host-facing ports and four 10 Gigabit Ethernet fabric interfaces
- Support for 32-MB shared buffers
- Supported on the Cisco Nexus 5000 Series Switches only

https://t.me/learningnets

Cisco Nexus B22HP Series Blade Fabric Extender for HP

- 16 x 10 Gigabit Ethernet internal host interfaces
- 8 x 10 Gigabit fabric interfaces (SFP+)
- Cisco Nexus 5000 Series parent switch
- Supported chassis:
  - HP BladeSystem c3000 enclosure
  - HP BladeSystem c7000 enclosure



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

The Cisco Nexus B22 Series Blade Fabric Extenders are designed to simplify data center server access architecture and operations in environments in which third-party blade servers are used.

The Cisco Nexus B22 Series Blade Fabric Extenders behave like remote line cards for a parent Cisco Nexus switch, together forming a distributed modular system. This architecture simplifies data center access operations and architecture. The architecture combines the management simplicity of a single high-density access switch with the cabling simplicity of integrated blade switches, and ToR access switches.

The Cisco Nexus B22 Series Blade Fabric Extenders provide the following benefits:

- **Highly scalable, consistent server access**: Distributed modular system creates a scalable server access environment with no reliance on STP, providing consistency between blade and rack servers.
- **Simplified operations**: One single point of management and policy enforcement using upstream Cisco Nexus switches eases the commissioning and decommissioning of blades through zero-touch installation and automatic configuration of FEXs.
- **Increased business benefits**: Consolidation, cabling reduction, 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 all contribute to reduced operating expenses (OpEx) and capital expenditures (CapEx).

Each member of the Cisco Nexus B22 Series Blade Fabric Extenders transparently integrates into the I/O module slot of a third-party blade chassis, drawing both power and cooling from the blade chassis itself.

The Cisco Nexus B22 Series Blade Fabric Extenders provide two types of ports: ports for blade server attachment (host interfaces) and uplink ports (fabric interfaces). Fabric interfaces, which are located on the front of the Cisco Nexus B22 Series Blade Fabric Extender module, are for connectivity to the upstream parent Cisco Nexus switch. The figure is a picture of the Cisco Nexus B22 Blade Fabric Extender for HP, showing its eight fabric interfaces.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 2000 Series Fabric Extenders
Summary

Features	2148T	2224TP	2248TP/TP-E	2232PP/TM-10GE
Fabric ports	4 X 10G SFP+	2 X 10G SFP+	4 X 10G SFP+	8 X 10G SFP+
Fabric port channels	1 X 4 ports max L2/L3 hash	1 X 2 ports max L2/L3/L4 hash	1 X 4 ports max L2/L3/L4 hash	1 X 8 ports max L2/L3/L4 hash
Host ports	48 X 1 Gb/s TP	24 X 100 Mb/s or 1 Gb/s TP	48 X 100 Mb/s or 1 Gb/s TP	32 X SFP/SFP+
Host port channels (ports per channel)	Not supported	24 maximum (8)	24 maximum (8)	16 maximum (8)
Max. FEX ports per 5000 chassis	12 X 48 = 576	12 X 24= 288	12 X 48 = 576	12 X 32 = 384
Max. FEX ports per 5500 chassis (L2)	24 X 48 = 1152	24 X 24 = 576	24 X 48 = 1152	24 X 32 = 768
Max. FEX ports per 5500 chassis (L3)	8 X 48 = 384	8 X 24 = 192	8 X 48 = 384	8 X 32 = 256
FCoE/DCB	No	No	No	Yes
Supports FET	No	Yes	Yes	Yes
Dimensions (in)	1 RU 1.72 X 17.3 X 20	1 RU 1.72 X 17.3 X 17.7	1 RU 1.72 X 17.3 X 17.7	1 RU 1.72 X 17.3 X 17.7
Power	165 W	95 W	110 W	270 W

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

The table compares the features of Cisco Nexus 2000 Series Fabric Extenders models.

The primary differences in these features lie in these three factors:

- Scalability and oversubscription of the internal architecture
- Number and type of host ports to fabric ports
- The ability to form host port channels and support the DCB protocol

https://t.me/learningnets

Summary

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

Summary

- The Cisco Nexus Family of products has been designed specifically for the data center and has innovative features to support the data center requirements.
- There are three chassis models in the Cisco Nexus 7000 Series that are designed to scale from 8.8- to 18.7-Tb/s throughput, providing the ability to support large-scale deployments.
- The Cisco Nexus 7000 Series Supervisor Module supports all three models and has a main controller board and CMP to help ensure continual management connectivity even during maintenance tasks.
- To provide flexibility, the Cisco Nexus 7000 Series uses a licensing methodology for features. This model enables customers to upgrade to additional features without having to install additional software.
- Each of the Cisco Nexus 7000 Series models supports up to five fabric modules. The Fabric 1 module provides up to 230-Gb/s per slot while the Fabric 2 module provides up to 550-Gb/s per slot if all five fabric modules are installed.

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

Summary (Cont.)

- The Cisco Nexus 7000 Series is a modular product supporting various I/O modules. These I/O modules can provide 1, 10, 40, or 100 Gb/s per port depending on the module that is selected.
- The Cisco Nexus 7000 Series supports four power redundancy modes.
- The Cisco Nexus 5000 Series Switches, 5000 Platform and 5500 Platform, provide 1 or 10 Gb/s per port and can be used as ToR, middle-of-row or EoR connectivity.
- The Cisco Nexus 5010 and 5020 Switches were the first generation of switches in this product range. These switches support Layer 2 connectivity.
- There are several expansion modules for the Cisco Nexus 5010 and 5020 Switches, with the Cisco Nexus 5010 supporting one expansion module and the Cisco Nexus 5020 supporting two.
- The Cisco Nexus 5500 Platform switches are the second generation switches in this product range. These switches support Layer 2 natively, and Layer 3 with the addition of a daughter card or expansion module.

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

https://t.me/learningnets

Summary (Cont.)

- There are several expansion modules available for the Cisco Nexus 5500 Platform switches, with the Cisco Nexus 5548 switch supporting one expansion module and the Cisco Nexus 5596 switch supporting up to three.
- Similar to the Cisco Nexus 7000 Series, the Cisco Nexus 5000 Series Switches support a licensing model for additional features.
- The Cisco Nexus 2000 Series Fabric Extenders provide additional 1 or 10 Gb/s ports in the form of an external I/O module for either the Cisco Nexus 5000 or 7000 Series Switches. All management and switching is provided by the parent switch, allowing customers to increase port counts without the increased management overheads.
- The Cisco Nexus 2000 Fabric Extenders come in the form of 1 Gb/s modules or 10 Gb/s modules.

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

https://t.me/learningnets

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

Lesson 3

Reviewing the Cisco MDS Product Family

Overview

The Cisco Multilayer Director Switch (Cisco MDS) Fibre Channel switches are used in the SAN environment to provide Fibre Channel network connectivity for devices. This lesson looks at their features and how those features would be used in the layered design model.

Objectives

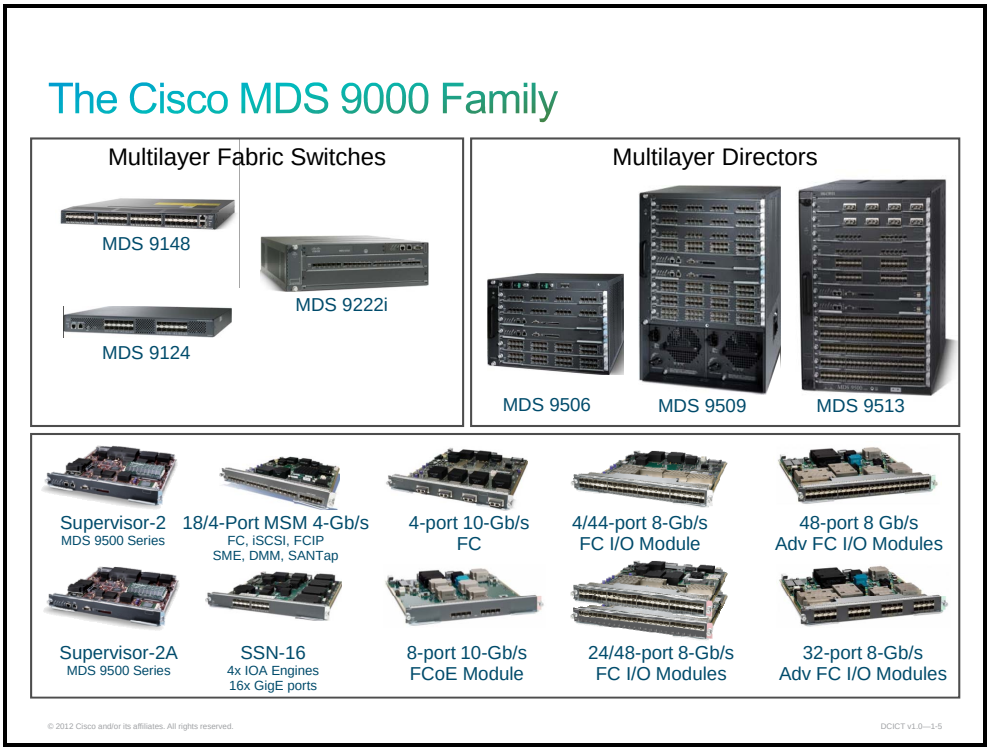
Upon completing this lesson, you will be able to describe the features of Cisco MDS Fibre Channel switches and their relationship to the layered design model. You will be able to meet these objectives:

- Describe the benefits to the data center of deploying the Cisco MDS 9000 Series Multilayer Switches
- Describe the capacities of the Cisco MDS 9500 Series Multilayer Directors chassis
- Describe the capabilities of the Cisco MDS 9500 Series supervisor modules
- Describe licensing options for the Cisco MDS 9000 Series Multilayer Switches
- Describe the capabilities of the Cisco MDS 9000 Series switching modules
- Describe the capabilities of the Cisco MDS 9500 Series power supply options
- Describe the capabilities of the Cisco MDS 9100 Series Multilayer Fabric Switches
- Describe the capabilities of the Cisco MDS 9222i Multiservice Modular Switch

https://t.me/learningnets

Cisco MDS 9000 Series Product Suite

The Cisco MDS 9000 Series includes both fabric and director class switches. The Cisco MDS 9000 Series is designed for SAN deployments in the data center, with a rich set of features to support the SAN requirements. This topic describes how data centers benefit from deployment of Cisco MDS products.



Multilayer switches are switching platforms with multiple layers of intelligent features, such as the following:

- Ultrahigh availability
- Scalable architecture
- Comprehensive security features
- Ease of management
- Advanced diagnostics and troubleshooting capabilities
- Seamless integration of multiple technologies
- Multiprotocol support

The Cisco MDS 9000 Family offers an industry-leading investment protection across a comprehensive product line, offering a scalable architecture with highly available hardware and software. The product line is based on the Cisco MDS 9000 Family operating system, and has a comprehensive management platform in Cisco Fabric Manager. The Cisco MDS 9000 Family offers various application I/O modules and a scalable architecture from an entry-level fabric switch to a director-class system.

The product architecture is forward- and backward-compatible with I/O modules, offering 1-, 2-, 4-, 8-, and 10-Gb/s Fibre Channel connectivity.

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

The Cisco MDS 9000 10-Gbps 8-Port FCoE Module provides up to 88 ports of line rate performance per chassis for converged network environments. The Fibre Channel over Ethernet (FCoE) module provides features that bridge the gap between the traditional Fibre Channel SAN and the evolution to an FCoE network implementation.

The Cisco MDS 9222i Multiservice Modular Switch uses the 18/4 architecture of the Cisco MDS 9000 Family 18/4-Port Multiservice Module (DS-9304-18K9) line card. In addition, it includes native support for Cisco MDS Storage Media Encryption (Cisco MDS SME).

The Cisco MDS 9148 Multiyear Fabric Switch is a new 8-Gb/s Fibre Channel switch providing 48 2-, 4-, and 8-Gb/s Fibre Channel ports. The base license supports 16-, 32-, and 48-port models but can be expanded using the 8-port license.

The Cisco MDS 9513 Multilayer Director chassis only uses the Cisco MDS 9500 Series Supervisor-2 Module or later. However, the initial Cisco MDS 9500 Series Supervisor-1 Module can be used in the Cisco MDS 9506 Multilayer Director and Cisco MDS 9509 Multilayer Director. All Cisco MDS 9500 Series chassis support all supervisor modules.

https://t.me/learningnets

Cisco MDS 9500 Series Chassis Options

Three chassis models are available in the Cisco MDS 9500 Series. This topic describes the capacities of the chassis for the Cisco MDS 9506, 9509, and 9513 Multilayer Directors.

Cisco MDS 9500 Series Multilayer Directors

- Nonblocking architecture:
  - Dual-redundant crossbar fabric
  - VOQ—resolve HOL blocking
- High bandwidth:
  - Up to 2.2-Tb/s aggregate internal bandwidth
  - Up to 160-Gb/s switch to switch—16 ISLs per port channel
- Low latency:
  - Less than 20 microsec per hop—link-rate-dependent
- Multiprotocol
  - Fibre Channel, FCoE, FICON, FCIP, iSCSI
- Scalable:
  - Store-and-forward architecture
  - Multiple VSAN support
- Secure:
  - Secure management access—SNMPv3 and RBAC
  - Port security—port and fabric binding
  - Data security—in transit and at rest
- Highly available:
  - Redundant fabric modules, supervisors, clocks, fans, power supplies, internal paths



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

Nonblocking architecture:

- Dual-redundant crossbar switch fabric
- Virtual output queues (VOQs)—resolve head-of-line (HOL) blocking

High bandwidth:

- Up to 2.2-Tb/s aggregate internal bandwidth on the Cisco MDS 9513 Multilayer Director switch
- Up to 160-Gb/s switch-to-switch—16 interswitch links (ISLs) per port channel

Low latency:

- Less than 20 microsec per hop—link-rate-dependent

Multiprotocol:

- Fibre Channel, FCoE, fiber connectivity (FICON), Fibre Channel over IP (FCIP), and Internet Small Computer Systems Interface (iSCSI)

Scalable:

- Store-and-forward architecture
- Up to 239 switches per fabric; up to 1000 Virtual Storage Area Networks (VSANs) per switch

 © 2012 Cisco Systems, Inc.

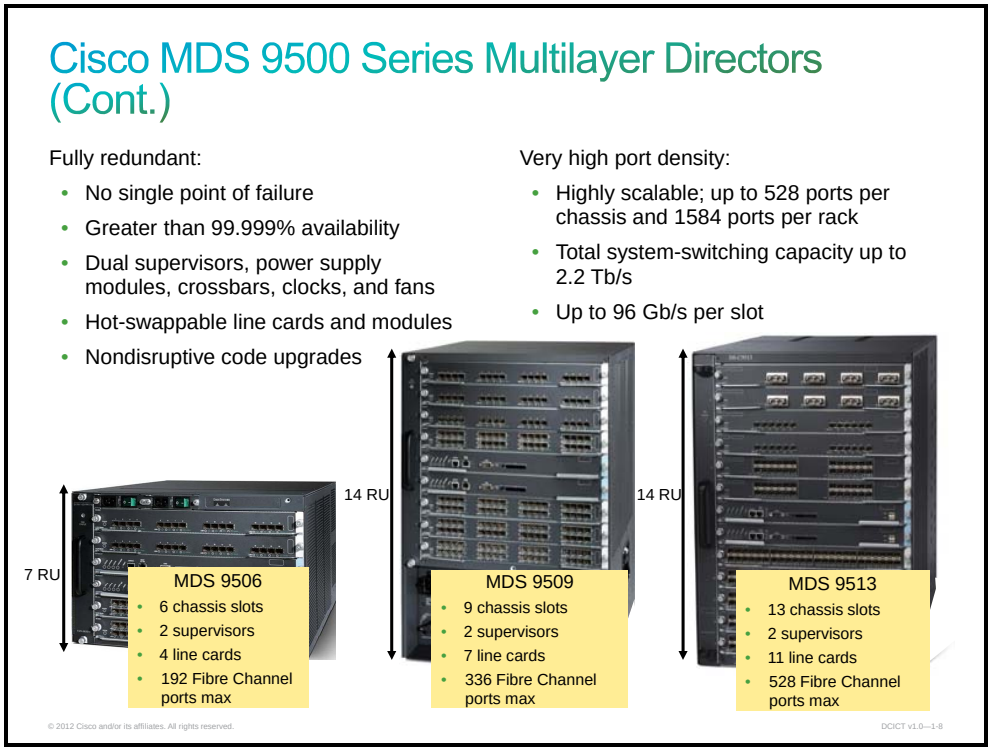
https://t.me/learningnets

Secure:

- Secure management access—Simple Network Management Protocol version 3 (SNMPv3) and role-based access control (RBAC)
- Port Security—port and fabric binding
- Data Security—In transit with Fibre Channel link encryption and IPsec
- Data Security—At rest with Cisco MDS SME

Highly available:

- Redundant fabric modules, supervisors, clocks, fans, power supply modules, internal paths



The Cisco MDS 9500 Series Multilayer Directors are enterprise-class multilayer director switches. They provide high availability, multiprotocol support, advanced scalability, security, nonblocking fabrics that are 10-Gb/s-ready, and a platform for storage management. The Cisco MDS 9500 Series allows you to deploy high-performance SANs with a lower total cost of ownership.

The Cisco MDS 9500 Series Multilayer Director Switches have a rich set of intelligent features and hardware-based services.

The chassis for the Cisco MDS 9500 Series switches are available in three sizes: Cisco MDS 9513, 14 rack units (14 RU), Cisco MDS 9509 14 RU, and Cisco MDS 9506 7 RU.

https://t.me/learningnets

Cisco MDS 9506 Chassis

The Cisco MDS 9506 Multilayer Director is a six-slot chassis and has slots for two supervisor modules and four switching or services modules. The power supply modules are located in the rear of the chassis, with the power entry modules (PEMs) in the front of the chassis for easy access.

Up to six Cisco MDS 9506 chassis can be installed in a standard 42-RU rack with up to 128 Fibre Channel ports per chassis. The maximum number of available Fibre Channel ports per rack is 768 in a single 7-foot (42 RU) rack, thus optimizing the use of valuable data center floor space. Additionally, cable management is facilitated by the single side position of both interface and power terminations.

Cisco MDS 9509 Chassis

The Cisco MDS 9509 Multilayer Director is a nine-slot chassis with redundant supervisor modules, up to seven switching or services modules, redundant power supply modules, and a removable fan module. Slots 5 and 6 are reserved for redundant supervisor modules, which provide control and switching capabilities, and local and remote management.

The Cisco MDS 9509 Multilayer Director supports up to 224 Fibre Channel ports in a single chassis.

There are two system-clock cards for added high availability. Dual redundant power supply modules are located at the front of the chassis; therefore, the Cisco MDS 9509 switch is only 18.8 inches deep.

Cisco MDS 9513 Chassis

The Cisco MDS 9513 Multilayer Director is a 13-slot chassis with redundant Supervisor-2 or Supervisor-2A Modules and up to 11 switching or services modules. It has redundant 6-KW power supply modules, a removable fan module at the front, and additional removable fan modules at the rear for the fabric modules. Slots 7 and 8 are reserved for redundant Supervisor-2 or -2A modules, which provide control and switching capabilities, and local and remote management.

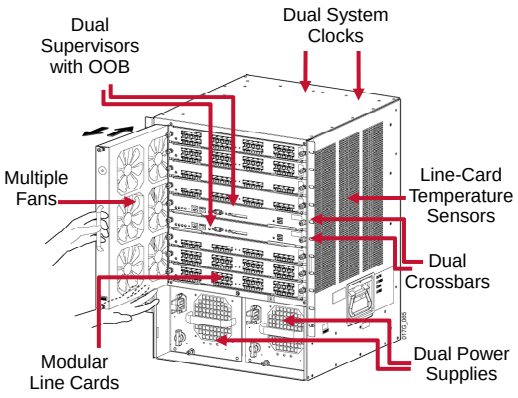
The Cisco MDS 9513 Multilayer Director supports up to 528 Fibre Channel ports in a single chassis. There are two new removable system-clock modules at the rear for added high availability. Dual-redundant 6-KW power supply modules are located at the rear of the chassis. The Cisco MDS 9513 Multilayer Director has a revised airflow system at the rear of the chassis: in at the bottom and out at the top.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

High Availability: Hardware

- Dual power supply modules
- Dual supervisors with OOB management
- Dual fabric crossbars
- Dual system clocks
- Multiple fans
- Hot-swappable modules
- Environmental monitoring



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

The Cisco MDS 9000 switch high availability goal is to exceed five 9s, or 99.999 percent, of uptime per year. This level of availability is equal to only 5 minutes of downtime per year and requires physical hardware redundancy, software availability, and logical network availability.

Hardware availability through redundancies and monitoring functions that are built into a Cisco MDS 9000 Series switch system include dual power supply modules, dual supervisors with out-of-band (OOB) management channels, dual fabric crossbars, dual system clocks, hot-swappable modules with power and cooling management, and environmental monitoring.

These necessary features provide a solid, reliable, director-class platform that is designed for mission-critical applications.

https://t.me/learningnets

Cisco MDS 9500 Series Supervisor Modules

The Cisco MDS 9500 Series switches are modular, and each model has two slots that are available for supervisor engines. This topic describes the capabilities of the Cisco MDS 9500 Series supervisor modules.

Supervisor-2 and -2A Module Features

- High-performance, integrated crossbar
- Enhanced crossbar arbiter
- PowerPC management processor
- Cisco MDS 9513 Multilayer Director requires Supervisor-2
- Supervisor-2A introduces FCoE support



Front Panel Interfaces
• Console Port
• Management Ethernet Port • 10/100/1000
• COM1 Port
• Compact Flash Slot
• USB Ports (2)

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

The Cisco MDS 9500 Series Supervisor-2 Module is an upgraded version of the Cisco MDS 9500 Series Supervisor-1 Module with additional flash memory, RAM and NVRAM memory, and redundant BIOS. It can be used in any Cisco MDS 9500 Series Multilayer Directors switch.

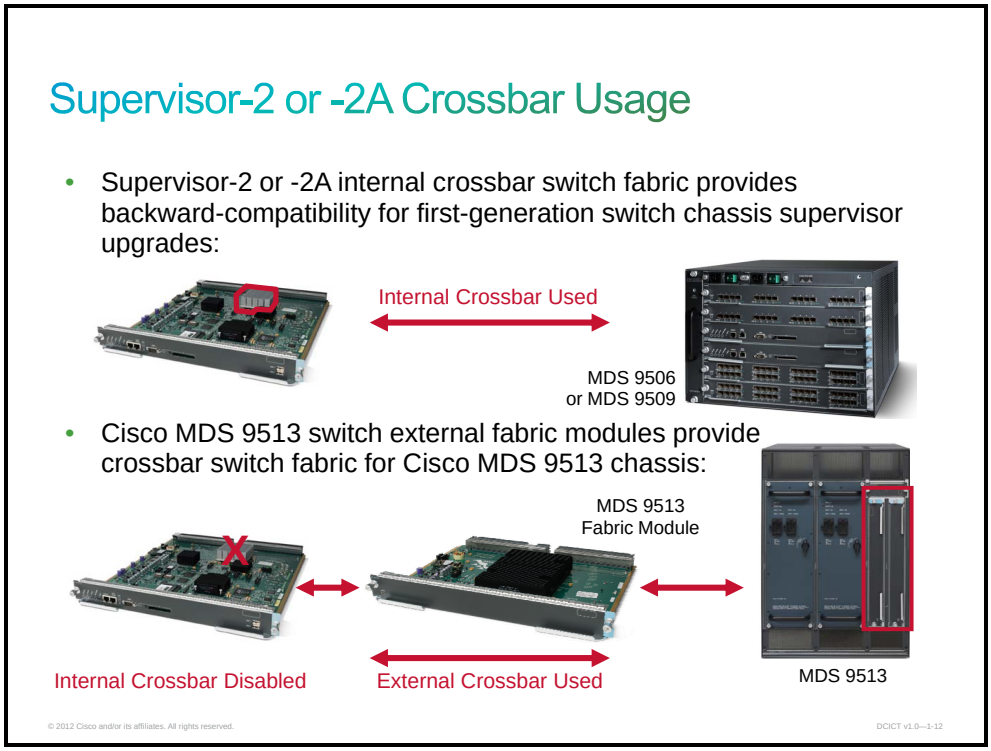
When it is used in a Cisco MDS 9506 or 9509 Multilayer Director switch, the integral crossbar is used. When it is used in the Cisco MDS 9513 Multilayer Director Switch, the integral crossbar is bypassed, and the Cisco MDS 9513 Crossbar Switching Fabric Modules are used instead.

The Supervisor-2 Module supports 1024 destination indexes providing up to 528 ports in a Cisco MDS 9513 Multilayer Director switch when only generation-2 or higher modules are used. If any generation-1 module is installed in the Cisco MDS 9513 Multilayer Director switch, then only 252 ports can be used.

The Cisco MDS 9500 Series Supervisor-2A Module is designed to integrate multiprotocol switching and routing, intelligent SAN services, and storage applications onto highly scalable SAN switching platforms. The Supervisor-2A Module enables intelligent, resilient, scalable, and secure high-performance multilayer SAN switching solutions. In addition to providing the same capabilities as the Supervisor-2 Module, the Supervisor-2A Module supports deployment of FCoE in the chassis of the Cisco MDS 9500 Series Multilayer Directors. The Cisco MDS 9000 Family lowers the total cost of ownership (TCO) for storage networking by combining robust and flexible hardware architecture, multiple layers of network and storage intelligence, and compatibility with all Cisco MDS 9000 Family switching modules. This powerful combination helps organizations build highly available, scalable storage networks with comprehensive security and unified management.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



The internal crossbar of the Supervisor-2 and -2A Modules provides backward-compatibility and is used for switching in first-generation chassis, such as the Cisco MDS 9506 or 9509 Multilayer Director switch.

In the Cisco MDS 9513 Multilayer Director chassis, the Supervisor-2 internal crossbar is disabled and the higher bandwidth external Cisco MDS 9513 Crossbar Switching Fabric Modules are used instead.

https://t.me/learningnets

Cisco MDS 9513 Crossbar Switching Fabric

Redundant crossbar fabric:

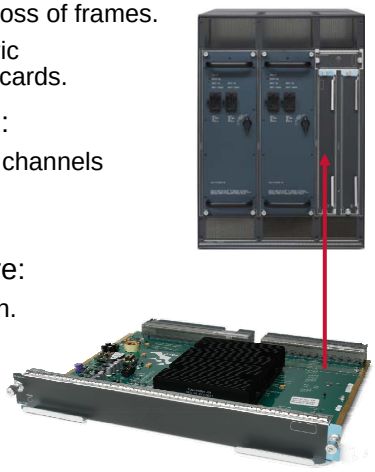
- Active/active operation balances the load across both crossbars.
- Rapid failover in case of failure ensures no loss of frames.
- In the event of failure, a single crossbar fabric still provides sufficient bandwidth for all line cards.

High-bandwidth nonblocking architecture:

- Each crossbar fabric supports dual 24-Gb/s channels from each line card.
- Total crossbar bandwidth = 2.2 Tb/s

High-performance centralized architecture:

- Ensures consistent latency across the switch.
- Supports up to 1024 indexes (destination interfaces).
- Enhanced high performance arbiter schedules frames at over 1 billion f/s.



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

Both Cisco MDS 9513 Crossbar Switching Fabric Modules are located at the rear of the chassis and provide a total aggregate bandwidth of 2.2 Tb/s.

Each fabric module is connected to each of the line cards via dual redundant 24-Gb/s channels making a total of 96 Gb/s per slot.

A single Cisco MDS 9513 Crossbar Switching Fabric Modules can support full bandwidth on all connected ports in a fully loaded Cisco MDS 9513 Multilayer Director switch without blocking.

The arbiter schedules frames at over 1 billion frames per second (f/s), ensuring that blocking will not occur even when the ports are fully utilized.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco MDS 9513 Crossbar Switching Fabric Modules

- Fabric 1 module provides 48 Gb/s per slot bandwidth.
- Fabric 2 module doubles active backplane bandwidth to 96 Gb/s per slot.
  - MDS 9000 24- and 48-Port 8-Gbps Fibre Channel Switching Modules require Fabric-2 for full bandwidth
  - MDS 9000 4/44-Port 8-Gbps Host-Optimized Fibre Channel Switching Module only requires 48 Gb/s, so it can use Fabric-1 or Fabric-2 modules.
  - Switch reload is required during upgrade from Cisco MDS SAN-OS 3.x to Cisco NX-OS 4.1 to support the 24- and 48-Port 8-Gbps Fibre Channel Switching Modules.
  - Fabric 2 module simplifies migration to 8-Gb/s modules.
- Fabric 3 Module provides 256 Gb/s per slot bandwidth



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

The Fabric 1 module provides up to 48 Gb/s per slot bandwidth.

The Fabric 2 module doubles active backplane bandwidth to 96 Gb/s per slot.

- Required for Cisco MDS 9000 24- and 48-port 8-Gbps Fibre Channel Switching Modules

Note	The Fabric 2 module is not required for Cisco MDS 9000 4/44-Port 8-Gbps Host-Optimized Fibre Channel Switching Module.

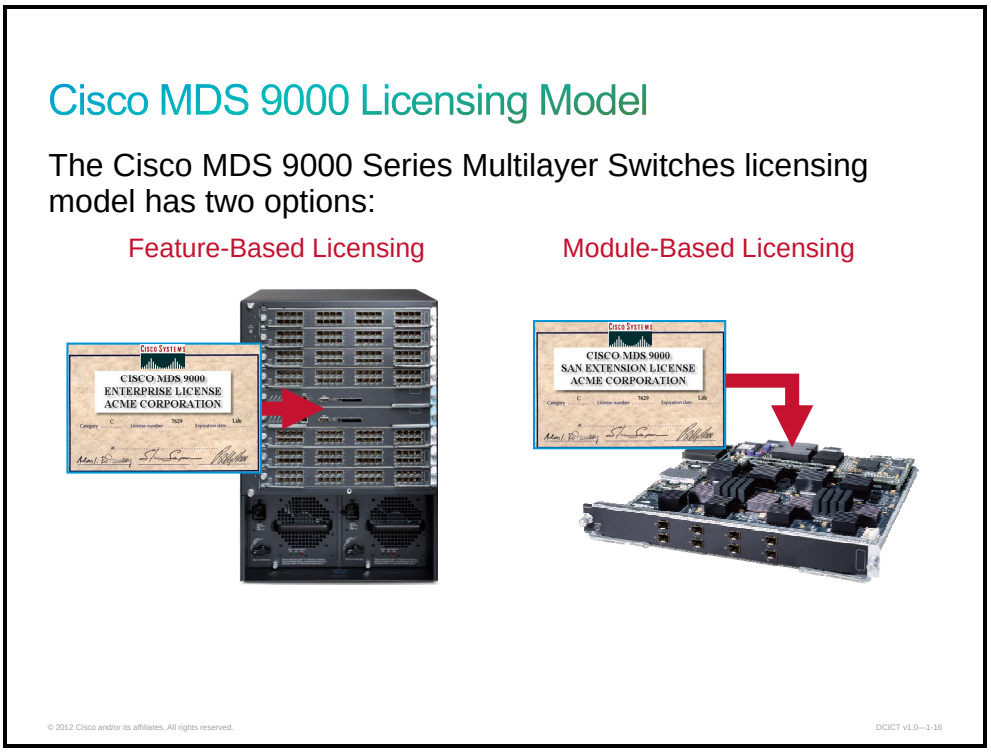
- Requires switch reload during upgrade from Cisco MDS SAN-OS Software Release 3.x to Cisco Nexus Operation System (Cisco NX-OS) Software Release 4.1 to support 24- and 48-port 8-Gbps Fibre Channel Switching Modules
- Simplifies migration to 8-Gb/s modules

Fabric 3 module provides up to 256 Gb/s per slot bandwidth.

https://t.me/learningnets

Cisco MDS 9000 Series Licensing Options

To provide flexibility in deploying new features on the switch, Cisco uses a licensing model for the Cisco Nexus Operating System (Cisco NX-OS) Software. This topic describes licensing options for the Cisco MDS 9000 Series Multilayer Switches.



Licensing allows access to specified premium features on the switch after you install the appropriate licenses. Licenses are sold, supported, and enforced for all releases of the Cisco NX-OS Software.

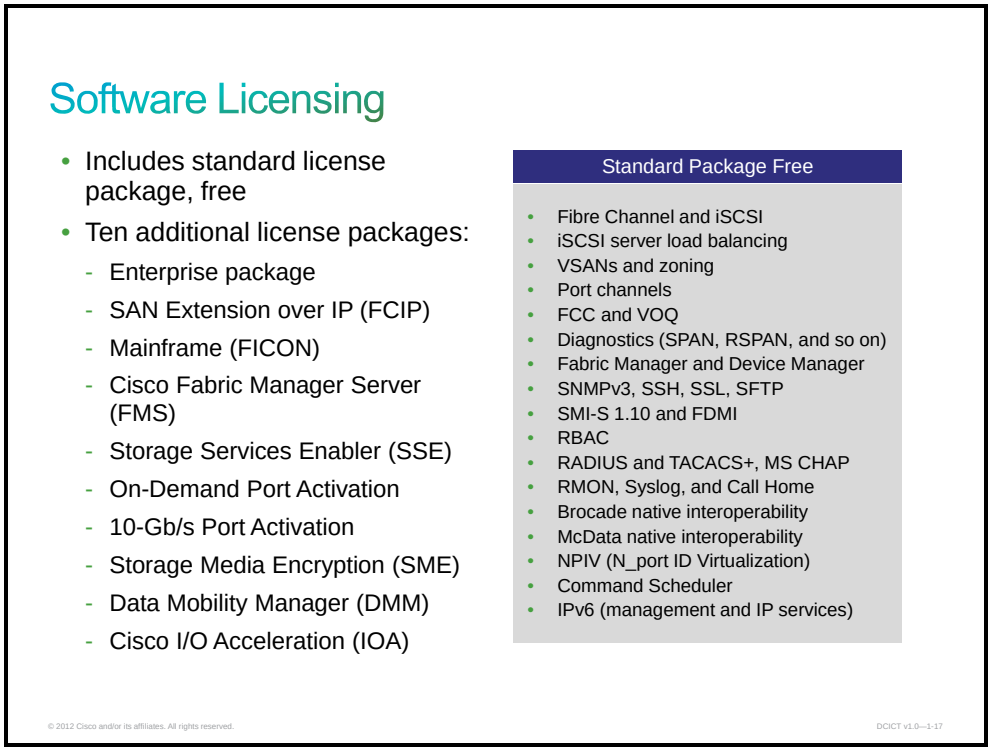
The licensing model that is defined for the Cisco MDS 9000 Series product line has two options:

- Feature-based licensing covers features that are applicable to the entire switch. The cost varies based on per-switch usage.
- Module-based licensing covers features that require additional hardware modules. The cost varies based on per-module usage.

Note The FCIP license that is bundled with Cisco MDS 9222i switches enables FCIP on the two fixed IP Services ports only. The features that are enabled on these ports by the bundled license are identical to the features that are enabled by the FCIP license on the Cisco MDS 9000 14/2-port Multiprotocol Services Module. If you install a module with IP ports in the empty slot on the Cisco MDS 9222i switch, you will need a separate FCIP license to enable FCIP on the IP ports of the additional I/O module.

Note Licensing on the Cisco MDS 9000 IP Storage Services Module (Cisco MDS 9000 16-Port Storage Services Node [SSN-16]) has the following limitations:
• Only one licensed feature can run on an SSN-16 engine at any time.
• On a given SSN-16 module, you can ”mix and match” the Cisco MDS Input/Output Accelerator (MDS 9000 IOA) license and SAN extension over IP license on the four service engines in any combination—for example, 4+0, 1+3, 2+2, 3+1, or 0+4, respectively.
• The SSN-16 module does not support “mix and match” for the Cisco MDS SME license.

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



The Cisco NX-OS Software is the underlying system software that powers the award-winning Cisco MDS 9000 Series Multilayer Switches. Cisco NX-OS is designed for SANs in the best traditions of Cisco IOS Software 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 provides many unique features that help the Cisco MDS 9000 Series 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 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.

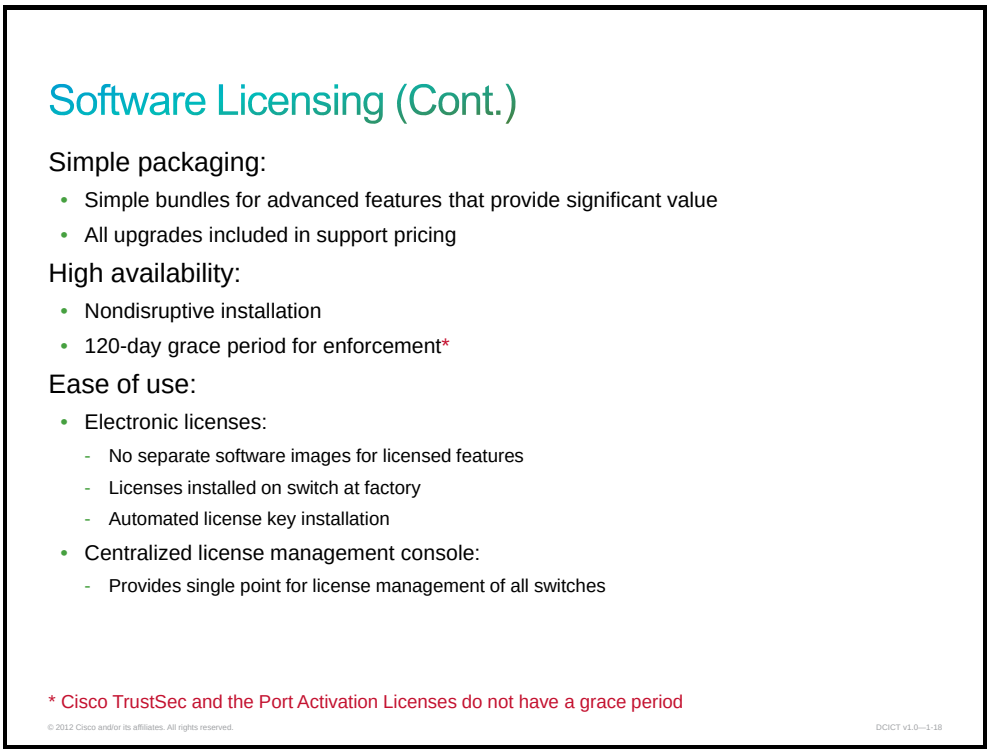
Most Cisco MDS 9000 Series software features are included in the base switch configuration. The standard software package includes the base set of features that Cisco believes are required by most customers for building a SAN. However, some features are logically grouped into add-on packages that must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

The Cisco NX-OS Software feature packages are as follows:

- **Enterprise package:** Adds a set of advanced features that are recommended for all enterprise SANs.
- **SAN Extension over IP package:** Enables FCIP for IP storage services and allows the customer to use the IP storage services to extend SANs over IP networks.
- **Mainframe package:** Adds support for the fiber connectivity (FICON) protocol. FICON VSAN support is provided to help ensure that there is true hardware-based separation of FICON and open systems. Switch cascading, fabric binding, and intermixing are also included in this package.

https://t.me/learningnets

- **Data Center Network Manager for SAN package:** Extends Cisco Fabric Manager by providing historical performance monitoring for network traffic hotspot analysis, centralized management services, and advanced application integration for greater management efficiency.
- **Storage Services Enabler package:** The Cisco MDS 9000 Storage Services Enabler (SSE) Package enables network-hosted storage applications to run on the Cisco MDS 9000 SSN-16.
- **On-Demand Port Activation:** Enables ports in bundles of eight as port requirements expand.
- **Storage Media Encryption:** Adds support for the Cisco MDS SME for a Cisco MDS 9000 SSN-16 engine or a Cisco MDS 9222i switch.
- **Data Mobility Manager:** Adds support for the Cisco Data Mobility Manager (Cisco DMM) feature on the Cisco MDS 9000 IP Storage Services Module (Cisco 9000 IP Module) in a Cisco MDS 9000 Series switch.
- **Cisco IOA:** Activates Cisco MDS 9000 IOA for the Cisco MDS 9000 SSN-16 module.
- **Extended Remote Copy (XRC)**: Enables support for FICON XRC acceleration on the Cisco MDS 9222i switch.



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 stores license keys on the chassis serial PROM (SPROM), so license keys are never lost even during a software reinstallation.

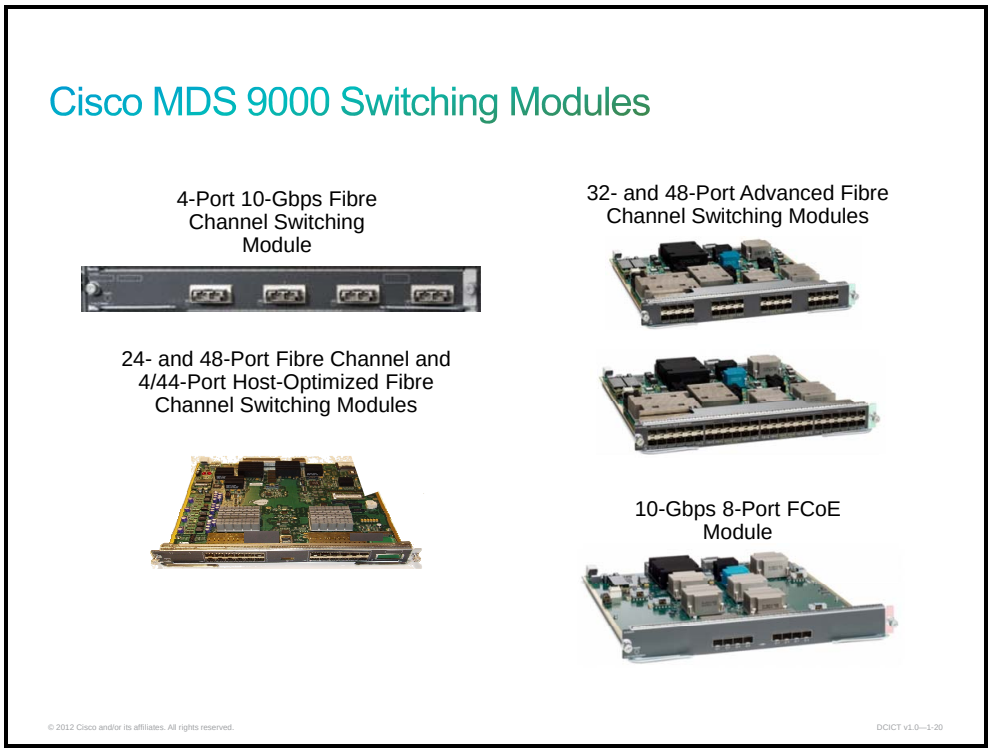
Cisco Data Center Network Manager for SAN includes a centralized license management console that provides a single interface for managing licenses across all Cisco MDS switches in the fabric. This single console reduces management overhead and prevents problems because of improperly maintained licensing. If an administrative error does occur with licensing, the switch provides a grace period before the unlicensed features are disabled. The grace period allows plenty of time to correct the licensing issue.

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

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

Cisco MDS 9000 Series Switching Modules

There are several modules available for the modular switches in the Cisco MDS 9000 Series product range. This topic describes the capacities of the Cisco MDS 9000 Series switching modules.



Cisco MDS 9000 4-Port 10-Gbps Fibre Channel Switching Module

The Cisco MDS 9000 4-Port 10-Gbps Fibre Channel Switching Module operates in full-rate mode. Each port on the module can deliver up to 10 Gb/s and operate as a port group.

The internal path to the forwarding ASIC provides 12 Gb/s, so more than enough bandwidth is available.

The 4-port 10-Gbps Fibre Channel Switching Module is suitable for any device that requires full 10-Gb/s bandwidth, for example, ISLs to other switches.

Cisco MDS 9000 Family 8-Gbps Fibre Channel Switching Modules

There are three Cisco MDS 9000 Family 8-Gbps Fibre Channel Switching Modules that provide higher performance and flexibility when you configure storage and server ports:

- The Cisco MDS 9000 24-Port 1/2/4/8-Gbps Fibre Channel Switching Module provides 8 Gb/s at 2:1 oversubscription and Full Rate (FR) bandwidth on each port at 1, 2, and 4 Gb/s.
- The Cisco MDS 9000 48-Port 1/2/4/8-Gbps Fibre Channel Switching Module provides 8 Gb/s at 4:1 oversubscription, 4 Gb/s at 2:1 oversubscription and FR bandwidth at 1 and 2 Gb/s on each port.
- The Cisco MDS 9000 4/44-Port 8-Gbps Host-Optimized Fibre Channel Switching Module provides four ports at 8 Gb/s with the remaining 44 ports at 4 Gb/s.

https://t.me/learningnets

Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules

Cisco MDS 9000 Family 8-Gpbs Advanced Fibre Channel Switching Modules are available in two configurations:

- The Cisco MDS 9000 32-Port 8-Gbps Advanced Fibre Channel Switching Module delivers line-rate performance across all ports and is ideal for high-end storage subsystems and for ISL connectivity.
- The Cisco MDS 9000 48-Port 8-Gbps Advanced Fibre Channel Switching Module provides higher port density and is ideal for connection of high-performance virtualized servers. With Arbitrated Local Switching enabled, this module supports 48 ports of line rate 8 Gb/s and is perfect for deploying dense virtual machine (VM) clusters with locally mapped storage. For traffic that is switched across the backplane, this module supports 1.5:1 oversubscription at 8-Gb/s Fibre Channel rate across all ports.

The 8-Gbps Advanced Fibre Channel Switching Modules are compatible with all Cisco MDS 9500 Series Multilayer Directors.

Advanced Fibre Channel Switching Module Features

- **Cisco FlexSpeed:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules are equipped with Cisco FlexSpeed technology, which enables ports on the Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules to be configured as either 1/2/4/8-Gb/s or 10-Gb/s Fibre Channel interfaces. The 10-Gb/s interfaces enable reduced cabling for ISLs because they provide a 50 percent higher data rate than 8-Gb/s interfaces. With integrated Cisco TrustSec encryption, the 10-Gb/s links provide secure, high-performance native Fibre Channel SAN Extension. Both 32- and 48-port modules support up to 24 10-Gb/s Fibre Channel interfaces. These modules enable consolidation of 1/2/4/8-Gb/s and 10-Gb/s ports into the same Fibre Channel switching module, conserving space on the Cisco MDS 9000 Series chassis.
- **Cisco Arbitrated Local Switching:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules provide line-rate switching across all the ports on the same module without performance degradation or increased latency for traffic that is exchanged with other modules in the chassis. This capability is achieved through Cisco MDS 9500 Series Multilayer Directors crossbar architecture with a central arbiter arbitrating fairly between local traffic and traffic to and from other modules. Local switching can be enabled on any Cisco MDS 9500 Series director-class chassis.
- **Integrated Hardware-Based VSANs and Inter-VSAN Routing (IVR):** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules enable deployment of large-scale consolidated SANs while maintaining security and isolation between applications. Integration into port-level hardware allows any port or any VM in a system or fabric to be partitioned into any VSAN. Integrated hardware-based IVR provides line-rate routing between any ports in a system or fabric without the need for external routing appliances.
- **Resilient High-Performance ISLs:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules support high-performance ISLs consisting of 8- or 10-Gb/s secure Fibre Channel. Advanced Fibre Channel switching modules also offer port channel technology. These modules provide up to 16 links spanning any port on any module within a chassis that is grouped into a logical link for added scalability and resilience. Up to 4095 buffer-to-buffer credits (BB_Credits) can be assigned to a single Fibre Channel port, providing industry-leading extension of storage networks. Networks may be extended to greater distances by tuning the BB_Credits to ensure that full link bandwidth is maintained.

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

- **Intelligent Fabric Services:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules provide integrated support for VSAN technology, Cisco TrustSec encryption, access control lists (ACLs) for hardware-based intelligent frame processing, and advanced traffic-management features to enable deployment of large-scale enterprise storage networks. The 8-Gbps Advanced Fibre Channel Switching Modules provide Fibre Channel Redirect (FC-Redirect) technology, which is a distributed flow redirection mechanism that can enable redirection of a set of traffic flows to an intelligent Fabric Service such as Cisco MDS 9000 IOA, Cisco DMM, and SME.
- **Advanced FICON Services:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules support 1/2/4/8-Gb/s FICON environments, including cascaded FICON fabrics, VSAN-enabled intermix of mainframe and open systems environments, and N-Port ID Virtualization (NPIV) for mainframe Linux partitions. FICON Control Unit Port (CUP) support enables in-band management of Cisco MDS 9000 Series switches from the mainframe management console.
- **Comprehensive Security Framework:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules support RADIUS and TACACS+, Fibre Channel Security Protocol (FC-SP), Secure FTP (SFTP), Secure Shell (SSH) Protocol, and Simple Network Management Protocol Version 3 (SNMPv3) implementing Advanced Encryption Standard (AES), VSANs, hardware-enforced zoning, ACLs, and per-VSAN role-based access control (RBAC). Cisco TrustSec Fibre Channel Link Encryption that is implemented by the 8-Gbps Advanced Fibre Channel switching modules secures sensitive data within or across data centers over high-performance 8- and 10-Gb/s Fibre Channel links.
- **Sophisticated Diagnostics:** Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules provide intelligent diagnostics, protocol decoding, and network analysis tools as well as integrated Call Home capability for added reliability, faster problem resolution, and reduced service costs.

Cisco MDS 9000 10-Gbps 8-Port FCoE Module

The Cisco MDS 9000 10-Gbps 8-Port FCoE Module can bridge the gap between the traditional Fibre Channel SAN and the evolution to FCoE. Enterprises can reduce capital expenditures (CapEx) and operating expenses (OpEx) in the data center by deploying FCoE while using the Cisco MDS 9000 10-Gbps 8-Port FCoE Module to protect their investment in existing Fibre Channel SANs and Fibre Channel-attached storage.

https://t.me/learningnets

Cisco MDS 9000 SSN-16

- 16x Gigabit Ethernet ports for FCIP WAN connectivity
  - Up to 3 FCIP tunnels per port
- 4 independent MDS 9000 IOA engines supporting:
  - FCIP hardware compression
  - FCIP hardware encryption
  - FCIP write acceleration
  - FCIP tape read and write acceleration with PortChannel support
  - Fibre Channel write acceleration
  - Fibre Channel tape read and write acceleration
  - Storage Media Encryption (SME)



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

The Cisco MDS 9000 16-Port Storage Services Node (SSN-16) hosts four independent service engines.

Each of the service engines can be activated individually and incrementally to scale as business requirements change or they can be configured to run separate applications.

Based on the single service engine originally in the Cisco MDS 9000 18/4-Port Multiservice Module, this four-to-one consolidation delivers dramatic hardware savings and frees valuable slots in the chassis of the Cisco MDS 9500 Multilayer Directors.

Supported applications include the following:

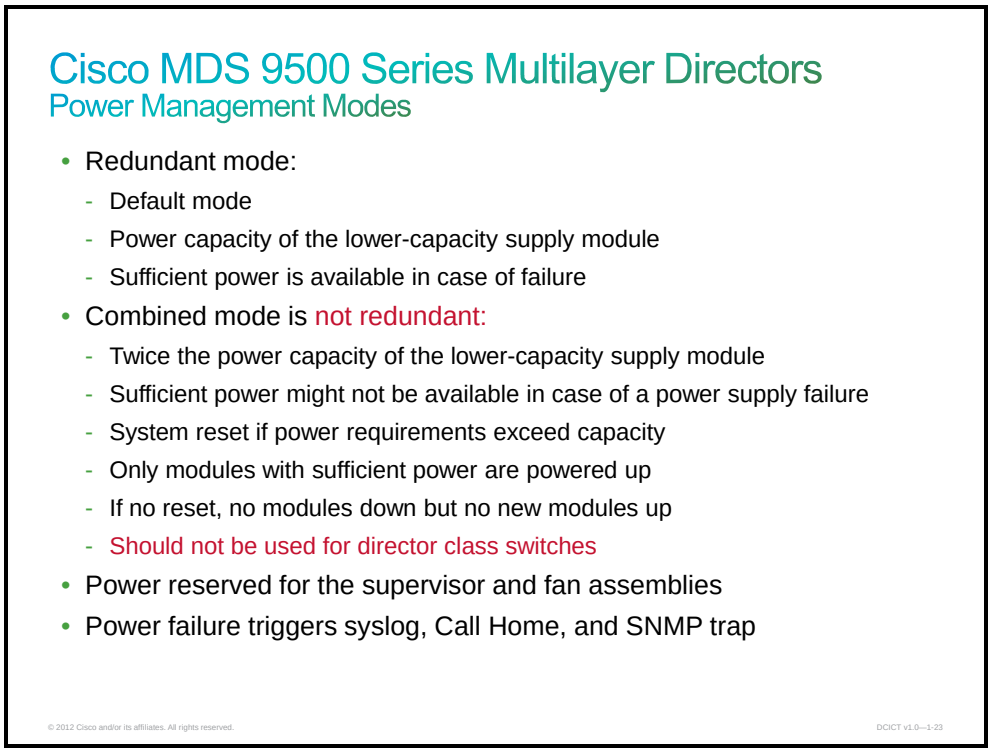
- Remote SAN extension with high-performance Fibre Channel over IP (FCIP)
- Metropolitan-area network (MAN) link optimization with Cisco MDS 9000 IOA
- Encryption of storage at rest with Cisco MDS SME

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco MDS 9500 Series Power Supply Options

Each switch in the Cisco MDS 9000 Series supports redundant power supplies. This topic describes the capabilities of Cisco MDS 9500 power supply options.



Power supply modules are configured in redundant mode by default, but they can also be configured in a combined, or nonredundant, mode:

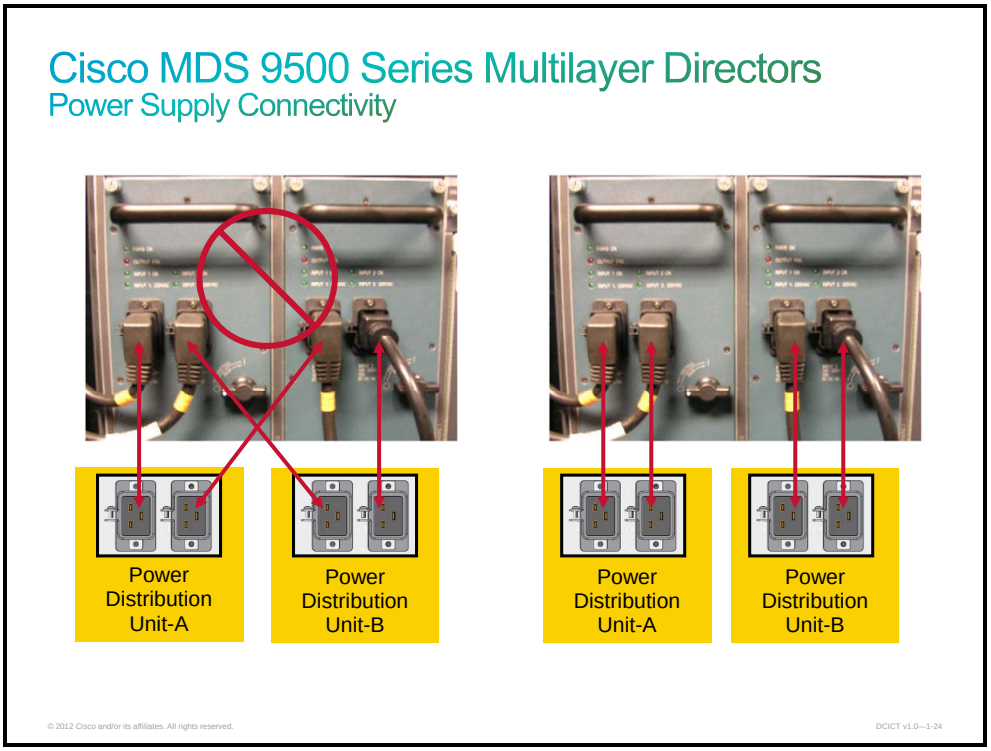
- In redundant mode, the chassis uses the power capacity of the lower-capacity power supply module so that sufficient power is available in a single power supply failure.
- In combined mode, the chassis uses twice the power capacity of the lower-capacity power supply modules. Sufficient power might not be available in a power supply failure in this mode. If there is a power supply failure and the real power requirements for the chassis exceed the power capacity of the remaining power supply modules, the entire system is reset automatically. This reset will prevent permanent damage to the power supply module.

In both modes, power is reserved for the supervisor and fan assemblies. Each supervisor module has roughly 220 W in reserve, even if there is only one that is installed, and the fan module has 210 W in reserve. If there is insufficient power, after supervisors and fans are powered, line card modules are given power from the top of the chassis down:

- After the reboot, only those modules that have sufficient power are powered up.
- If the real power requirements do not trigger an automatic reset, no module is powered down. Instead, no new module is powered up.
- In all cases of power supply failure or removal, a syslog message is printed, a Call Home message is sent (if configured), and an SNMP trap is sent.

Note Combined mode should not be used for director-class switches.

https://t.me/learningnets



The 6000-W AC power supply modules for the Cisco MDS 9513 Multilayer Director are designed to provide output power for the modules and fans. Each power supply module has two AC power connections and provides power as follows:

- One AC power connection at 110 VAC (no output)
- Two AC power connections at 110 VAC (2900-W output)
- One AC power connection at 220 VAC (2900-W output)
- Two AC power connections at 220 VAC (6000-W output)

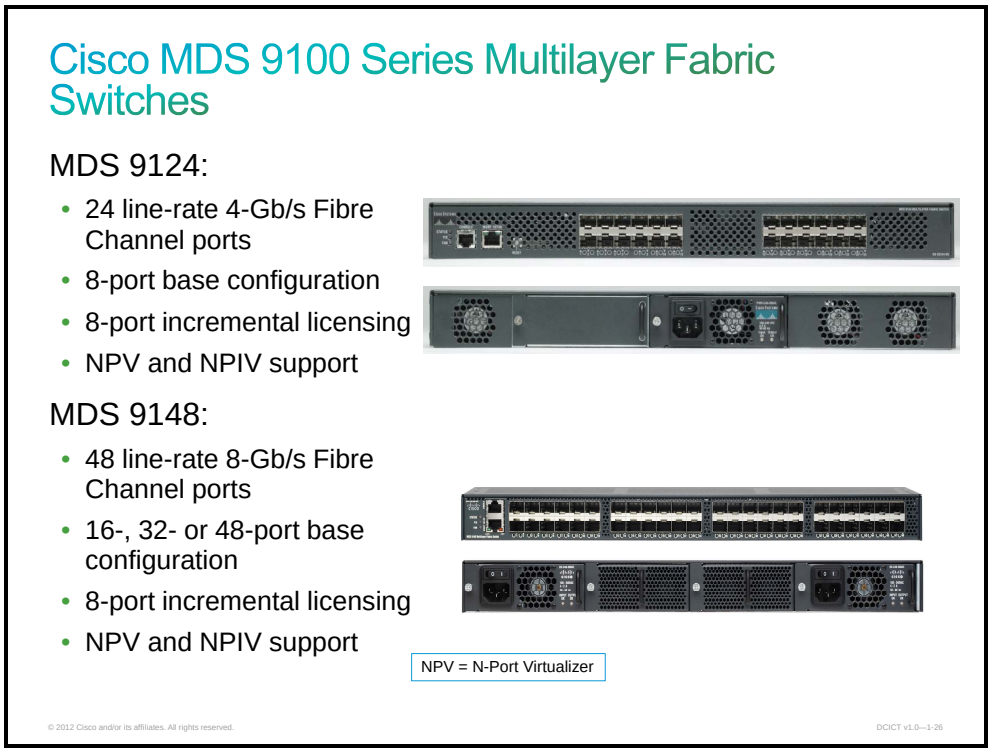
The figure illustrates incorrect and correct connection sequences to external power distribution units. When the power supply modules are configured in redundant mode (default), the wrong connection sequence can leave the switch with only 2900 W, which can result in the chassis shutting down line card modules. This event would require manual configuration to set combined mode, upon which 5800-W power capacity is available. There is no automatic provision to configure combined mode if there is a loss of an external power source.

The correct connection sequence provides the complete 6000-W capacity if an external power source fails.

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

Cisco MDS 9100 Series Switches

The Cisco MDS 9100 Series Multilayer Fabric Switches are edge fabric switches for the SAN infrastructure. This topic describes the capabilities of Cisco MDS 9100 Series Multilayer Fabric Switches.



Cisco MDS 9124 Multilayer Fabric Switch

The Cisco MDS 9124 Fabric Switch supports 24 line-rate 4-Gb/s ports. There are four ports per port group and 64 BB_Credits per port group.

You can assign up to 61 BB_Credits to a single port in each port group if you need to support long-distance ISLs.

The Cisco MDS 9124 Fabric Switch supports Port Licensing supporting 8-, 16-, or 24-port configurations.

The base configuration has eight ports that are enabled and a single power supply. An optional second power supply can be fitted for high availability.

Cisco MDS 9148 Multilayer Fabric Switch

This Cisco MDS 9148 Fabric Switch provides an affordable and scalable storage networking solution for small, midsize, and large enterprise customers. The switch provides line-rate 8-Gb/s ports that are based on a purpose-built "switch-on-a-chip" ASIC with high performance, high density, and enterprise-class availability.

The switch provides flexibility, high availability, security, and ease of use, at an affordable price in a compact, 1-rack unit (1-RU) form factor. The Cisco MDS 9148 Fabric Switch has the flexibility to expand from 16 to 48 ports in eight-port increments, and offers the densities that are required to scale from an entry-level departmental switch to a top-of-rack (TOR) switch for edge connectivity in enterprise SANs. The Cisco MDS 9148 Fabric Switch offers a nonblocking architecture, with all 48 1-, 4-, and 8-Gb/s ports operating at line rate concurrently.

https://t.me/learningnets

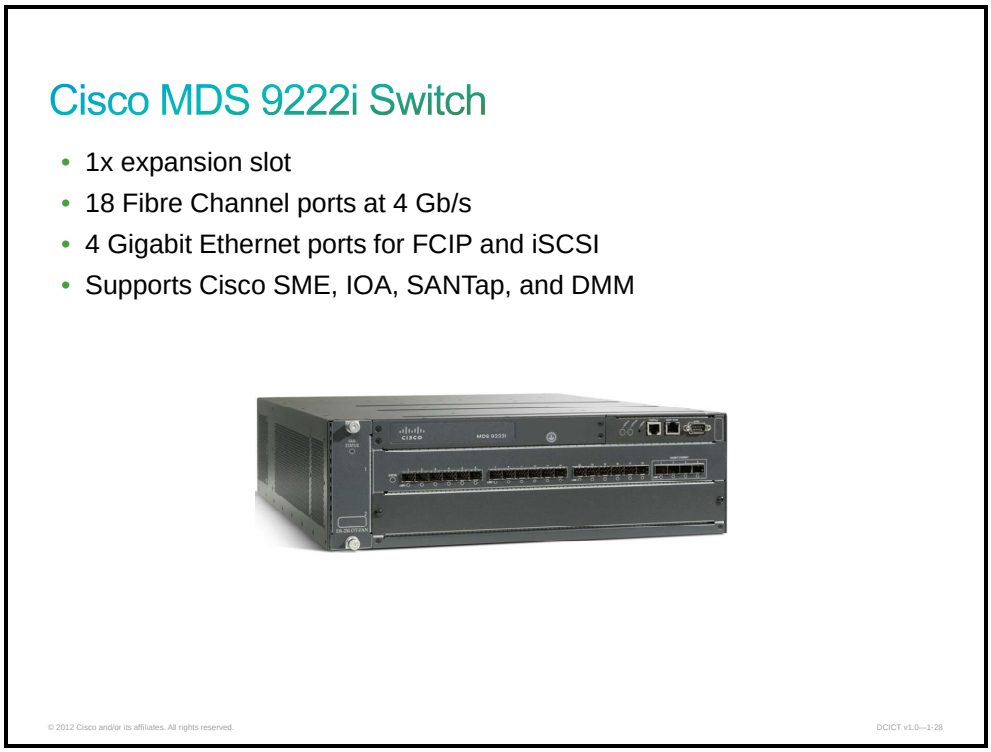
The Cisco MDS 9148 Fabric Switch supports a quick configuration, zero-touch, immediately active (plug-and-play) features. There are task wizards that allow it to be deployed quickly and easily in networks of any size. Powered by Cisco NX-OS Software, it includes advanced storage networking features and functions and is compatible with Cisco MDS 9500 Series Multilayer Directors and Cisco MDS 9200 Series Multilayer Switches, providing transparent, end-to-end service delivery in core-edge deployments.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco MDS 9222i Switch

The Cisco MDS 9222i Multiservice Modular Switch is a semimodular switch with one fixed module and a flexible slot for any I/O module that is available. This topic describes the capabilities of the Cisco MDS 9222i switch.



The Cisco MDS 9222i Multiservice Modular Switch delivers multiprotocol and distributed multiservice convergence, offering high-performance SAN extension and disaster recovery solutions. In addition, it supports intelligent Fabric Services such as Cisco MDS Storage Media Encryption (Cisco MDS SME) and cost-effective multiprotocol connectivity. The Cisco MDS 9222i switch has a compact form factor, and the modularity of the expansion slot supports advanced capabilities normally only available on director-class switches. The Cisco MDS 9222i switch is an ideal solution for departmental and remote branch-office SANs requiring the features that are present in a director, but at a lower cost of entry.

Product Highlights:

- High-density Fibre Channel switch, which scales up to 66 Fibre Channel ports
- Integrated hardware-based virtual fabric isolation with VSANs and Fibre Channel Routing with IVR
- Remote SAN extension with high-performance FCIP
- Long distance over Fibre Channel with extended buffer-to-buffer credits
- Multiprotocol and Mainframe support (Fibre Channel, FCIP, iSCSI, and FICON)
- IPv6 capable
- Platform for intelligent fabric applications such as Cisco MDS SME
- Cisco In-Service Software Upgrade (ISSU)
- Comprehensive network security framework
- Provides hosting, assisting, and acceleration of storage applications such as volume management, data migration, data protection, and backup with Cisco MDS 9000 switch Cisco 9000 IP Module
- Supports Cisco SANTap and DMM

https://t.me/learningnets

Summary

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

Summary

- The Cisco MDS 9000 Series switch is designed for SAN environments in the data center, providing a range of products to suit small- to large-scale deployments.
- The Cisco MDS 9500 Series chassis includes three models, the 6-, 9- and 13-slot chassis.
- The Cisco MDS 9500 Series chassis is modular, and has two slots available for supervisor engines. There are two current supervisor modules, the Supervisor-2 and the Supervisor-2A. The Supervisor-2A is required if FCoE support is to be enabled on the chassis.
- To ensure flexibility when upgrading to additional features, the Cisco MDS 9000 Series makes use of a licensing model. This licensing model provides seamless nondisruptive upgrades to additional feature sets.
- There are several modules that are supported in the Cisco MDS 9500 Series chassis and the Cisco MDS 9222i switch. These modules support line rates up to 10 Gb/s depending on the module that is being deployed.

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

Summary (Cont.)

- The Cisco MDS 9000 Series uses redundant power supply modules. There are several redundancy modes available, and customers should be aware of the cabling requirements for the Cisco MDS 9513 Multilayer Director chassis power supply modules.
- The Cisco MDS 9100 Series of switches are fabric switches designed for the access layer of the SAN infrastructure. The two models available today are the Cisco MDS 9124 Fabric Switch and the Cisco MDS 9148 Fabric Switch. The Cisco MDS 9148 Fabric Switch provides 8 Gb/s connectivity.
- The Cisco MDS 9222i switch is a semimodular switch providing connectivity at the access layer, small core, or for site-to-site connectivity through the use of features such as FCIP. It has one fixed module slot and one modular module slot.

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

https://t.me/learningnets

Lesson 4

Monitoring the Cisco Nexus 7000 and 5000 Series Switches

Overview

The Cisco Nexus switches have no default configuration and are only aware of an admin user. When the switch is initially switched on, an initial configuration needs to be performed. This lesson explains how to perform that initial configuration and validate common features that are used.

Objectives

Upon completing this lesson, you will be able to perform an initial configuration and validate common features of the Cisco Nexus 7000 and 5000 Series Switches. You will be able to meet these objectives:

- Describe how to connect to the console port of the Cisco Nexus 7000 and 5000 Series Switches
- Describe the how to run the initial setup script for the Cisco Nexus 7000 and 5000 Series Switches
- Describe how to connect to the CMP on the Cisco Nexus 7000 Series Switches
- Describe how to use SSH to connect to the management VRF on the Cisco Nexus 7000 and 5000 Series Switches
- Describe the ISSU capabilities of the Cisco Nexus 7000 and 5000 Series Switches
- Verify VLANs on the Cisco Nexus 7000 and 5000 Series Switches
- Describe the control, management, and data planes on the Cisco Nexus 7000 and 5000 Series Switches
- Describe CoPP on the Cisco Nexus 7000 and 5000 Series Switches
- Describe important CLI commands on the Cisco Nexus 7000 and 5000 Series Switches

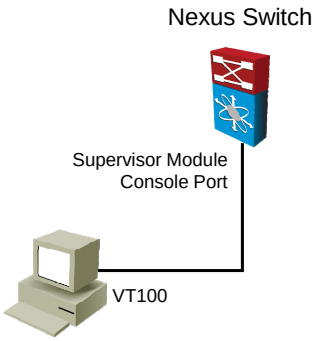
https://t.me/learningnets

Connecting to the Console Port

This topic describes how to connect to the console port on the Cisco Nexus 7000 and 5000 Series Switches.

Cisco Nexus Switch Console Port

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



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

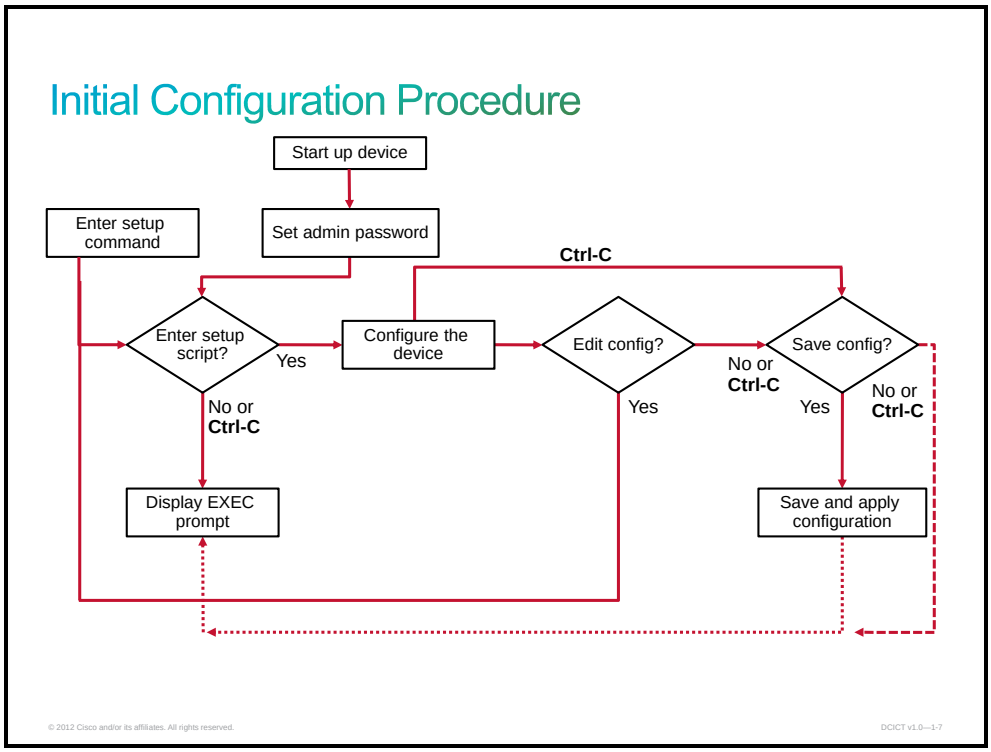
Initial connectivity to a switch is via the console port. The console requires a rollover RJ-45 cable. The terminal setup steps are listed in this figure. After the switch has booted up, it automatically goes into the default setup script the first time that the switch is started and configured.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Running the Initial Setup Script

There is no default configuration on Cisco Nexus switches. Therefore, on initial startup, the switch goes into setup mode. This topic describes how to run the initial setup script on the Cisco Nexus 7000 and 5000 Series Switches.



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

The setup utility is used mainly to configure the system initially when no configuration exists. However, it 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 any step, the setup utility changes the configuration using that default and not the configured value.

Note	Be sure to configure the IP version 4 (IPv4) route, the default network IPv4 address, and the default gateway IPv4 address to enable Simple Network Management Protocol (SNMP) access.

https://t.me/learningnets

Initial Configuration

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

Enter the password for "admin": 1234QWer
  Confirm the password for "admin": 1234QWer

 ---- Basic System Configuration Dialog VDC: 2 ----

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

 *Note: setup is mainly used for configuring the system initially, when
no configuration is present. So setup always assumes system defaults
and not the current system configuration values.

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

to skip the remaining dialogs.


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

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

The admin user is the only user that the switch knows about initially. On all Cisco NX-OS devices, strong passwords are the default. When configuring the password of the admin user, you need to use a minimum of eight characters—alpha, upper, and lower case, and numerical.

Once the admin user password is configured, you are asked if you wish to enter the “basic system configuration dialog.” Answering **yes** takes you through the basic setup.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Initial Configuration (Cont.)

```
Do you want to enforce secure password standard (yes/no) [y]: y
  Create another login account (yes/no) [n]: n
  Configure read-only SNMP community string (yes/no) [n]: n
  Configure read-write SNMP community string (yes/no) [n]: n
  Enter the switch name : PROD-A
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: y
    Mgmt0 IPv4 address : 172.16.1.201
    Mgmt0 IPv4 netmask : 255.255.255.0
  Configure the default gateway? (yes/no) [y]: y
    IPv4 address of the default gateway : 172.16.1.1
  Configure advanced IP options? (yes/no) [n]: n
  Enable the telnet service? (yes/no) [n]: n
  Enable the ssh service? (yes/no) [y]: y
    Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
    Number of rsa key bits <1024-2048> [1024]: 1024
  Configure default interface layer (L3/L2) [L3]: L3
  Configure default switchport interface state (shut/noshut) [shut]: shut
```

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

This figure shows the initial configuration default questions. If a question has a default answer in square brackets [xx], then the Enter key can be used to accept the default parameter.

Initial Configuration (Cont.)

```
The following configuration will be applied:
  password strength-check
  switchname PROD-A
interface mgmt0
ip address 172.16.1.201 255.255.255.0
no shutdown
vrf context management
ip route 0.0.0.0/0 172.16.1.1
exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
  no system default switchport
  system default switchport shutdown
Would you like to edit the configuration? (yes/no) [n]: <enter>
Use this configuration and save it? (yes/no) [y]: y
```

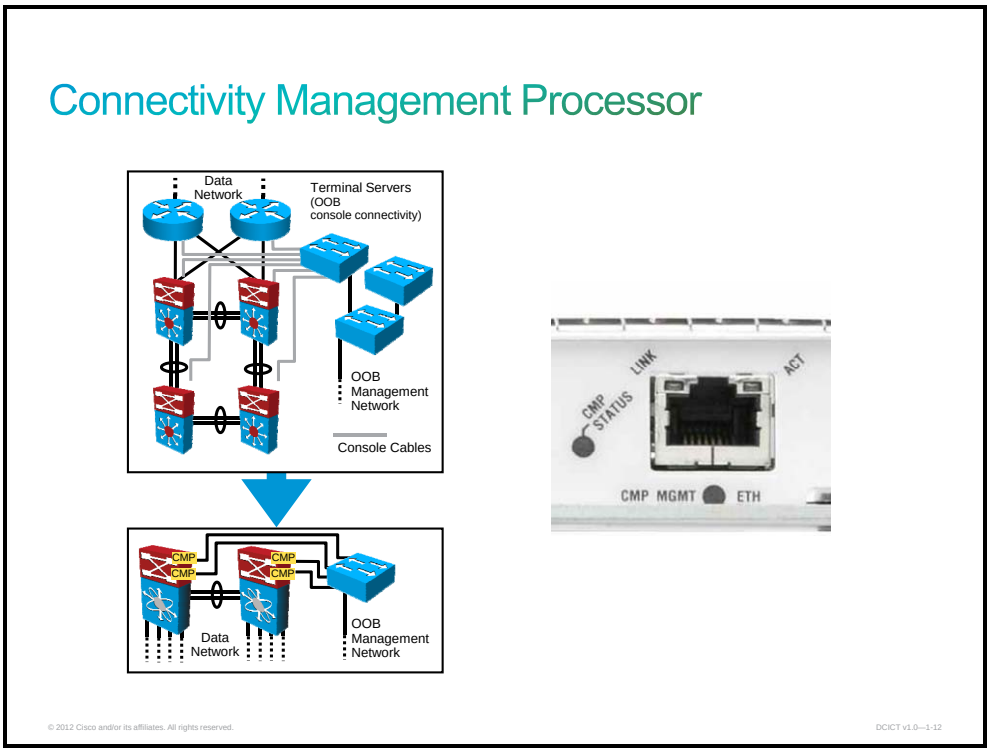
© 2012 Cisco and/or its affiliates. All rights reserved.

Once the initial parameters have been entered, the switch provides the initial configuration for you to verify. If it is correct, then there is no need to edit the configuration; you just need to accept the “Use this configuration and save it?” question by pressing Enter to save the configuration.

https://t.me/learningnets

Connecting to the Cisco Nexus 7000 CMP

For management high availability, the Cisco Nexus 7000 Series Switches have a connectivity management processor that is known as the CMP. This topic describes how to connect to the CMP.

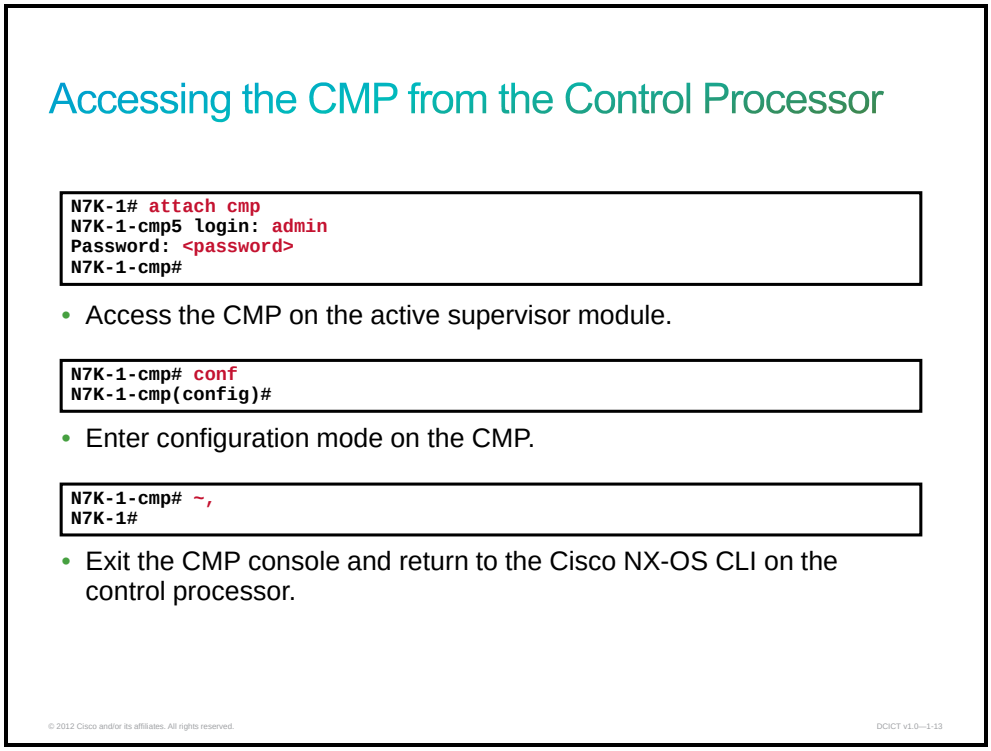


The CMP provides out-of-band (OOB) management and monitoring capability independent from the primary operating system. The CMP enables lights-out Remote Monitoring (RMON) and management of the supervisor module, all other modules, and the Cisco Nexus 7000 Series system without the need for separate terminal servers.

Key features of the CMP include the following:

- **Dedicated operating environment:** Independent remote system management monitoring capabilities.
- **Monitoring of supervisor status and initiation of resets:** Removes the need for separate terminal server devices for OOB management.
- **System reset while retaining OOB Ethernet connectivity:** Complete visibility during the entire boot process.
- **Capability to initiate a complete system restart:** No local operator intervention required.
- **Login authentication:** Provides secure access to the OOB management environment.
- **Access to supervisor logs:** Access to critical log information enables rapid detection and prevention of potential system problems.
- **Control capability:** Ability to take complete console control of the supervisor.
- **Dedicated front-panel LEDs:** CMP status is clearly identified separately from the supervisor.

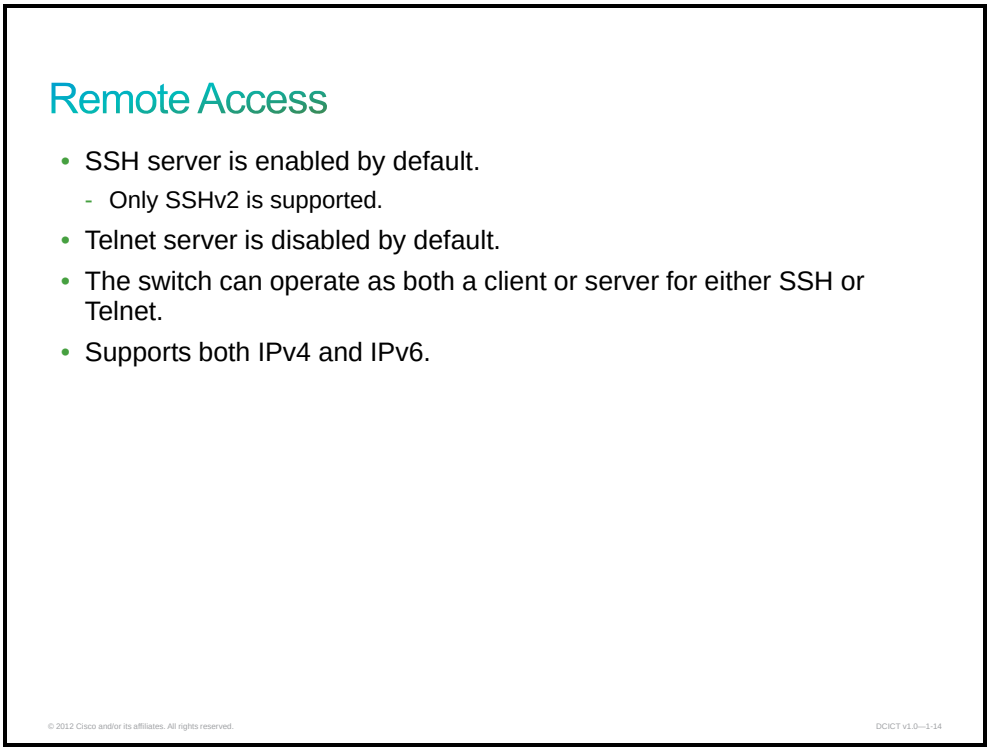
© 2012 Cisco Systems, Inc.
https://t.me/learningnets



You can access the CMP from the active supervisor module. Before you begin, ensure that you are in the default virtual device context (VDC).

When the control processor and CMP are both operational, you can log into the CMP through the control processor using your Cisco NX-OS username and password or the admin username and password. If the control processor is configured with RADIUS or TACACS, then your authentication is also managed by RADIUS or TACACS. If the control processor is operational, the CMP accepts logins from users with network-admin privileges. The CMPs use the same authentication mechanism to configure the control processor (that is, RADIUS, TACACS, or local). The control processor automatically synchronizes the admin password with the active and standby CMP. You are then able to use the admin username and password when a control processor is not operational.

https://t.me/learningnets



The Secure Shell (SSH) Protocol server feature enables an SSH client to make a secure, encrypted connection to a Cisco Nexus Series switch. SSH uses strong encryption for authentication. The SSH server in the Cisco Nexus Series switch interoperates with publicly and commercially available SSH clients.

By default, the SSH server is enabled on the Cisco Nexus Series switch. The Cisco Nexus Series switch supports only SSH version 2 (SSHv2).

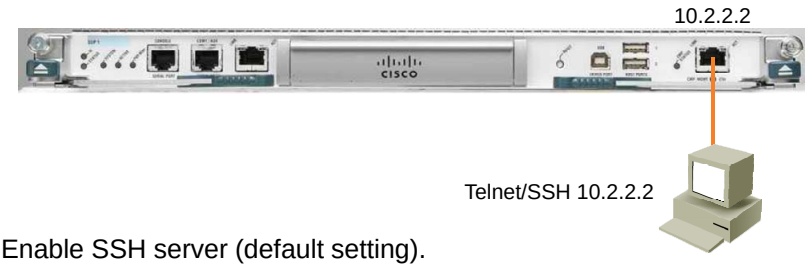
The Telnet Protocol enables a user at one site to establish a TCP/IP connection to a login server at another site. However, all parameters and keystrokes are passed in cleartext over the network, unlike SSH, which is encrypted. The Telnet server is disabled by default on the Cisco Nexus Series Switch.

To reach a remote system, both Telnet and SSH can accept any of the following:

- IP version 4 (IPv4) address
- IP version 6 (IPv6) address
- Device name that is resolved via the local host file or a distributed name server

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

Remote Access to the CMP



- Enable SSH server (default setting).

```
N7K-1-cmp(config)# ssh server enable
```

- Enable Telnet server.

```
N7K-1-cmp(config)# telnet server enable
```

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

You can access the CMP from the following:

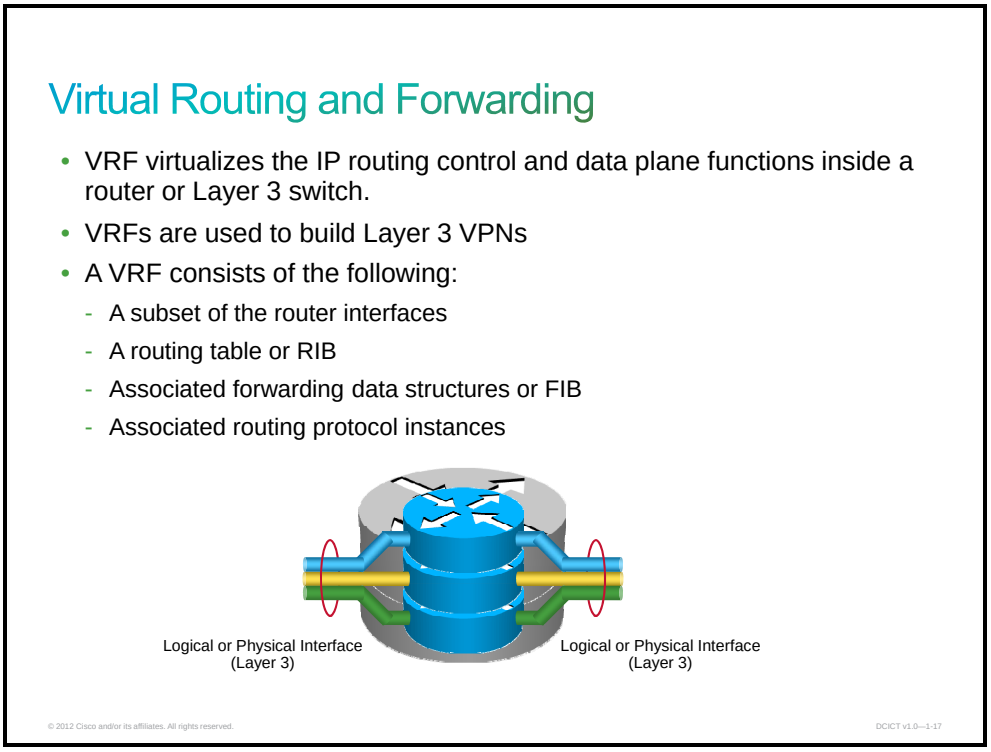
- Control processor
- SSH
- Telnet

To access the CMP by SSH or Telnet, you must enable those sessions on the CMP. (By default, the SSH server session is enabled.)

https://t.me/learningnets

Connecting to the Switch Using SSH to Connect to the Management VRF

The Cisco Nexus Series switch has two default virtual routing and forwarding instances (VRFs)—the management VRF and the default VRF. The management port on the Cisco Nexus Series switch automatically goes into the management VRF. This topic explains how to use SSH connect to the management VRF on Cisco Nexus 7000 and 5000 Switches.



To provide logical Layer 3 separation within a Layer 3 switch or router, the data plane and control plane functions of the device are segmented into different Layer 3 VPNs. This process is similar to the way that a Layer 2 switch segments the Layer 2 control and data plane into different VLANs.

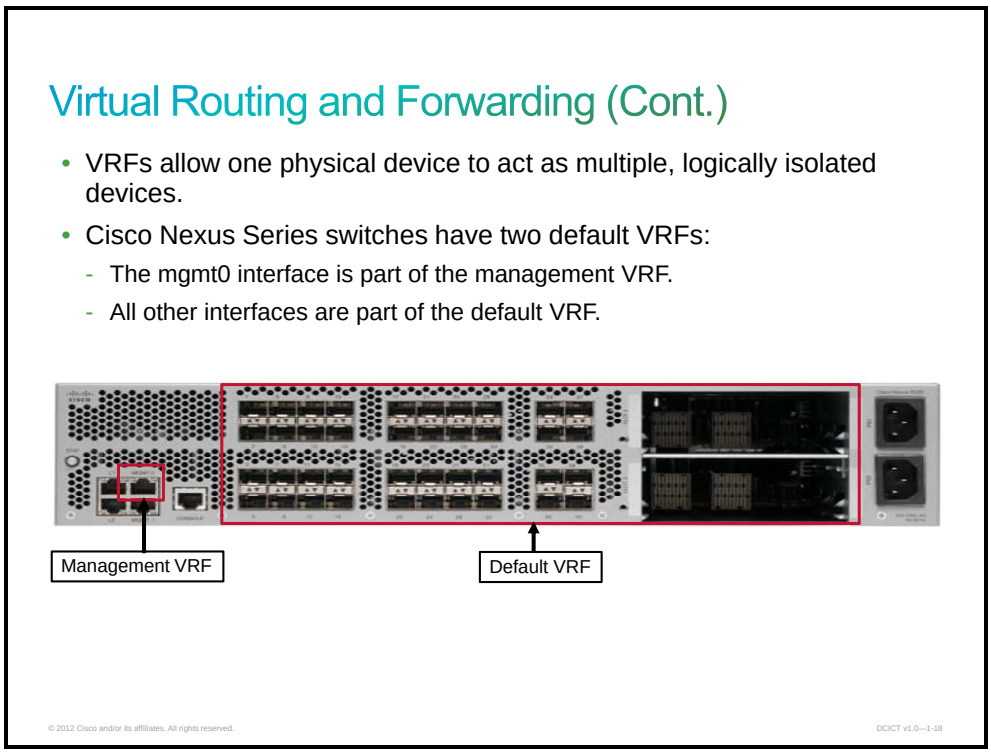
The core concept in Layer 3 VPNs is a VRF instance. This instance consists of all the data plane and control plane data structures and processes that together define the Layer 3 VPN.

A VRF includes the following components:

- **A subset of the Layer 3 interfaces on a router or Layer 3 switch:** Similar to how Layer 2 ports are assigned to a particular VLAN on a Layer 2 switch, the Layer 3 interfaces of the router are assigned to a VRF. Because the elementary component is a Layer 3 interface, this component includes software interfaces, such as subinterfaces, tunnel interfaces, loopback interfaces, and switch virtual interfaces (SVIs).
- **A routing table or Routing Information Base (RIB):** Traffic between Layer 3 interfaces that are in different VRFs should remain separated. Therefore, a separate routing table is necessary for each VRF. The separate routing table ensures that traffic from an interface in one VRF cannot be routed to an interface in a different VRF.
- **A Forwarding Information Base (FIB):** The routing table or RIB is a control plane data structure. From it, an associated FIB is calculated to be used in actual packet forwarding. The FIB also needs to be separated by the VRF.

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

- Routing protocol instances: To ensure control plane separation between the different Layer 3 VPNs, implement routing protocols on a per-VRF basis. To accomplish this task, you can run an entirely separate process for the routing protocol in the VRF. Or you can use a subprocess or routing protocol instance in a global process that is in charge of the routing information exchange for the VRF.



There are two VRFs that are enabled on the Cisco Nexus Series switch. The two VRFs cannot be deleted.

The management interface is permanently part of the management VRF. All remaining interfaces are, by default, part of the default VRF. The VRFs allow the management traffic to be logically isolated from the traffic on the switch ports.

Note Only VRF-Lite is supported on the Cisco Nexus 5500 Platform switches

https://t.me/learningnets

Verifying Management Reachability

Verify that the switch has reachability to the network.

```
Nexus-7010-PROD-A# ping 172.16.1.42 vrf management
PING 172.16.1.42 (172.16.1.42): 56 data bytes
64 bytes from 172.16.1.42: icmp_seq=0 ttl=254 time=0.771 ms
64 bytes from 172.16.1.42: icmp_seq=1 ttl=254 time=0.544 ms
64 bytes from 172.16.1.42: icmp_seq=2 ttl=254 time=0.702 ms
64 bytes from 172.16.1.42: icmp_seq=3 ttl=254 time=0.62 ms
64 bytes from 172.16.1.42: icmp_seq=4 ttl=254 time=0.699 ms

--- 172.16.1.42 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.544/0.667/0.771 ms
Nexus-7010-PROD-A#
```

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

It is important to verify reachability to any Layer 3 interfaces that are configured on the Cisco Nexus Series switch. Use the **ping** command to ensure that the address of the default gateway, or any other device in the network, is reachable.

Unless otherwise specified, the ping and all other Layer 3 traffic originate from the default VRF. Therefore, it is important to specify which VRF is to be used.

Initiating Remote Connections

Initiate an SSH session.

```
Nexus-7010-PROD-A# ssh 172.16.1.41 vrf management
The authenticity of host '172.16.1.41 (172.16.1.41)' can't be
established.
RSA key fingerprint is e6:de:4d:cf:08:ea:80:cd:e9:de:34:6b:df:44:e5:30.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '172.16.1.41' (RSA) to the list of known
hosts.
Nexus 5000 Switch
Password:
```

```
Nexus-7010-PROD-A# ssh admin@172.16.1.41 vrf management
Nexus 5000 Switch
Password:
```

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

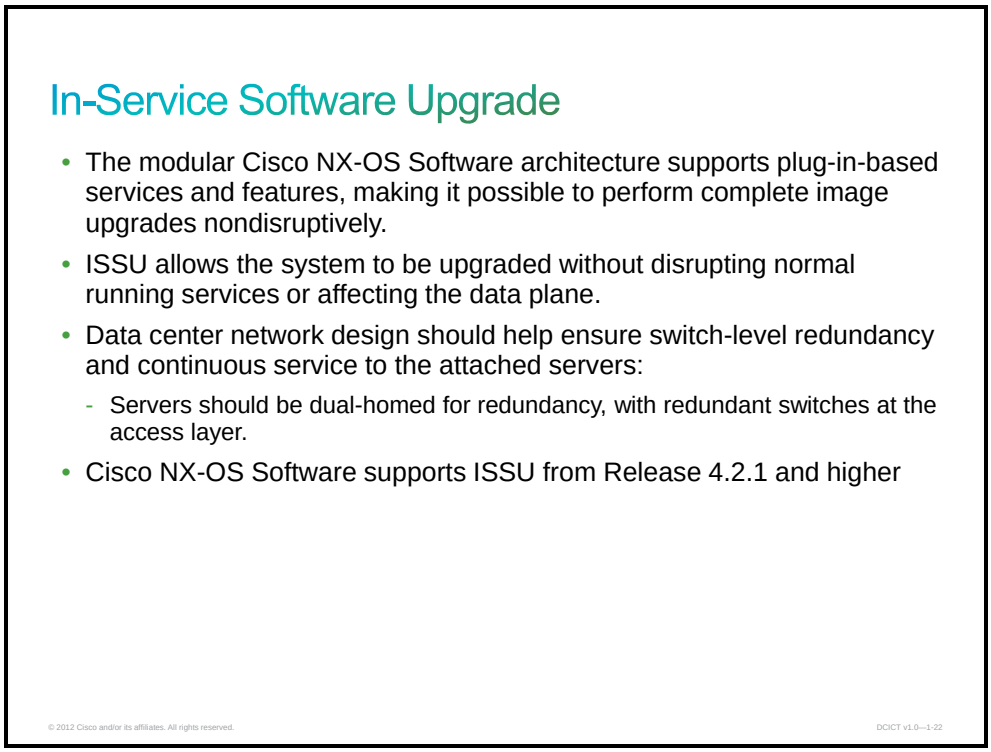
The figure shows how to initiate an SSH session from the local switch.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Reviewing the ISSU on the Cisco Nexus Switches

In-Service Software Upgrades (ISSUs) allow administrators to upgrade the Cisco Nexus switch nondisruptively. This topic describes the ISSU capabilities of the Cisco Nexus 7000 and 5000 Series Switches.



An ISSU is initiated manually either through the CLI by an administrator or via the management interface of the Cisco Data Center Network Manager software platform. When an ISSU is initiated, it updates the following components on the system as needed:

- Kickstart image
- Supervisor BIOS
- System image
- Fabric extender (FEX) image
- I/O module BIOS and image

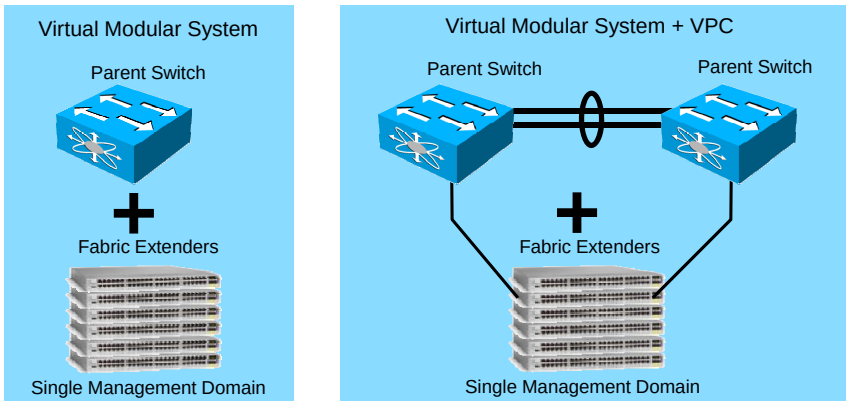
Following data center network design best practices should help ensure switch-level redundancy and continuous service to the attached servers. Such practices include dual-homing servers for redundancy with redundant switches at the access layer.

Once initiated, the ISSU installer service begins the ISSU cycle. The upgrade process is composed of several phased stages that are designed to minimize overall system impact with no impact to data traffic forwarding.

https://t.me/learningnets

In-Service Software Upgrade (Cont.)

- Cisco NX-OS Software is automatically downloaded and upgraded from the parent switch to each fabric extender.
- The control plane will be brought online within 80 seconds



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

When an ISSU is performed on the parent switch, the Cisco NX-OS Software automatically downloads and upgrades each of the attached Cisco Nexus 2000 Series Fabric Extenders. This process results in each Cisco Nexus 2000 Series Fabric Extender being upgraded to the Cisco NX-OS Software release of the parent switch after the parent switch ISSU completes.

The Cisco Nexus 5000 Series Switches have only a single supervisor. The ISSU should be hitless for the switch, its modules, and attached FEXs while the ISSU upgrades the system, kickstart, and BIOS images.

During the ISSU, control plane functions of the switch that is undergoing ISSU are temporarily suspended and configuration changes are disallowed. The control plane is brought online again within 80 seconds to allow protocol communications to resume.

In FEX virtual port channel (vPC) configurations, the primary switch is responsible for upgrading the FEX. The peer switch is responsible for holding onto its state until the ISSU process completes. When one peer switch is undergoing an ISSU, the other peer switch locks the configuration until the ISSU is completed.

When a Layer 3 license is installed, the switches in the Cisco Nexus 5500 Platform do not support an ISSU.

Hot-swapping a Layer 3 module in the Cisco Nexus 5000 Series switch is not supported during an ISSU operation.

Note	Topology must be stable during ISSU.

Additionally, always refer to the latest release notes for ISSU guidelines.

https://t.me/learningnets

ISSU Impact

```
N5K-A#  show install all impact kickstart n5000-uk9-kickstart.5.1.3.N1.1.bin
     system n5000-uk9.5.1.3.N1.1.bin

Verifying image bootflash:/n5000-uk9-kickstart.5.1.3.N1.1.bin for boot variable
"kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n5000-uk9.5.1.3.N1.1.bin for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[###########         ]  50%
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n5000-uk9.5.1.3.N1.1.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/n5000-uk9-kickstart.5.1.3.N
1.1.bin.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n5000-uk9.5.1.3.N1.1.bin.
[####################] 100% -- SUCCESS
```

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

After you acquire, download, and copy the new Cisco NX-OS Software files to your upgrade location, which is typically the bootflash, issue the **show install all impact** command.

ISSU Impact (Cont.)

show install all impact output continued

```
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             New-Version  Upg-Required
------  ----------  ----------------------  ----------------------  ------------
     1      system             5.1(3)N1(1)             5.1(3)N1(1)            no
     1   kickstart             5.1(3)N1(1)             5.1(3)N1(1)            no
     1        bios      v3.5.0(02/03/2011)      v3.5.0(02/03/2011)            no
     1      SFP-uC                 v1.0.0.0                v1.0.0.0            no
     1   power-seq                    v1.0                    v1.0            no
     3   power-seq                    v2.0                    v2.0            no
     1          uC                 v1.2.0.1                v1.2.0.1            no
```

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

This command allows the administrator to determine which components of the system will be affected by the upgrade before performing the upgrade of the Cisco NX-OS Software.

https://t.me/learningnets

Verifying the Files on the Cisco Nexus 5000

Ensure that the upgrade files are resident on the bootflash.

```
N5K-A# dir
       3246    Jan 31 02:17:32 2009  duck1
       3341    Feb 02 14:53:50 2009  duck2
        503    Jan 05 08:18:12 2012  license_SSI154207FW_9.lic
       3457    Feb 04 20:06:42 2009  mts.log
   34342400    Jan 05 08:11:30 2012  n5000-uk9-kickstart.5.1.3.N1.1.bin
  147395647    Jan 05 08:12:34 2012  n5000-uk9.5.1.3.N1.1.bin
       4096    Jan 01 08:15:26 2009  vdc_2/
       4096    Jan 01 08:15:26 2009  vdc_3/
       4096    Jan 01 08:15:26 2009  vdc_4/

Usage for bootflash://
  299311104 bytes used
 1349312512 bytes free
 1648623616 bytes total
N5K-A#
```

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

Use the **dir** command to verify that the required Cisco NX-OS files are present on the bootflash of the switch to be upgraded (or downgraded). Note that each Cisco NX-OS Software version has two separate files—a kickstart file and the system file—that must have the same version.

The kickstart image is contained in an independent file. The system image, BIOS, and code for the modules, including the FEX, are all contained in a single system file.

There can be multiple versions of Cisco NX-OS Software on the bootflash in addition to other files.

Note	By default, the kickstart image contains the word *kickstart* in the filename. The system file does not contain the word *system*.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying the Files on the Cisco Nexus 7000

- Before you perform an ISSU, check active and standby supervisors to ensure that there is enough room on bootflash to hold the files.
- Before you perform an ISSU, check the bootflash on active and standby supervisor to verify that your files are in place.

```
N7K-1# dir bootflash:
        315     Oct 11 00:34:26 2008  MDS20081010111742129.lic
        273     Jan 15 05:38:38 2008  TBM12234289_115853568.lic
      16384     Jan 15 05:18:37 2008  lost+found/
   87081184     Jan 15 05:38:59 2008  n7000-s1-dk9.4.1.3.bin
   92054694     Feb 12 00:13:06 2009  n7000-s1-dk9.4.1.4.bin
   23584768     Jan 15 05:39:05 2008  n7000-s1-kickstart.4.1.3.bin
   23776768     Feb 12 00:14:32 2009  n7000-s1-kickstart.4.1.4.bin
       4096     Oct 16 21:40:09 2008  vdc_2/
       4096     Oct 17 02:00:06 2008  vdc_3/
       4096     Oct 20 21:29:28 2008  vdc_4/

Usage for bootflash://sup-local
  352321536 bytes used
 1457577984 bytes free
 1809899520 bytes total
```

Old Versions

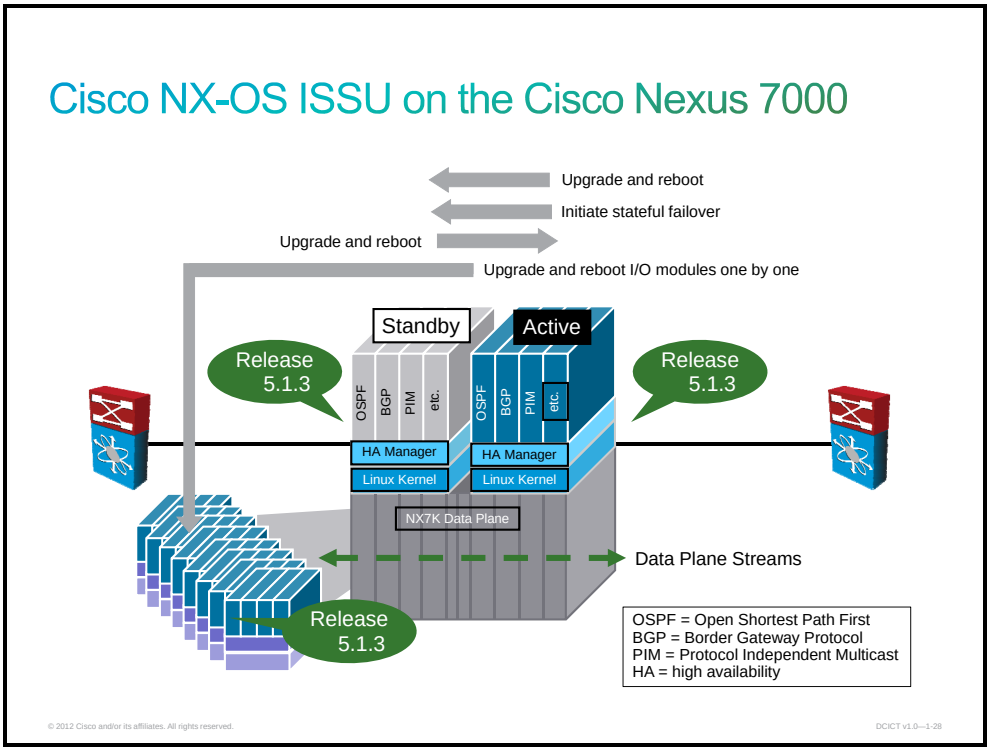
New Versions

- Use **show install all impact** command to verify impact of upgrade.

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

Before you perform a Cisco Nexus ISSU, verify that kickstart and system images are in the bootflash of both supervisor modules.

https://t.me/learningnets



In a Cisco Nexus 7000 Series chassis with dual supervisors, you can use the Cisco NX-OS ISSU feature to upgrade the system software while the system continues to forward traffic. A Cisco NX-OS ISSU uses the existing features of Cisco Nonstop Forwarding (Cisco NSF) with Stateful Switchover (SSO) to perform the software upgrade with no system downtime.

When an ISSU is initiated, the Cisco NX-OS ISSU updates (as needed) the following components on the system:

- Supervisor BIOS, kickstart image, and system image
- Module BIOS and image
- CMP BIOS and image

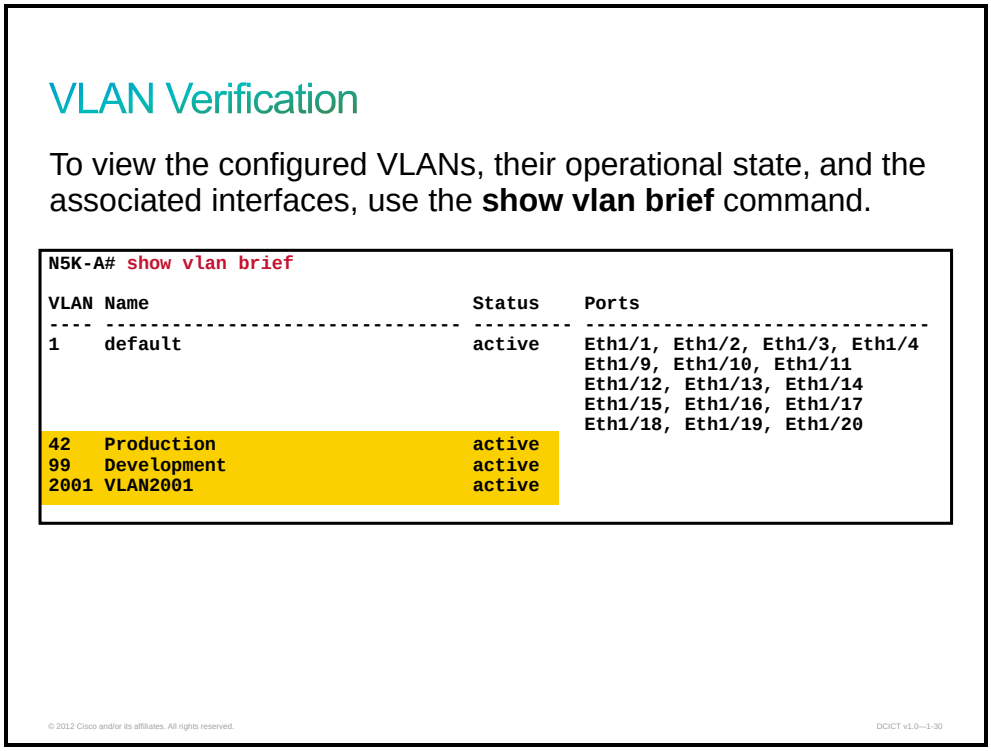
A Cisco NX-OS ISSU has the following limitations and restrictions:

- It does not change any configuration settings or network connections during the upgrade. Any changes in the network settings may cause a disruptive upgrade.
- In some cases, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:
  - A single supervisor system with kickstart or system image changes
  - A dual-supervisor system with incompatible system software images
- Configuration mode is blocked during the Cisco ISSU to prevent any changes.

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

Verifying VLANs

This topic explains how to verify VLANs on the Cisco Nexus 7000 and 5000 Series Switches.



To view the configured VLANs, their operational states, and the associated interfaces, use the **show vlan brief** command.

VLANs that have been administratively defined are listed in addition to the default VLAN. If the VLAN was created successfully and not disabled for any reason, it shows a status of active.

All ports are associated to the default VLAN until they are defined by the administrator as participating in another VLAN.

https://t.me/learningnets

VLAN Verification (Cont.)

The Cisco Nexus switches support up to 4094 VLANs in each VDC in accordance with the IEEE 802.1Q standard.

- 80 VLANs in the high end of the VLAN range are reserved for internal use by the system and cannot be used.

```
Nexus-7010-PROD-A# show vlan internal usage

VLANs                DESCRIPTION
-------------------  -----------------
3968-4031            Multicast
4032-4035,4048-4059  Online Diagnostic
4036-4039,4060-4087  ERSPAN
4042                 Satellite
3968-4095            Current
Nexus-7010-PROD-A#
```

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

A switch port belongs to a VLAN. Unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Consider each VLAN to be a logical network. Packets that are destined for a station that does not belong to the same VLAN must be forwarded through a router.

This table describes the VLAN ranges.

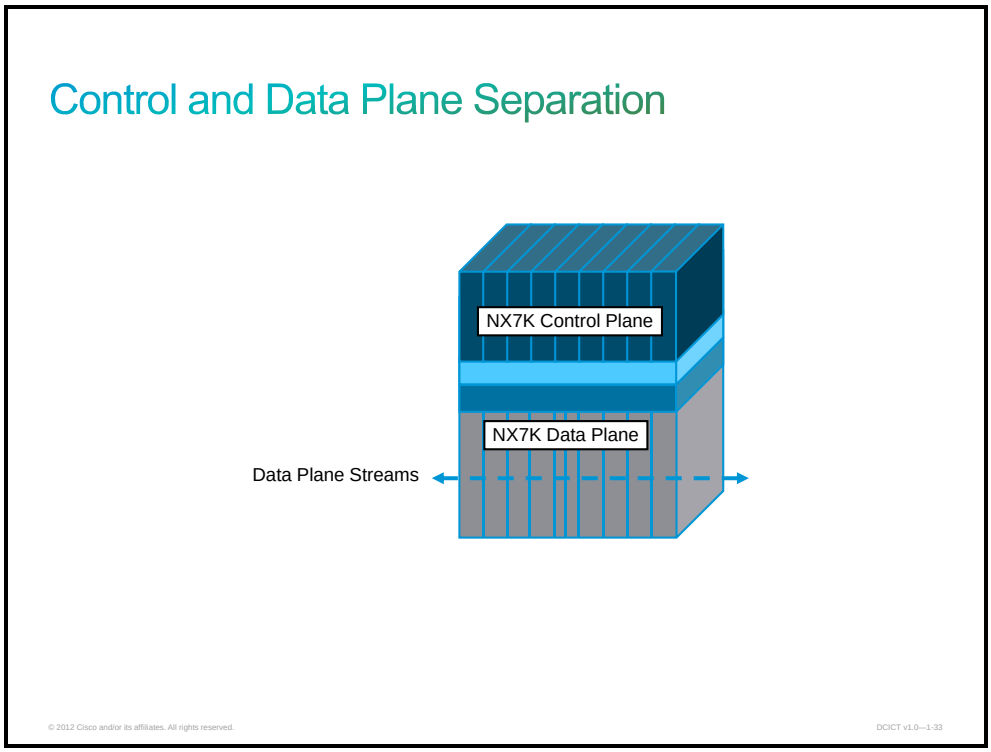
VLAN Numbers	Range	Use
1	Normal	Cisco default. You can use this VLAN, but you cannot modify or delete it.
2–1005	Normal	You can create, use, modify, and delete these VLANs.
1006–3967 and 4048–4093	Extended	You can create, name, and use these VLANs. You cannot change the following parameters: ▪ The state is always active. ▪ The VLAN is always enabled. You cannot shut down these VLANS.
3968–4047 and 4094	Internally allocated	These 80 VLANs and VLAN 4094 are allocated for internal device use. You cannot create, delete, or modify any VLANs within the block that is reserved for internal use.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

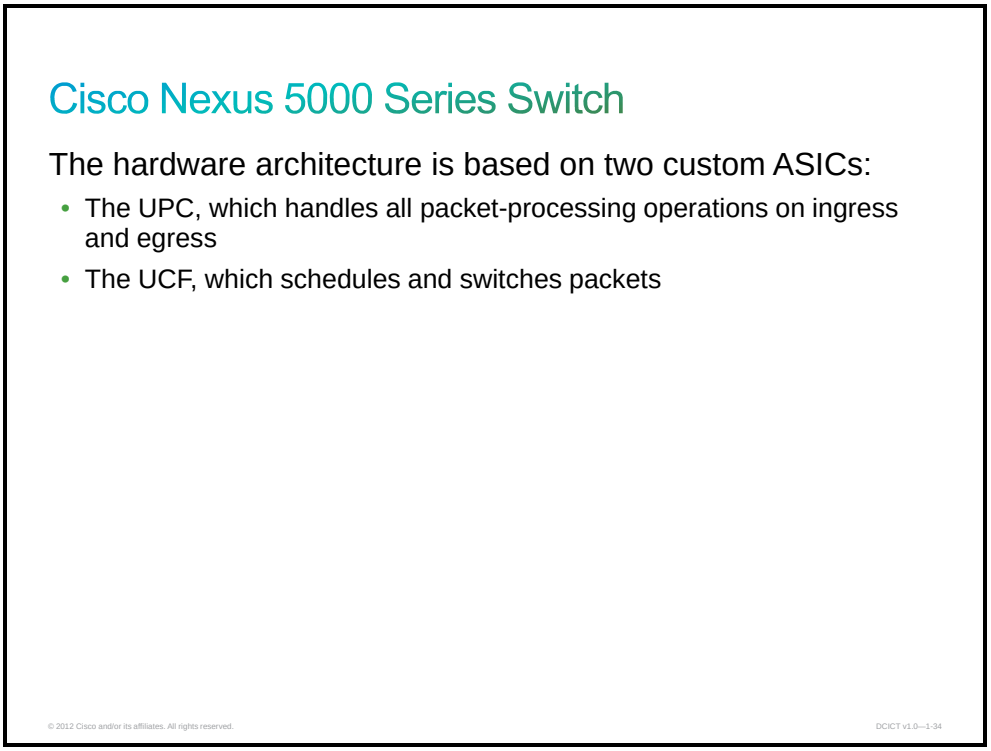
Examining the Operational Plane

This topic describes control, management, and data planes on Nexus 7000 and 5000 Series Switches.



Cisco NX-OS Software provides isolation between control and data forwarding planes within the device. This isolation means that a disruption within one plane does not disrupt the other.

https://t.me/learningnets



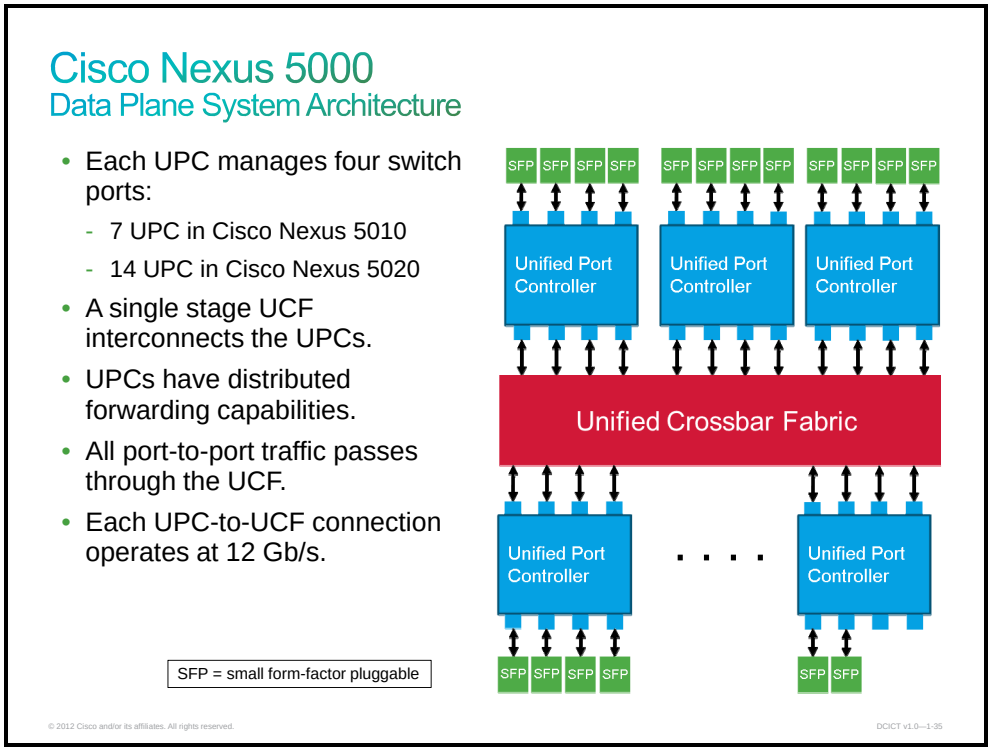
The Cisco Nexus 5000 Series Switches use a scalable cut-through input queuing switching architecture. The architecture is implemented primarily by two ASICs developed by Cisco:

- A set of unified port controllers (UPCs) that perform data plane processing
- A unified crossbar fabric (UCF) that cross-connects the UPCs

The UPC ASIC manages all packet-processing operations on ingress and egress interfaces. UPCs provide distributed packet forwarding capabilities. Four switch ports are managed by each UPC.

A single-stage UCF ASIC interconnects all UPCs. The UCF schedules and switches packets from ingress to egress UPCs and is simply known as “the fabric.” All port-to-port traffic passes through the UCF.

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



The figure represents the architectural diagram of the Cisco Nexus 5000 Series Switches data plane. It describes a distributed forwarding architecture.

Each UPC manages four 10 Gigabit Ethernet ports and makes forwarding decisions for the packets that are received on those ports. After a forwarding decision is made, the packets are queued in virtual output queues (VOQs) where they wait to be granted access to the UCF. Because of the cut-through characteristics of the architecture, packets are queued and dequeued before the full packet contents have been received and buffered on the ingress port. The UCF is responsible for coupling ingress UPCs to available egress UPCs. The UCF internally connects each 10 Gigabit Ethernet, FCoE-capable interface through fabric interfaces running at 12 Gb/s. This 20 percent overspeed helps ensure line-rate throughput regardless of the packet manipulation that is performed in the ASICs.

Unified Port Controller

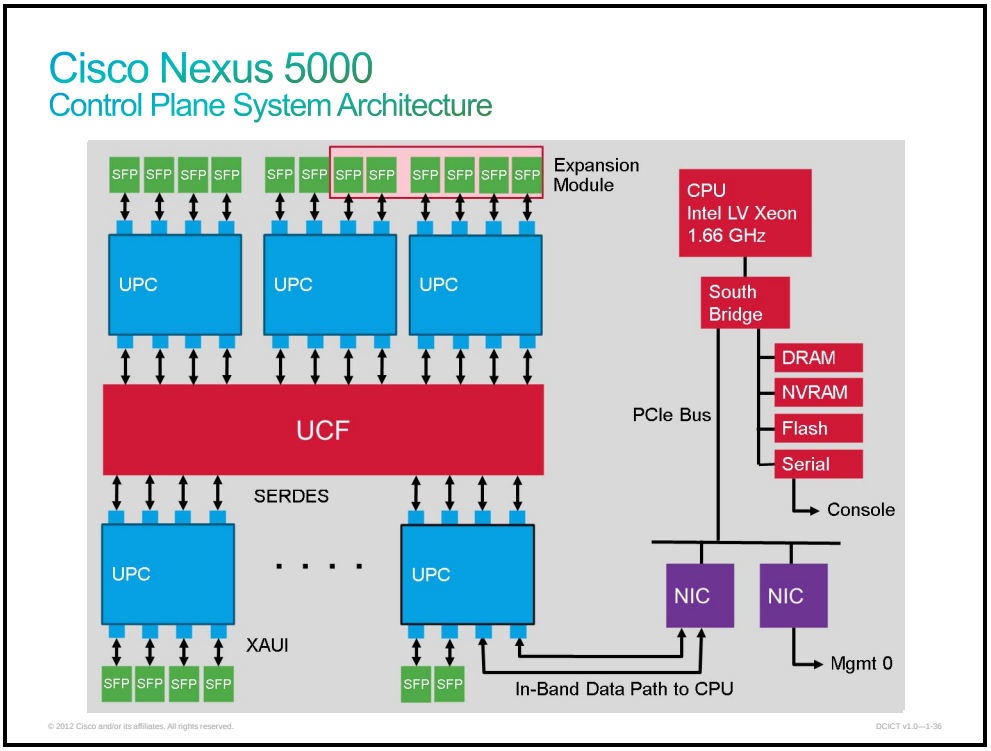
The UPC manages all packet-processing operations within the Cisco Nexus 5000 Series Switches. It is a Layer 2 multipathing (L2MP) device with the capability to operate simultaneously and at wire speed with the following protocols:

- Classic Ethernet
- Fibre Channel
- FCoE

On the ingress side, UPC manages the physical details of different media as it maps the received packets to a unified internal packet format. The UPC then makes forwarding decisions that are based on protocol-specific forwarding tables that are stored locally in the ASIC. On the egress side, the UPC remaps the unified internal format to the format that is supported by the egress medium and Layer 2 protocol and transmits the packet.

Each external-facing 10-Gb/s interface on a UPC can be wired to serve as two Fibre Channel interfaces at 1, 2, and 4 Gb/s for an expansion module. Therefore, a single UPC can connect up to eight Fibre Channel interfaces through expansion modules.

https://t.me/learningnets



The figure represents the architectural diagram of the Cisco Nexus 5000 Series Switches control plane.

In the control plane, the Cisco Nexus 5000 Series Switch runs the Cisco NX-OS Software on a single-core 1.66-GHz Intel LV Xeon CPU with 2 GB of DRAM. The supervisor complex is connected to the data plane in-band through two internal ports running 1-Gb/s Ethernet.

The system is managed in-band, through the out-of-band 10-, 100-, and 1000-Mb/s management port or via the serial console port.

The control plane is responsible for managing all control traffic. Data frames bypass the control plane and are managed by the UCF and the UPCs respectively. Bridge protocol data units (BPDUs), fabric login (FLOGI) frames and other control protocol-related frames are managed by the control plane supervisor. Additionally, the control plane is responsible for running the Cisco NX-OS Software.

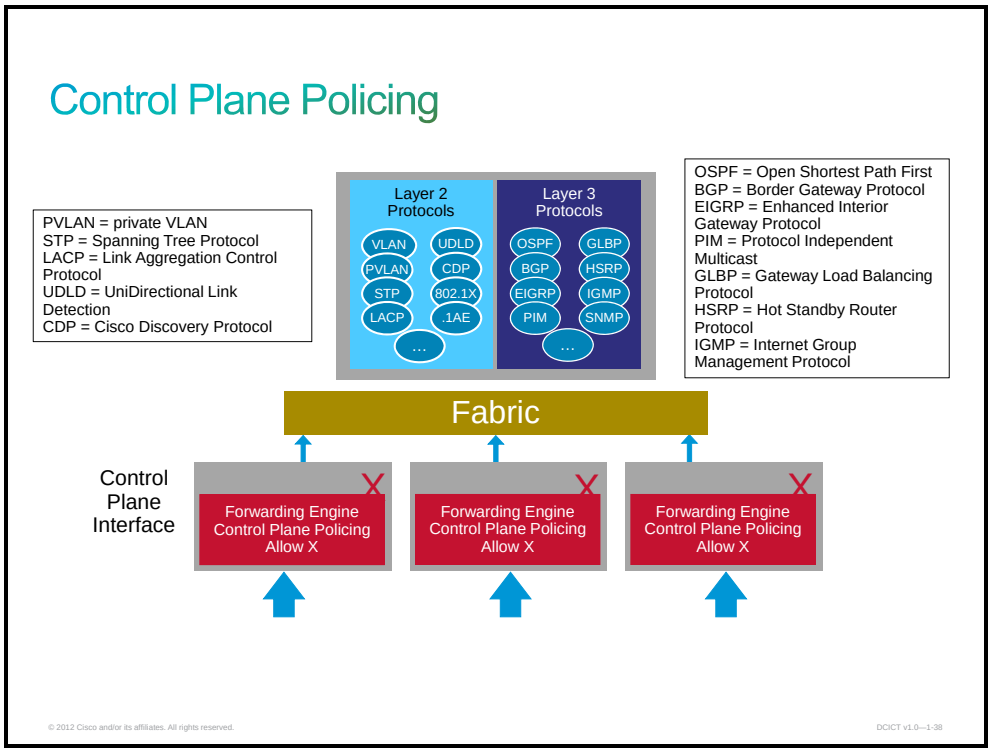
The control plane consists of the following:

- Single-core 1.66-GHz Intel LV Xeon CPU
- Serializer/deserializer (SERDES)
- Intel South Bridge ASIC that controls memory, flash, and serial console access
- PCI Express (PCIe) bus
- Network interface card (NIC)
- Memory
  - 2-GB Double Date Rate (DDR) 400
  - 1-GB USB Flash
  - 2-MB BIOS
  - 2-MB of NVRAM

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

Reviewing Cisco NX-OS Default Control Plane Policing

This topic describes Cisco NX-OS default Control Plane Policing (CoPP) on Nexus 7000 and 5000 Series Switches.



CoPP protects the control plane and separates it from the data plane.

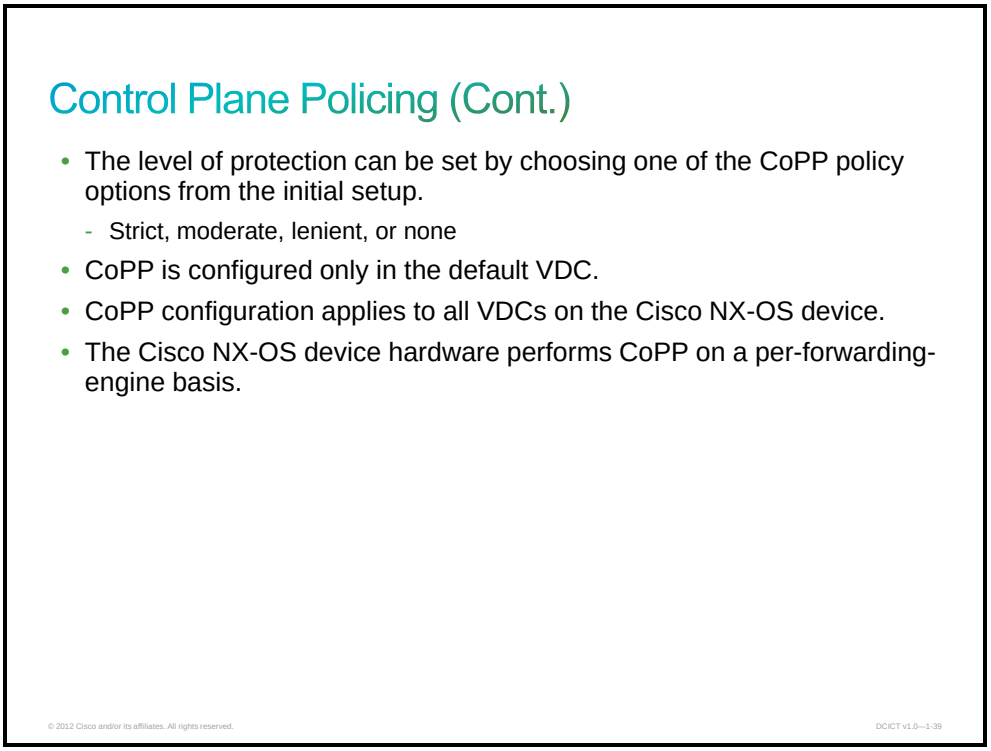
The Cisco NX-OS device provides CoPP to prevent denial of service (DoS) attacks from impacting performance. Such attacks, which can be perpetrated either inadvertently or maliciously, typically involve high rates of traffic targeting the route processor itself.

The supervisor module divides the traffic that it manages into three functional components or planes:

- **Data plane**: Manages all the data traffic.
- **Control plane**: Manages all routing protocol control traffic. These packets are destined to router addresses and are called control plane packets.
- **Management plane**: Runs the components that are meant for Cisco NX-OS device management purposes such as the CLI and SNMP.

The supervisor module has both the management plane and control plane and is critical to the operation of the network. Any disruption or attacks to the supervisor module may result in serious network outages.

https://t.me/learningnets



Default Policing Policies

The setup utility allows building an initial configuration file using the system configuration dialog. After the initial setup, the Cisco NX-OS Software installs the default CoPP-system-policy policy to protect the supervisor module from DoS attacks. The level of protection can be set by choosing one of the following CoPP policy options from the initial setup:

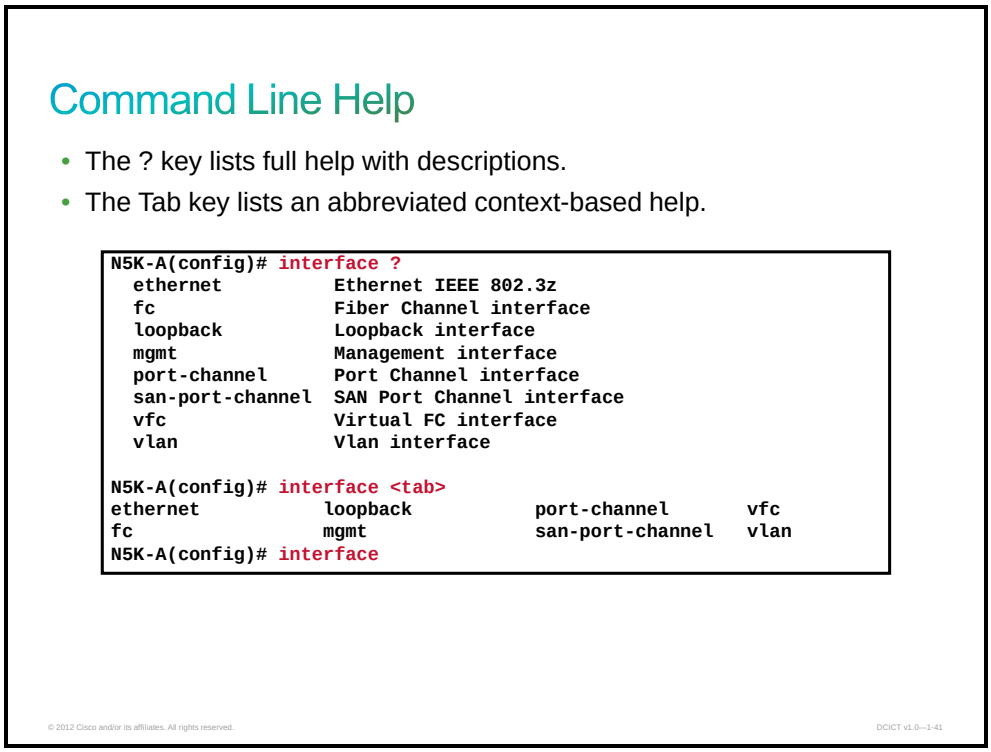
- **Strict:** This policy is 1 rate and 2 color and has a committed burst size (Bc) value of 250 microsec (except for the important class, which has a Bc value of 1000 microsec).
- **Moderate:** This policy is 1 rate and 2 color and has a Bc value of 310 microsec (except for the important class, which has a Bc value of 1250 microsec). These values are 25 percent greater than the strict policy.
- **Lenient:** This policy is 1 rate and 2 color and has a Bc value of 375 microsec (except for the important class, which has a Bc value of 1500 microsec). These values are 50 percent greater than the strict policy.
- **None:** No control plane policy is applied.

The Cisco NX-OS device hardware performs CoPP on a per-forwarding-engine basis. CoPP does not support distributed policing. Therefore, you should choose rates so that the aggregate traffic does not overwhelm the supervisor module.

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

Using Important CLI Commands

This topic describes important CLI commands on the Cisco Nexus 7000 and 5000 Series Switches.



For CLI help, pressing the **?** key displays full parser help with accompanying descriptions. Pressing the Tab key displays a brief list of the commands that are available at the current branch, but there are no accompanying descriptions.

https://t.me/learningnets

Command Completion

- The Tab key still provides command completion.
- Reserved names can also be tab completed.

```
N5K-A# conf
Enter configuration commands, one per line.  End with CNTL/Z.
N5K-A(config)# vrf context ?
  WORD                   VRF name (Max Size 32)
  management (no abbrev)  Configurable VRF name

N5K-A(config)# vrf context man<tab>
N5K-A(config)# vrf context management
N5K-A(config-vrf)#
```

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

Pressing the Tab key provides command completion to a unique but partially typed command string. In the Cisco NX-OS Software, command completion may also be used to complete a system-defined name and, in some cases, user-defined names.

In the figure, the Tab key is used to complete the name of a VRF on the command line.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Command Line

- The Ctrl-Z key combination exits configuration mode but retains the user input on the command line.

```
N5K-A(config-if)# show <CTRL-Z>
N5K-A# show
```

- The **where** command displays the current context and login credentials.

```
N5K-A(config-if)# where
  conf; interface Ethernet1/19      admin@N5K-A
N5K-A(config-if)#
```

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

The Ctrl-Z key combination lets you exit configuration mode but retains the user input on the command line.

The **where** command displays the current configuration context and login credentials of the user.

show running-config

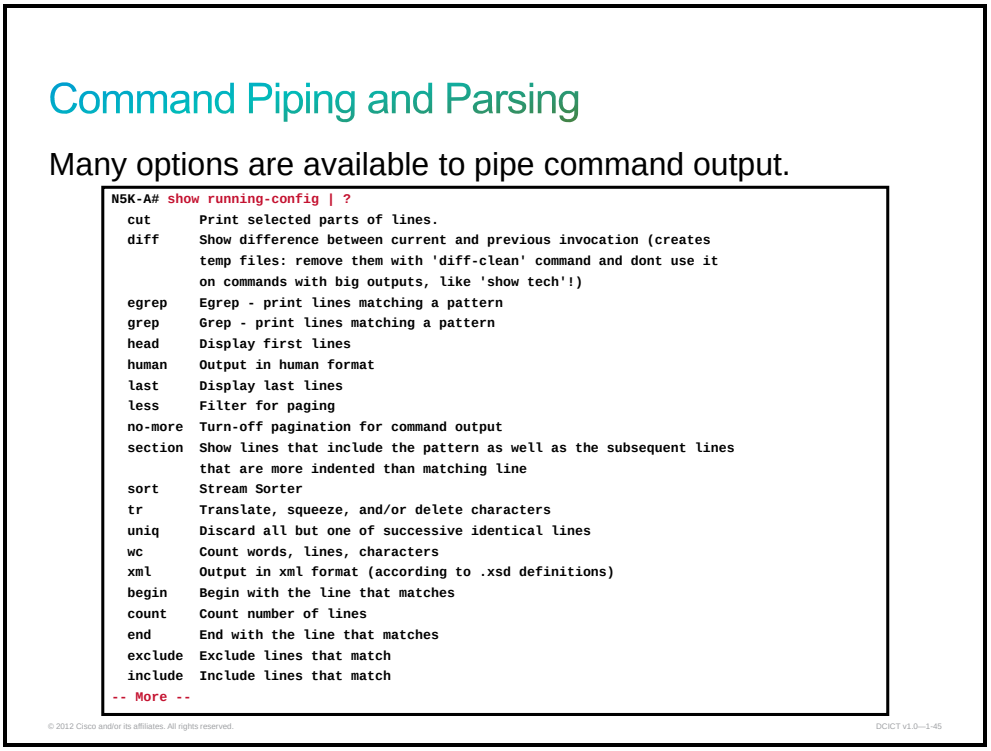
A subset of the **running-config** can be displayed.

```
N5K-A# show running-config ?
  <CR>
  >                       Redirect it to a file
  >>                      Redirect it to a file in append mode
  aaa                     Display aaa configuration
  aclmgr                  Show running config for aclmgr
  adjmgr                  Display adjmgr information
  all                     Current operating configuration with defaults
  arp                     Display arp information
  assoc                   Original ID to Translated ID Association
  callhome                Display callhome configuration
  cdp                     Display cdp configuration
  cert-enroll             Display certificates configuration
  cfs                     Display cfs configurations
  copp                    Show running config for copp
  diagnostic              Display diagnostic information
  diff                    Show the difference between running and startup
                          configuration
  exclude                 Exclude running configuration of specified features
  exclude-provision       Hide config for offline pre-provisioned interfaces
  expand-port-profile     Expand port profile
  fcoe_mgr                Display fcoe_mgr configuration
  icmpv6                  Display icmpv6 information
  igmp                    Display igmp information
-- More --
```

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

The **show running-config** command displays the current running configuration of the switch. If you need to view a specific features portion of the configuration, you can do so by running the command with the feature specified.

https://t.me/learningnets



The pipe parameter can be placed at the end of all CLI commands to qualify, redirect, or filter the output of the command. Many advanced pipe options are available including the following:

- Get regular expression (grep) and extended get regular expression (egrep)
- less
- no-more
- head and tail

The pipe option can also be used multiple times or recursively.

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

Command Piping and Parsing (Cont.)

Advanced grep with multiple qualifiers is allowed.

```
N5K-A# show spanning-tree | egrep ?
  WORD          Search for the expression
  count         Print a total count of matching lines only
  ignore-case   Ignore case difference when comparing strings
  invert-match  Print only lines that contain no matches for <expr>
  line-exp      Print only lines where the match is a whole line
  line-number   Print each match preceded by its line number
  next          Print <num> lines of context after every matching line
  prev          Print <num> lines of context before every matching line
  word-exp      Print only lines where the match is a complete word

N5K-A# show spanning-tree | egrep ignore-case next 4 rstp
  Spanning tree enabled protocol rstp
  Root ID    Priority    24577
             Address     64a0.e743.03c2
             Cost        2
             Port        129 (Ethernet1/1)
```

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

The egrep parameter allows you to search for a word or expression and print a word count. It also has other options.

In the figure, the **show spanning-tree | egrep ignore-case next 4 rstp** command is used. The **show spanning-tree** command output is piped to **egrep**, where the qualifiers indicate to display the lines that match the regular expression **rstp**, and to not be case sensitive. When a match is found, the **next 4** lines following the match of **rstp** are also to be displayed.

https://t.me/learningnets

More Options

```
--More--
Most commands optionally preceded by integer argument k.  Defaults in brackets.
Star (*) indicates argument becomes new default.
-------------------------------------------------------------------------------
<space>                 Display next k lines of text [current screen size]
z                       Display next k lines of text [current screen size]*
<return>                Display next k lines of text [1]*
d or ctrl-D             Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt>   Exit from more
s                       Skip forward k lines of text [1]
f                       Skip forward k screenfuls of text [1]
b or ctrl-B             Skip backwards k screenfuls of text [1]
'                       Go to place where previous search started
=                       Display current line number
/<regular expression>   Search for kth occurrence of regular expression [1]
n                       Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd>       Execute <cmd> in a subshell
v                       Start up /usr/bin/vi at current line
ctrl-L                  Redraw screen
:n                      Go to kth next file [1]
:p                      Go to kth previous file [1]
:f                      Display current file name and line number
.                       Repeat previous command
```

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

When a command output exceeds the length of the screen, the --More-- prompt is displayed.

When a --More-- prompt is presented, press **h** to see a list of available options.

To break out of the **more** command, press the letter **Q**.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

show running-config (Cont.)

Only the QoS subset of the **running-config** command

```
N5K-A# show running-config ipqos

!Command: show running-config ipqos
!Time: Tue Feb 17 14:40:22 2009

version 5.1(3)N1(1)
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
system qos
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type qos input fcoe-default-in-policy
  service-policy type network-qos fcoe-default-nq-policy
```

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

The figure shows the output of the **show running-config ipqos** command. This command displays only the portion of the running-config that relates to the ipqos configuration.

show running-config (Cont.)

Only the QoS subset of the **running-config** command including the default parameters

```
N5K-A# show running-config ipqos all

!Command: show running-config ipqos all
!Time: Tue Feb 17 14:41:42 2009

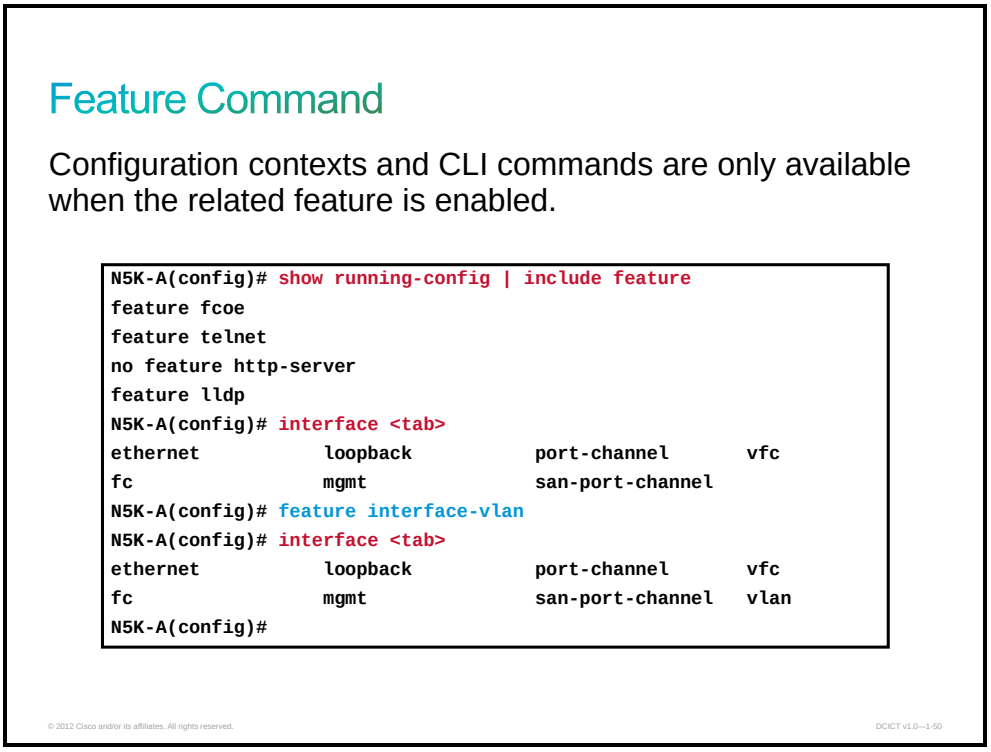
version 5.1(3)N1(1)
class-map type qos class-fcoe
class-map type qos match-any class-default
class-map type qos match-any class-all-flood
class-map type qos match-any class-ip-multicast
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-default
  match qos-group 0
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos default-in-policy
  class class-default
    set qos-group 0
policy-map type qos fcoe-default-in-policy
  class class-fcoe
    set qos-group 1
--More--
```

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

The figure shows the output of the **show running-config ipqos all** command. The addition of the **all** parameter to this command now displays the portion of the running-config that relates to the ipqos configuration. In addition, default values that are not normally visible in the configuration appear.

https://t.me/learningnets

The **all** parameter can be used on the **show running-config** command with or without another parameter. For example, the **show running-config all** command would display the entire running-configuration with all default values included.



The following types of interfaces are supported via the main chassis of the switch:

- **Physical**: 1 and 10 Gigabit Ethernet, and 1-, 2-, 4-, and 8 Gb Fibre Channel via expansion modules. Additional interface types may be present depending on the FEX model that is connected.
- **Logical**: Ethernet and SAN port channels, SVI, virtual Fibre Channel (vFC), and tunnels.

Because of the **feature** command on the Cisco Nexus platform, interfaces, protocols, and their related configuration contexts are not visible until the feature is enabled on the switch.

In the figure, the **show running-config | include feature** command is used to determine the features that are currently enabled. In the default configuration, only **feature telnet** and **feature lldp** are enabled. Using the **interface** command, followed by pressing the Tab key for help, only ethernet, mgmt, and ethernet port channel interface types appear in the Cisco NX-OS Software command output.

Enabling **feature interface-vlan** and repeating the **interface** command, SVI interface type has appeared as a valid interface option. Similarly, enabling **feature fcoe** makes available all Fibre Channel interface types—Fibre Channel, SAN Port Channel, and vFC.

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

Verifying Interfaces

The **show interface brief** command provides a summary of the interface status of all interfaces on the switch and lists the most relevant interface parameters.

```
Nexus-7010-PROD-A# show interface brief

--------------------------------------------------------------------------------
Port   VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0  --           up     172.16.1.201                            1000     1500

--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Por
t
Interface                                                                    Ch
#
--------------------------------------------------------------------------------
Eth2/1        1       eth  access down    Link not connected         auto(D) --
Eth2/2        1       eth  access down    Link not connected         auto(D) --
Eth2/3        1       eth  access down    Link not connected         auto(D) --
Eth2/4        1       eth  access down    Link not connected         auto(D) --
Eth2/5        1       eth  access up      none                        10G(D) --
Eth2/6        1       eth  access down    Link not connected         auto(D) --
Eth2/7        1       eth  access up      none                        10G(D) --
Eth2/8        1       eth  access up      none                        10G(D) --
Eth2/9        1       eth  access down    SFP not inserted           auto(D) --
Eth2/10       1       eth  access down    SFP not inserted           auto(D) --
Eth2/11       1       eth  access up      none                        10G(D) --
Eth2/12       1       eth  access up      none                        10G(D) --
Eth3/1        --      eth  routed down    Administratively down      auto(S) --
Eth3/3        --      eth  routed down    Administratively down      auto(S) --
Eth3/5        --      eth  routed down    SFP not inserted           auto(S) --
---more---
```

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

The **show interface brief** command lists the most relevant interface parameters for all interfaces on a Cisco Nexus Series switch. For interfaces that are down, it also lists the reason that the interface is down. Information about speed and rate mode (dedicated or shared) is also displayed.

https://t.me/learningnets

Cisco NX-OS Interfaces

- Cisco NX-OS supports various types of interfaces.
- All physical Ethernet interfaces on a Cisco Nexus switch are designated as **interface ethernet *slot/port*** regardless of interface type and speed.

```
Nexus-7010-PROD-A# show interface ethernet 2/11
Ethernet2/11 is up
  Dedicated Interface
  Hardware: 1000/10000 Ethernet, address: c464.13ba.fe92 (bia c464.13ba.fe92)
  Description: To N7K Prod_B int Eth 2/23
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  Port mode is access
  full-duplex, 10 Gb/s, media type is 10G
  Beacon is turned off
  Auto-Negotiation is turned on
  Input flow-control is off, output flow-control is off
  Rate mode is dedicated
  Switchport monitor is off
  EtherType is 0x8100
  Last link flapped 1d18h
  Last clearing of "show interface" counters 1d20h
  30 seconds input rate 72 bits/sec, 0 packets/sec
  30 seconds output rate 304 bits/sec, 0 packets/sec
---more---
```

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

Cisco NX-OS Software supports the following types of interfaces:

- **Physical:** Ethernet (10/100/1000/10G)
- **Logical:** port channel, loopback, null, SVI, tunnels, subinterfaces
- **Management:** Management

All Ethernet interfaces are named Ethernet. There is no differentiation in the naming convention for the different speeds.

The **show interface** command displays the operational state of any interface, including the reason why that interface might be down.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Default Interface Settings

The default settings for any type of interface may be displayed.

```
N5K-A# show running-config interface ethernet 1/1 all

!Command: show running-config interface Ethernet1/1 all
!Time: Tue Feb 17 15:19:36 2009

version 5.1(3)N1(1)

interface Ethernet1/1
  description To N7K Eth 2/5
  priority-flow-control mode auto
  lldp transmit
  lldp receive
  no switchport block unicast
  no switchport block multicast
  no hardware multicast hw-hash
  no hardware vethernet mac filtering per-vlan
  cdp enable
  switchport
  switchport mode access
  no switchport monitor
  no switchport dot1q ethertype
  no switchport priority extend
  spanning-tree port-priority 128
  spanning-tree cost auto
--More--
```

Includes default parameters

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

Interfaces on the Cisco Nexus Series switches have many default settings that are not visible in the **show running-config** command. The default configuration parameters for any interface may be displayed using the **show running-config interface all** command.

https://t.me/learningnets

Verifying Modules

Module 1 is the main chassis.

- Expansion modules and FEX may also be present.

```
N5K-A# show module
Mod Ports  Module-Type                       Model                 Status
--- -----  -------------------------------- --------------------- ------------
1    32     O2 32X10GE/Modular Universal Pla N5K-C5548UP-SUP        active *
3    0      O2 Non L3 Daughter Card          N55-DL2                ok

Mod  Sw              Hw      World-Wide-Name(s) (WWN)
---  --------------  ------  --------------------------------------------------
1    5.1(3)N1(1)     1.0     20:01:54:7f:ee:5c:6e:c0 to 20:04:54:7f:ee:5c:6e:c0
3    5.1(3)N1(1)     1.0     --

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    547f.ee5c.6ea8 to 547f.ee5c.6ec7         FOC154849HZ
3    0000.0000.0000 to 0000.0000.000f         FOC15475JZJ
N5K-A#
```

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

The **show module** command provides a hardware inventory of the interfaces and module that are installed in the chassis. The main system is referred to as Module 1 on the Cisco Nexus 5000 Series Switch. Depending on the model of the switch, the number of ports that are associated with this module will vary. This number is visible under the Module-Type column.

Any expansion modules are also be listed. The number and type of expansion modules vary by system. When FEXs are attached, each FEX is also considered an expansion module of the main system.

Modules that are inserted into the expansion module slots will be numbered according to which slot they are inserted in. In the figure, the 8-port Fibre Channel Module is installed in slot number 2.

Fabric extender slot numbers are assigned by the administrator during FEX configuration.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying Modules (Cont.)

```
Nexus-7010-PROD-A# show module
Mod  Ports  Module-Type                          Model              Status
---  -----  ---------------------------------- ------------------ ----------
2    32     1/10 Gbps Ethernet Module           N7K-F132XP-15      ok
3    32     10 Gbps Ethernet XL Module          N7K-M132XP-12L     ok
5    0      Supervisor module-1X                N7K-SUP1           ha-standby
6    0      Supervisor module-1X                N7K-SUP1           active *

---output removed---

Nexus-7010-PROD-A# conf
Enter configuration commands, one per line.  End with CNTL/Z.
Nexus-7010-PROD-A(config)# show module
Mod  Ports  Module-Type                          Model              Status
---  -----  ---------------------------------- ------------------ ----------
2    32     1/10 Gbps Ethernet Module           N7K-F132XP-15      ok
3    32     10 Gbps Ethernet XL Module          N7K-M132XP-12L     ok
5    0      Supervisor module-1X                N7K-SUP1           ha-standby
6    0      Supervisor module-1X                N7K-SUP1           active *

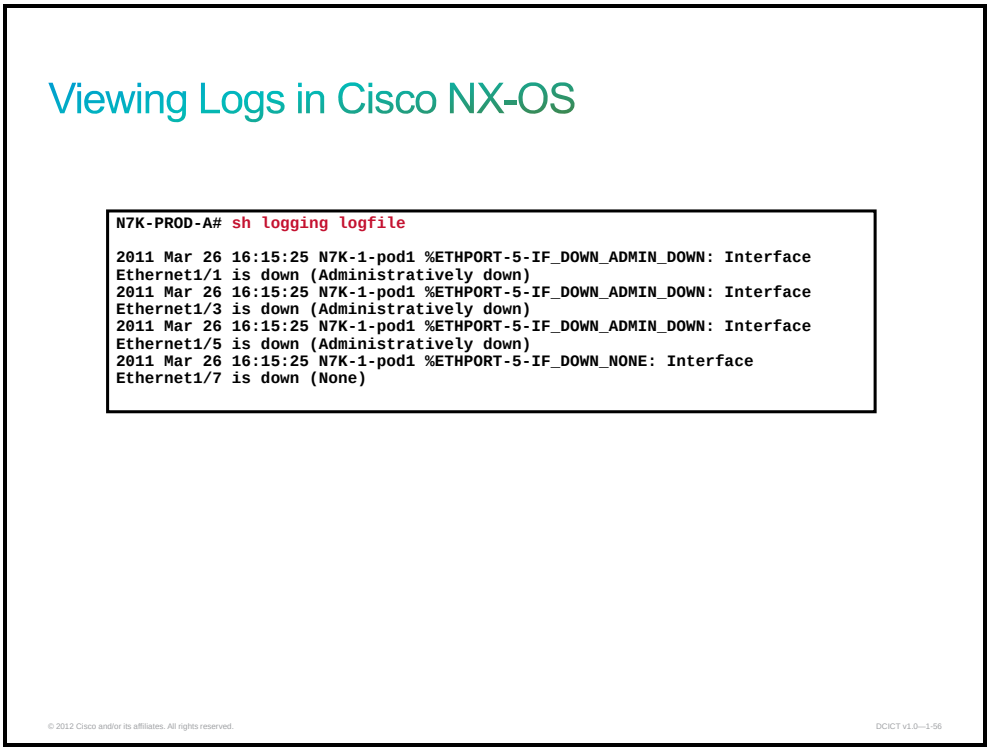
Mod  Sw              Hw
---  --------------  ------
2    6.0(2)          1.2
3    6.0(2)          1.3
5    6.0(2)          2.1
6    6.0(2)          1.8


---output removed---
```

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

The **show module** command on the Cisco Nexus 7000 Series Switches identifies which modules are inserted and operational.

https://t.me/learningnets



By default, the device outputs messages to terminal sessions and to a log file.

You can display or clear messages in the log file. Use the following commands:

- The **show logging last number-lines** command displays the last number of lines in the logging file.
- The **show logging logfile** [***start-time yyyy mmm dd hh:mm:ss***] [***end-time yyyy mmm dd hh:mm:ss***] command displays the messages in the log file that have a time stamp within the span entered.
- The **clear logging *logfile*** command clears the contents of the log file.

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

Viewing NVRAM logs

- The most recent 100 system messages that are priority 0, 1, or 2 are logged into NVRAM on the supervisor module.
- After a switch reboots, these messages can be displayed.

```
Nexus-7010-PROD-A# show logging nvram
2011 Dec 20 19:40:09 N7K-1-pod1 %$ VDC-2 %$ %STP-2-ROOTGUARD_CONFIG_CHANGE:
Root
 guard enabled on port Ethernet1/1.
2011 Dec 20 19:40:09 N7K-1-pod1 %$ VDC-2 %$ %STP-2-ROOTGUARD_CONFIG_CHANGE:
Root
 guard enabled on port Ethernet1/3.
2011 Dec 20 19:42:56 N7K-1-pod1 %$ VDC-2 %$ %STP-2-ROOTGUARD_CONFIG_CHANGE:
Root
 guard disabled on port Ethernet1/1.
2011 Dec 20 19:42:56 N7K-1-pod1 %$ VDC-2 %$ %STP-2-ROOTGUARD_CONFIG_CHANGE:
Root
 guard disabled on port Ethernet1/3.
2011 Dec 21 10:10:00 N7K-1-pod1 %$ VDC-2 %$ %STP-2-
VPC_PEERSWITCH_CONFIG_ENABLED
: vPC peer-switch configuration is enabled. Please make sure to configure
spanni
ng tree "bridge" priority as per recommended guidelines to make vPC peer-
switch
operational.
---output removed---
```

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

To view the NVRAM logs, use the **show logging nvram [last number-lines]** command.

The **clear logging nvram** command clears the logged messages in NVRAM.

https://t.me/learningnets

Summary

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

Summary

- Initial connectivity to the Cisco Nexus Series switches is via the console port.
- At the console port, the administrator can perform an initial configuration which includes network connectivity parameters so that a remote connection can be created from a remote PC.
- The Cisco Nexus 7000 Series Switches have an additional management interface called the Connectivity Management Processor.
- The Cisco Nexus Series switches put the management port into a management VRF by default, with SSH enabled and Telnet disabled in the default settings.
- The Cisco Nexus Series switch provides ISSU capabilities for seamless upgrades of the operating system.
- The Cisco Nexus Series switches use VLANs for separating traffic at a Layer 2 level.

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

Summary (Cont.)

- To provide maximum availability, the control, management, and data planes on the Cisco Nexus Series switches are separated.
- CoPP is Control Plane Policing on the Cisco Nexus Series switches, and is used to restrict control plane traffic from overpowering the CPU.
- There are various **show** commands that should be used when you are verifying and monitoring the Cisco Nexus 7000 and 5000 Series Switches.

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

https://t.me/learningnets

Lesson 5

Describing vPCs and Cisco FabricPath in the Data Center

Overview

Cisco FabricPath and virtual port channels (vPCs) can be used to provide a Layer 2 multipathing solution for the data center. This lesson looks at the benefits of these two solutions and how they are implemented.

Objectives

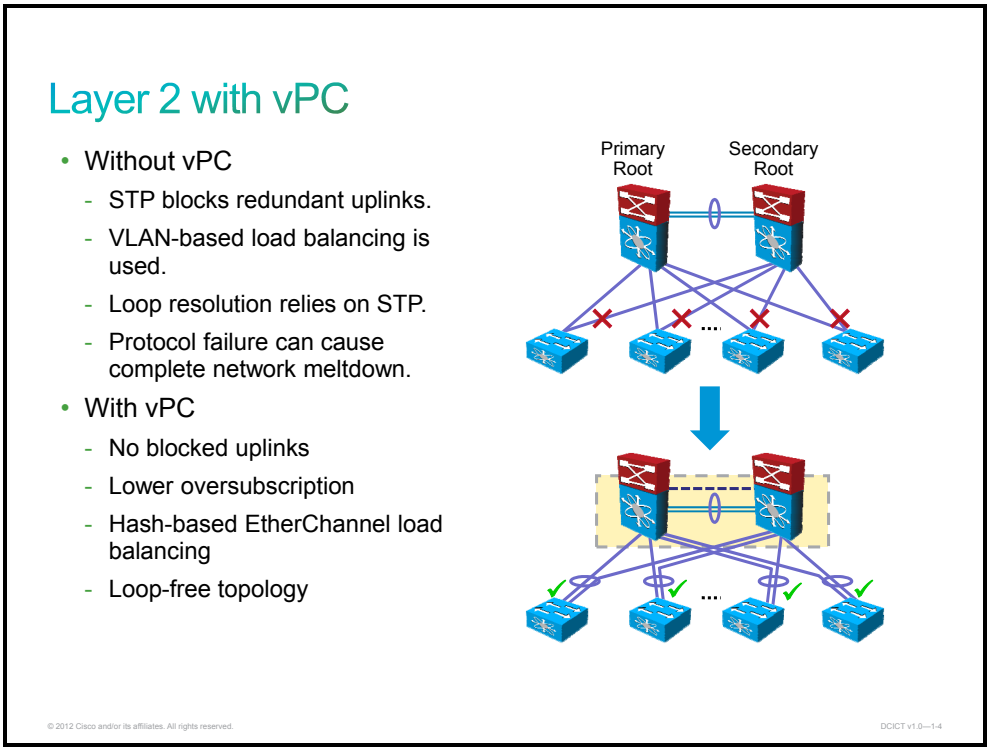
Upon completing this lesson, you will be able to describe vPCs and Cisco Fabric Path and how to verify their operation on Cisco Nexus 7000 and 5000 Series Switches. You will be able to meet these objectives:

- Describe vPCs
- Verify vPCs on the Cisco Nexus 7000 and 5000 Series Switches
- Describe Cisco FabricPath on the Cisco Nexus 7000 Series Switches and Nexus 5500 Platform switches
- Verify Cisco FabricPath on the Cisco Nexus 7000 Series Switches and 5500 Platform switches

https://t.me/learningnets

Virtual Port Channels

The Cisco vPC technology allows dual-homing of a downstream device to two upstream devices, yet makes the upstream devices look like one switch from a Spanning Tree Protocol (STP) perspective. This topic describes vPCs on the Cisco Nexus 7000 and 5000 Series Switches.



Virtualization technologies such as VMware ESX Server and clustering solutions such as Microsoft Cluster Service currently require Layer 2 Ethernet connectivity to function properly. Use of virtualization technologies in data centers, as well as across data center locations, leads organizations to shift from a highly-scalable Layer 3 network model to a highly-scalable Layer 2 model. This shift is causing changes in the technologies that are used to manage large Layer 2 network environments. These changes include migration away from STP as a primary loop management technology toward new technologies such as vPC and Cisco FabricPath.

In early Layer 2 Ethernet network environments, it was necessary to develop protocol and control mechanisms to limit the disastrous effects of a topology loop in the network. STP was the primary solution to this problem by providing a loop detection and loop management capability for Layer 2 Ethernet networks. This protocol has gone through a number of enhancements and extensions. While STP scales to very large network environments, it still has one suboptimal principle. This principle is that to break loops in a network, only one active path is allowed from one device to another. This principle is true regardless of how many actual connections might exist in the network.

An early enhancement to Layer 2 Ethernet networks was port-channel technology. This enhancement meant that multiple links between two participating devices could use all the links between the devices to forward traffic. Traffic is forwarded by using a load-balancing algorithm that equally balances traffic across the available interswitch links (ISLs). At the same time, the algorithm manages the loop problem by bundling the links as one logical link. This logical construct keeps the remote device from forwarding broadcast and unicast frames back to the logical link, breaking the loop that actually exists in the network. Port-channel technology

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

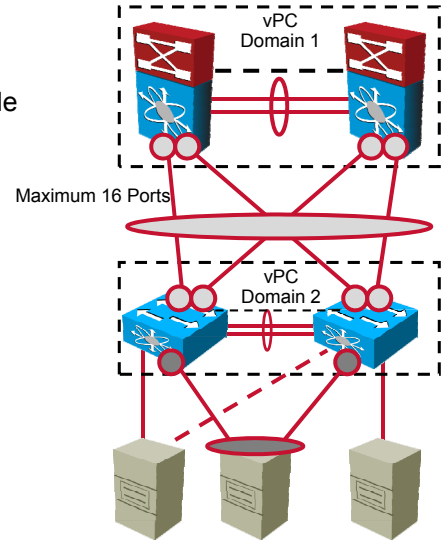
has another primary benefit. It can potentially manage a link loss in the bundle in less than a second, with very little loss of traffic and no effect on the active spanning-tree topology.

The biggest limitation in classic port-channel communication is that the port channel operates only between two devices. In large networks, the support of multiple devices together is often a design requirement to provide some form of hardware failure alternate path. This alternate path is often connected in a way that would cause a loop, limiting the benefits that are gained with port channel technology to a single path. To address this limitation, the Cisco NX-OS Software platform provides a technology called vPC. A pair of switches acting as a vPC peer endpoint looks like a single logical entity to port-channel-attached devices. However, the two devices that act as the logical port channel endpoint are still two separate devices. This environment combines the benefits of hardware redundancy with the benefits of port channel loop management. The other main benefit of migration to an all-port-channel-based loop management mechanism is that link recovery is potentially much faster. STP can recover from a link failure in approximately 6 seconds, while an all-port-channel-based solution has the potential for failure recovery in less than a second.

https://t.me/learningnets

Dual-Sided vPC

- vPC is supported on both the Cisco Nexus 5000 and Cisco Nexus 7000 Series Switches.
- vPC can be deployed in multiple layers of the data center simultaneously.
  - Server to access
  - Access to aggregation
- Dual-sided vPC enables a unique 16-way port channel.
  - Can be scaled to 32-way port channels with F-series modules



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

vPCs are supported on both the Cisco Nexus 7000 and Cisco Nexus 5000 Series Switches. The benefits that are provided by the vPC technology apply to any Layer 2 switched domain. Therefore, a vPC is commonly deployed in both the aggregation and access layers of the data center.

A vPC can be used to create a loop-free logical topology between the access and aggregation layer switches, which increases the bisectional bandwidth and improves network stability and convergence. It can also be used between servers and the access layer switches to enable server dual-homing with dual-active connections.

When the switches in the access and aggregation layers both support the vPC, a unique 16-way port channel can be created between the two layers. This scenario is commonly referred to as a dual-sided vPC. This design provides up to 160 Gb/s of bandwidth from a pair of access switches to the aggregation layer.

Note	If Cisco Nexus 7000 Series switches with F1-series modules are used on both sides of a dual-sided vPC, a 32-way port channel can be created to support up to 320 Gb/s of bandwidth between the access and aggregation layers.

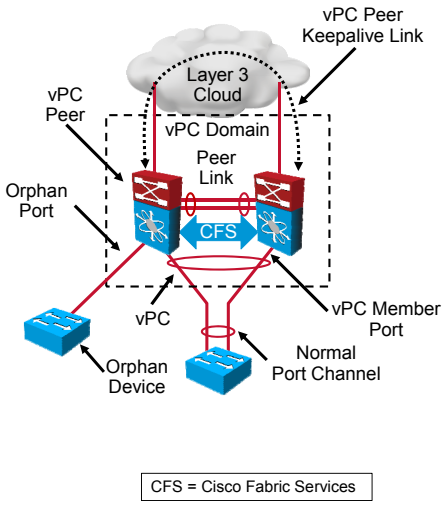
 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

vPC Architecture Components

The vPC architecture consists of the following components:

- vPC peers: A pair of vPC-enabled switches
- vPC peer link: Link that carries control traffic between vPC peer devices
- Cisco Fabric Services: A protocol, that is used for state synchronization and configuration validation between vPC peer devices



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

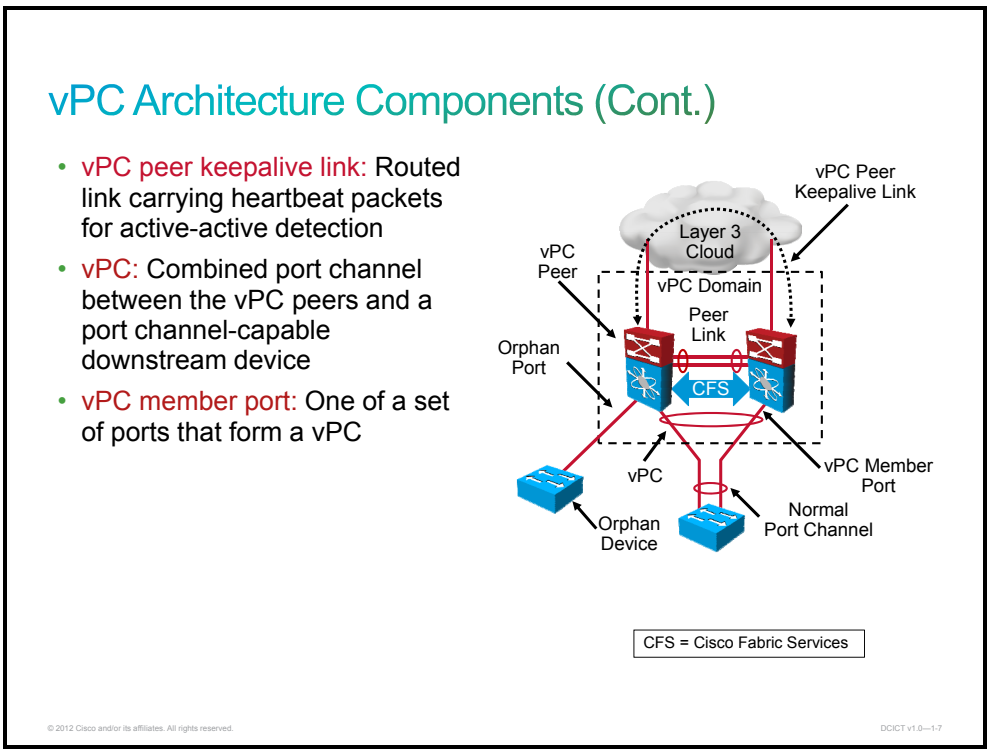
A pair of Cisco Nexus switches that uses the vPC appears to other network devices as a single logical Layer 2 switch. However, the two switches remain two separately managed switches with an independent management and control plane. The vPC architecture includes modifications to the data plane of the switches to ensure optimal packet forwarding. It also includes control plane components to exchange state information between the switches and allow the two switches to appear as a single logical Layer 2 switch to downstream devices.

The vPC architecture consists of the following components:

- **vPC peers:** The core of the vPC architecture is a pair of Cisco Nexus switches. This pair of switches acts as a single logical switch, which allows other devices to connect to the two chassis using a Multichassis EtherChannel (MEC).
- **vPC peer link:** The vPC peer link is the most important connectivity element in the vPC system. This link creates the illusion of a single control plane. It does so by forwarding bridge protocol data units (BPDUs) and Link Aggregation Control Protocol (LACP) packets to the primary vPC switch from the secondary vPC switch. The peer link is also used to synchronize MAC address tables between the vPC peers and to synchronize Internet Group Management Protocol (IGMP) entries for IGMP snooping. The peer link provides the necessary transport for multicast traffic and for the traffic of orphaned ports. If a vPC device is also a Layer 3 switch, the peer link also carries Hot Standby Router Protocol (HSRP) packets.
- **Cisco Fabric Services:** The Cisco Fabric Services protocol is a reliable messaging protocol that is designed to support rapid stateful configuration message passing and synchronization. The vPC peers use the Cisco Fabric Services protocol to synchronize data plane information and implement necessary configuration checks. vPC peers must synchronize the Layer 2 Forwarding table between the vPC peers. This way, if one vPC peer learns a new MAC address, that MAC address is also programmed on the Layer 2 Forwarding (L2F) table of the other peer device. The Cisco Fabric Services protocol travels on the peer link and does not require any configuration by the user. To help ensure that the peer link communication for the Cisco Fabric Services over Ethernet (FSoE) protocol is always available, spanning tree has been modified to keep the peer-link ports always

https://t.me/learningnets

forwarding. The Cisco FSoE protocol is also used to perform compatibility checks to validate the compatibility of vPC member ports to form the channel, to synchronize the IGMP snooping status, to monitor the status of the vPC member ports, and to synchronize the Address Resolution Protocol (ARP) table.

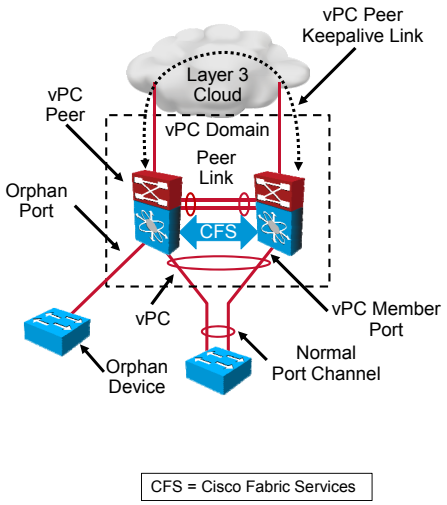


- **vPC peer keepalive link:** The peer keepalive link is a logical link that often runs over an out-of-band (OOB) network. It provides a Layer 3 communications path that is used as a secondary test to determine whether the remote peer is operating properly. No data or synchronization traffic is sent over the vPC peer keepalive link, just IP packets that indicate that the originating switch is operating and running vPC. The peer keepalive status is used to determine the status of the vPC peer when the vPC peer link goes down. In this scenario, it helps the vPC switch determine whether the peer link itself has failed, or if the vPC peer has failed entirely.
- **vPC:** A vPC is a MEC, a Layer 2 port channel that spans the two vPC peer switches. The downstream device that is connected on the vPC sees the vPC peer switches as a single logical switch. The downstream device does not need to support the vPC itself. It connects to the vPC peer switches using a regular port channel, which can either be statically configured or negotiated through LACP.
- **vPC member port:** This port is a port that is on one of the vPC peers that is a member of one of the vPCs that is configured on the vPC peers.

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

vPC Architecture Components (Cont.)

- vPC Domain: A pair of vPC peers and associated vPC components
- Orphan Device: A device that is connected to a vPC peer using a non-vPC link
- Orphan Port: Port on a vPC peer that connects to an orphan device
  - The term “orphan port” is also used for a vPC member port that connects to a device that has lost connectivity to the other vPC peer.



CFS = Cisco Fabric Services

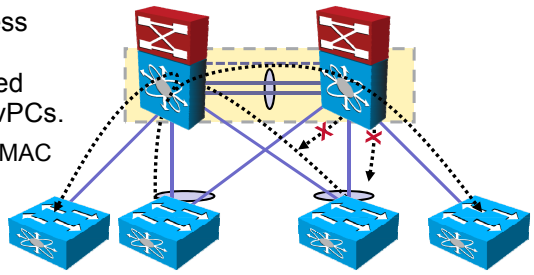
© 2012 Cisco and/or its affiliates. All rights reserved.

- **vPC domain:** The vPC domain includes both vPC peer devices, a vPC peer keepalive link, a vPC peer link, and all port channels in the vPC domain that are connected to the downstream devices. A numerical vPC domain ID identifies the vPC. You can have only one vPC domain ID on each device.
- **Orphan device:** The term “orphan device” refers to any device that is connected to a vPC domain using regular links instead of connecting through a vPC.
- **Orphan port:** The term “orphan port” refers to a switch port that is connected to an orphan device. The term is also used for vPC ports whose members are all connected to a single vPC peer. This situation can occur if a device that is connected to a vPC loses all its connections to one of the vPC peers.

https://t.me/learningnets

vPC Data Plane Operation

- The vPC peer link carries the following traffic only:
  - vPC control traffic
  - Flooded traffic (broadcast, multicast, unknown unicast)
  - Traffic for orphan ports
- Regular switch MAC address learning is replaced with Cisco Fabric Services-based MAC address learning for vPCs.
  - Non-vPC ports use regular MAC address learning.
- Frames that enter a vPC peer switch from the peer link cannot exit the switch on a vPC member port.



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

vPC technology has been designed to limit the use of the peer link specifically to switch management traffic and the occasional traffic flow from a failed network port. The peer link does not carry regular traffic for vPCs. This feature allows the vPC solution to scale, because the bandwidth requirements for the peer link are not directly related to the total bandwidth requirements of the vPCs.

The vPC peer link is designed to carry several types of traffic. To start with, it carries vPC control traffic, such as Cisco FSoE, BPDUs, and LACP messages. In addition, it carries traffic that needs to be flooded, such as broadcast, multicast, and unknown unicast traffic. It also carries traffic for orphan ports.

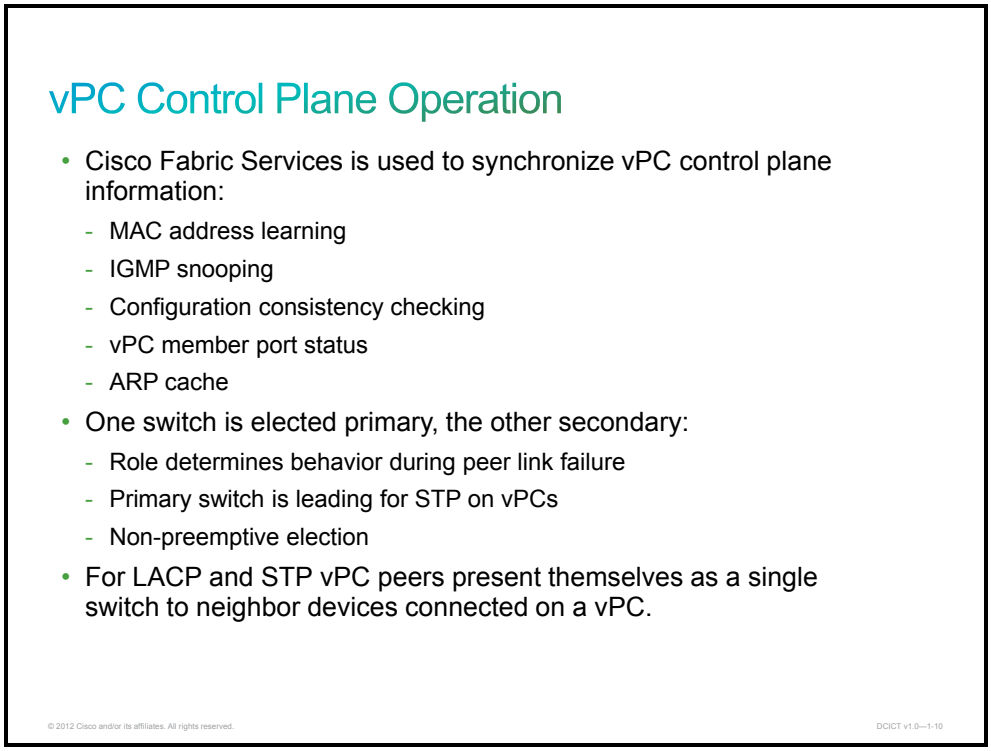
The term orphan port is used for two types of ports. An orphan port refers to any Layer 2 port on a vPC peer switch that does not participate in vPC. These ports use normal switch forwarding rules. Traffic from these ports can use the vPC peer link as a transit link to reach orphan devices that are connected to the other vPC peer switch. An orphan port can also refer to a port that is a member of a vPC, but for which the peer switch has lost all the associated vPC member ports. When a vPC peer switch loses all member ports for a specific vPC, it will forward traffic that is destined for that vPC to the vPC peer link. In this special case, the vPC peer switch will be allowed to forward the traffic that is received on the peer link to one of the remaining active vPC member ports.

To implement the specific vPC forwarding behavior, it is necessary to synchronize the L2F tables between the vPC peer switches through Cisco Fabric Services instead of depending on the regular MAC address learning. Cisco Fabric Services-based MAC address learning only applies to vPC ports and is not used for ports that are not in a vPC.

One of the most important forwarding rules for vPC is that a frame that enters the vPC peer switch from the peer link cannot exit the switch from a vPC member port. This principle prevents frames that are received on a vPC from being flooded back onto the same vPC by the other peer switch. The exception to this rule is traffic that is destined for an orphaned vPC member port.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



Cisco Fabric Services is used as the primary control plane protocol for vPC. It performs several functions:

- vPC peers must synchronize the Layer 2 MAC address table between the vPC peers. For example, one vPC peer learns a new MAC address on a vPC. That MAC address is also programmed on the L2F table of the other peer device for that same vPC. This MAC address learning mechanism replaces the regular switch MAC address learning mechanism and prevents traffic from being forwarded across the vPC peer link unnecessarily.
- The synchronization of IGMP snooping information is performed by Cisco Fabric Services. L2F of multicast traffic with vPC is based on modified IGMP snooping behavior that synchronizes the IGMP entries between the vPC peers. In a vPC implementation, IGMP traffic entering a vPC peer switch through a vPC triggers hardware programming for the multicast entry on both vPC member devices.
- Cisco Fabric Services is also used to communicate essential configuration information to ensure configuration consistency between the peer switches. Similar to regular port channels, vPCs are subject to consistency checks and compatibility checks. During a compatibility check, one vPC peer conveys configuration information to the other vPC peer to verify that vPC member ports can actually form a port channel. In addition to compatibility checks for the individual vPCs, Cisco Fabric Services is also used to perform consistency checks for a set of switchwide parameters that need to be configured consistently on both peer switches.
- Cisco Fabric Services is used to track vPC status on the peer. When all vPC member ports on one of the vPC peer switches go down, Cisco Fabric Services is used to notify the vPC peer switch that its ports have become orphan ports. All traffic that is received on the peer link for that vPC should now be forwarded via the local vPC.
- Starting from Cisco NX-OS Software Release 5.0(2) for the Cisco Nexus 5000 Series Switches and Cisco NX-OS Software version 4.2(6) for the Cisco Nexus 7000 Series Switches, Layer 3 vPC peers synchronize their respective ARP tables. This feature is transparently enabled and helps ensure faster convergence time upon reload of a vPC

https://t.me/learningnets

switch. When two switches are reconnected after a failure, they use Cisco Fabric Services to perform bulk synchronization of the ARP table.

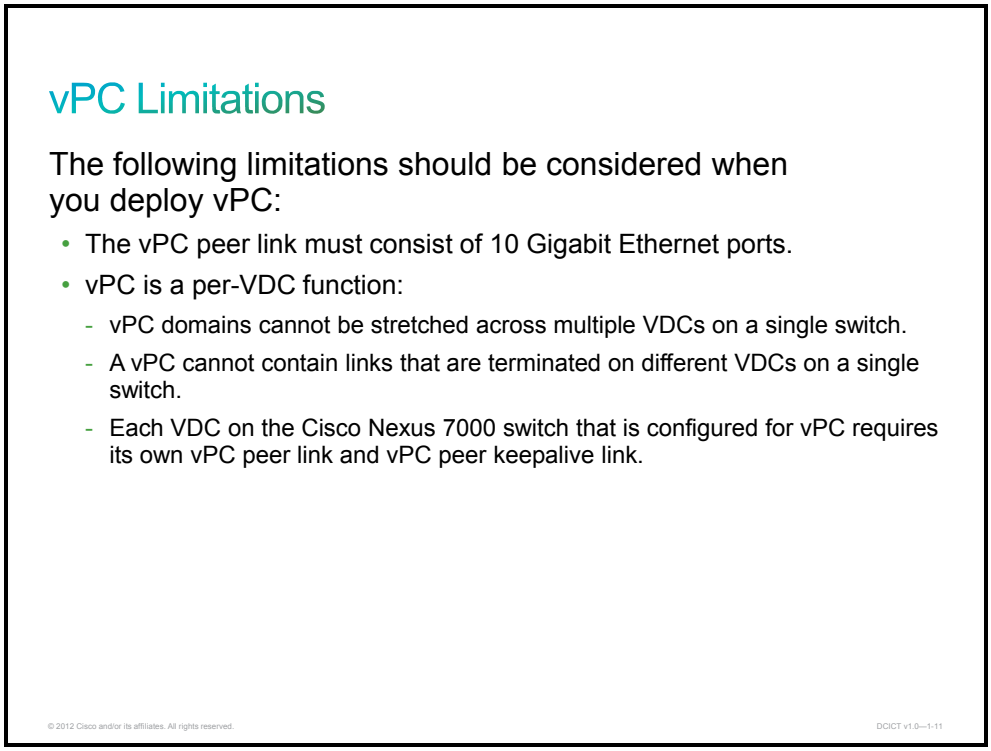
Between the pair of vPC peer switches, an election is held to determine a primary and secondary vPC device. This election is nonpreemptive. The vPC primary or secondary role is primarily a control plane role. This role determines which of the two switches will primarily be responsible for the generation and processing of spanning-tree BPDUs for the vPCs.

Note Starting from Cisco NX-OS Software Release 5.0(2) for the Cisco Nexus 5000 Series Switches and Cisco NX-OS Software Release 4.2(6) for the Cisco Nexus 7000 Series Switches, the vPC peer-switch option can be implemented. This option allows both the primary and secondary to generate BPDUs for vPCs independently. The two switches will use the same spanning-tree bridge ID to ensure that devices that are connected on a vPC still see the vPC peers as a single logical switch.

Both switches actively participate in traffic forwarding for the vPCs. However, the primary and secondary roles are also important in certain failure scenarios, most notably in a peer link failure. When the vPC peer link fails, the vPC peer switches attempt to determine through the peer keepalive mechanism if the peer switch is still operational. If it is operational, the operational secondary switch suspends all vPC member ports. The secondary also shuts down all switch virtual interfaces (SVIs) that are associated with any VLANs that are configured as allowed VLANs for the vPC peer link.

For LACP and STP, the two vPC peer switches present themselves as a single logical switch to devices that are connected on a vPC. For LACP, this appearance is accomplished by generating the LACP system ID from a reserved pool of MAC addresses, which are combined with the vPC domain ID. For STP, the behavior depends on the use of the peer-switch option. If the peer-switch option is not used, the vPC primary is responsible for generating and processing BPDUs and uses its own bridge ID for the BPDUs. The secondary switch relays BPDU messages, but does not generate BPDUs itself for the vPCs. When the peer-switch option is used, both the primary and secondary switches send and process BPDUs. However, they will use the same bridge ID to present themselves as a single switch to devices that are connected on a vPC.

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



Consider the following limitations when you deploy a vPC:

- Only 10 Gigabit Ethernet ports can be used for the vPC peer link. It is recommended that you use at least two 10 Gigabit Ethernet ports in dedicated mode on two different I/O modules.
- vPC is a per-VDC function on the Cisco Nexus 7000 Series Switches. vPCs can be configured in multiple virtual device contexts (VDCs), but the configuration is entirely independent. A separate vPC peer link and vPC peer keepalive link is required for each of the VDCs. vPC domains cannot be stretched across multiple VDCs on the same switch, and all ports for a given vPC must be in the same VDC.

https://t.me/learningnets

vPC Limitations (Cont.)

The following limitations should be considered when you deploy vPC:

- A vPC domain cannot consist of more than two peer switches or VDCs.
- You cannot configure more than one vPC domain per switch or VDC.
- A vPC is a Layer 2 port channel.
  - Dynamic routing to vPC peers across a vPC or across the vPC peer link is not supported.
  - Static routing across a vPC to an FHRP addresses is supported.
  - Dynamic routing across a vPC between two Layer 3 switches that are not participating in vPC is supported.

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

A vPC domain by definition consists of a pair of switches that is identified by a shared vPC domain ID. It is not possible to add more than two switches or VDCs to a vPC domain.

Only one vPC domain ID can be configured on a single switch or VDC. It is not possible for a switch or VDC to participate in more than one vPC domain.

A vPC is a Layer 2 port channel. The vPC technology does not support the configuration of Layer 3 port channels. Dynamic routing from the vPC peers to routers connected on a vPC is not supported. It is recommended that routing adjacencies be established on separate routed links.

Static routing to First Hop Routing Protocol (FHRP) addresses is supported. The FHRP enhancements for a vPC enable routing to a virtual FHRP address across a vPC.

A vPC can be used as a Layer 2 link to establish a routing adjacency between two external routers. The routing restrictions for the vPC only apply to routing adjacencies between the vPC peer switches and routers that are connected on a vPC.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying vPCs

This topic explains how to verify vPCs on the Cisco Nexus 7000 and 5000 Series Switches.

Verifying vPC

To verify vPC operation, use the **show vpc brief** command.

```
Nexus-7010-PROD-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 103
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 2
Peer Gateway                      : Enabled
Peer gateway excluded VLANs       : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po7    up     1,113-114

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- ------                     ------------
71   Po71   up     success     success                    1,113-114

72   Po72   up     success     success                    1,113-114
```

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

Several commands can be used to verify the operation of vPC. The primary command to be used in initial verification is the **show vpc brief** command. This command displays the following operational information:

- vPC domain ID
- Peer-link status
- Keepalive message status
- Status of the successful configuration consistency
- Status of whether the peer link is formed
- Status of the individual vPCs that are configured on the switch (including the result of the consistency checks)

https://t.me/learningnets

Verifying vPC (Cont.)

To check for potential vPC configuration consistency problems, use the **show vpc consistency-parameters** command.

```
Nexus-7010-PROD-A# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
Interface-vlan admin up     2     1                      1
Interface-vlan routing      2     1                      1
capability
Allowed VLANs               -     1,113-114              1,113-114
Local suspended VLANs       -     -                      -
Nexus-7010-PROD-A#
```

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

If the **show vpc brief** command displays failed consistency checks, you can use the **show vpc consistency-parameters** command to find the specific parameters that caused the consistency check to fail. The **global** option on this command allows you to verify the consistency of the global parameters between the two peer switches. The **vpc** or **interface** option can be used to verify consistency between the port channel configurations for vPC member ports.

After you enable the vPC feature and configure the peer link on both vPC peer devices, Cisco Fabric Services messages provide a copy of the configuration on the local vPC peer device configuration to the remote vPC peer device. The system then determines whether any of the crucial configuration parameters differ on the two devices.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying vPC (Cont.)

To check for potential vPC configuration consistency problems, use the **show vpc consistency-parameters** command.

```
Nexus-7010-PROD-A# show vpc consistency-parameters vpc 71

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Port Type               1     Default                Default
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-67, 8047, 0-23-4-ee-be-67, 8047,
                                   0, 0), (8000,          0, 0), (8000,
                                  54-7f-ee-5c-6e-fc, 46, 54-7f-ee-5c-6e-fc, 46,
                                   0, 0)]                 0, 0)]
mode                        1     active                 active
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
vPC card type               1     Orion                  Orion
Allowed VLANs               -     1,113-114              1,113-114
Local suspended VLANs       -     -                      -
```

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

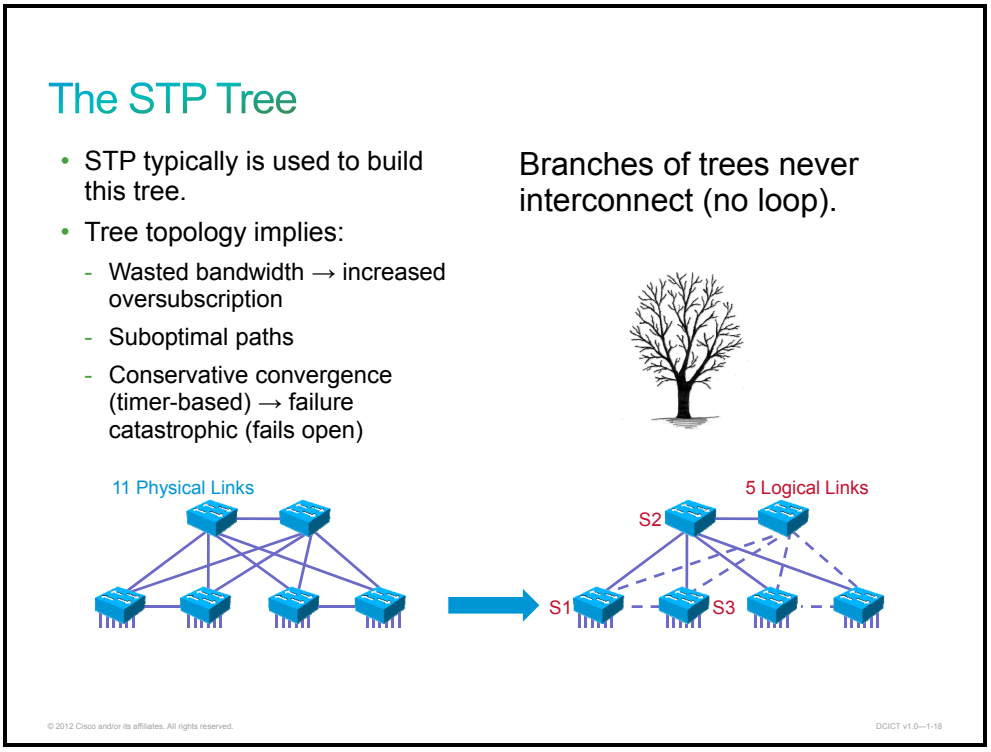
The configuration parameters in this section must be configured identically on both devices of the vPC peer link or the vPC moves into suspend mode. The devices automatically check for compatibility for some of these parameters on the vPC interfaces. The per-interface parameters must be consistent per interface, and the global parameters must be consistent globally. The following parameters are checked:

- Port-channel mode: on, off, or active
- Link speed per channel
- Duplex mode per channel
- Trunk mode per channel (including native VLAN, VLANs that are allowed on trunk, and the tagging of native VLAN traffic)
- STP mode
- STP region configuration for Multiple Spanning Tree (MST)
- Enabled or disabled state per VLAN
- STP global settings, including Bridge Assurance setting, port type setting, and loop guard settings
- STP interface settings, including port type setting, loop guard, and root guard
- Maximum Transmission Unit (MTU)

https://t.me/learningnets

Cisco FabricPath

Cisco FabricPath is a feature that can be enabled on the Cisco Nexus switches to provide routing functionality at a Layer 2 level in the data center. This topic describes the Cisco FabricPath feature on the Cisco Nexus 7000 Series Switches and Nexus 5500 Platform switches.



To support the Layer 2 domain and high availability, switches are normally interconnected. The STP then runs on the switches to create a tree-like structure that is loop-free. To provide a loop-free topology, spanning tree builds the tree and then blocks certain ports to ensure that traffic cannot loop around the network endlessly.

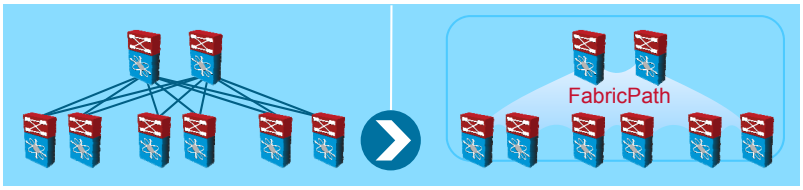
This tree topology implies that certain links are unused, traffic does not necessarily take the optimal path, and when a failure does occur, the convergence time is based around timers.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco FabricPath: An Ethernet Fabric

- Connect a group of switches using an arbitrary topology.
- With a simple CLI, aggregate them into a fabric.
- An open protocol based on L3 technology provides fabric-wide intelligence and ties the elements together.



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

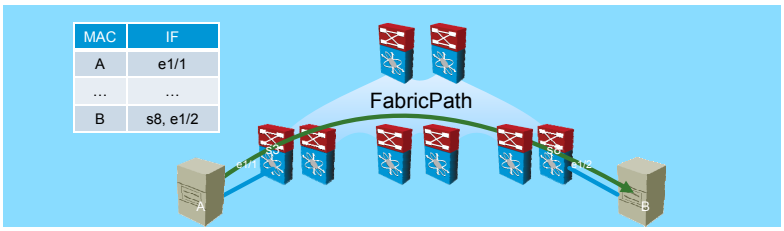
Cisco FabricPath is an innovative Cisco NX-OS feature that is designed to bring the stability and performance of routing to Layer 2. It brings the benefits of Layer 3 routing to Layer 2 switched networks to build a highly resilient and scalable Layer 2 fabric.

Cisco FabricPath switching allows multipath networking at the Layer 2 level. The Cisco FabricPath network still delivers packets on a best-effort basis (which is similar to the Classical Ethernet network), but the Cisco FabricPath network can use multiple paths for Layer 2 traffic. In a Cisco FabricPath network, you do not need to run the STP. Instead, you can use Cisco FabricPath across data centers, some of which have only Layer 2 connectivity, with no need for Layer 3 connectivity and IP configurations.

https://t.me/learningnets

Cisco FabricPath: An Ethernet Fabric (Cont.)

- Shortest path is any-to-any.
- Single address lookup at the ingress edge identifies the exit port across the fabric.
- Traffic is switched using the shortest path that is available.
- Reliable Layer 2 and Layer 3 connectivity is provided using any-to-any (Layer 2 as if it was within the same switch, no STP inside).



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

Externally, a fabric looks like a single switch, yet internally there is a protocol that adds fabric-side intelligence. This intelligence ties the elements of the Cisco FabricPath infrastructure together. The protocol provides the following:

- Any-to-any optimal, low latency connectivity
- High bandwidth and high resiliency
- Open management and troubleshooting

Cisco FabricPath provides scalability at a Layer 2 level.

Frames are forwarded along the shortest path to their destination, reducing the latency of the exchanges between end stations when compared to a spanning tree-based solution.

MAC addresses are learned selectively at the edge, allowing the network to scale beyond the limits of the MAC address table of individual switches.

The any-to-any shortest path provides the following:

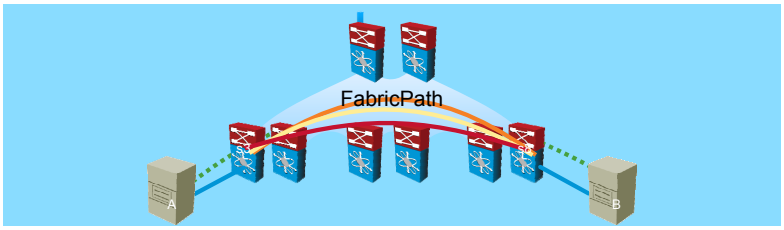
- A single address lookup at the ingress edge that identifies the exit port across the fabric
- Traffic is switched using the shortest path available
- Reliable Layer 2 and Layer 3 connectivity between any devices

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco FabricPath: An Ethernet Fabric (Cont.)

- ECMP
- Multipathing (up to 16 links active between any 2 devices)
- Traffic is redistributed across remaining links in case of failure, providing fast convergence



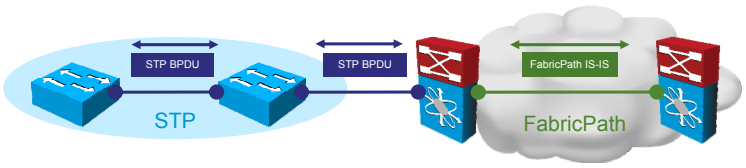
© 2012 Cisco and/or its affiliates. All rights reserved.

Because Equal-Cost Multipath (ECMP) can be used at the data plane, the network can use all the links that are available between any two devices. The first-generation hardware supporting Cisco FabricPath can perform 16-way ECMP, which, when combined with 16-port 10-Gb/s port channels, represents a potential bandwidth of 2.56 Tb/s between switches.

https://t.me/learningnets

Cisco FabricPath Control Plane

- Cisco FabricPath IS-IS replaces STP as the control-plane protocol in Cisco FabricPath network
- Introduces link-state protocol with support for ECMP for Layer 2 forwarding
- Exchanges reachability of switch IDs and builds forwarding trees
- Improves failure detection, network reconvergence, and high availability
- Minimal IS-IS knowledge required—no user configuration by default



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

With Cisco FabricPath, you use the Layer 2 Intermediate System-to-Intermediate System (IS-IS) protocol for a single control plane that functions for unicast, broadcast, and multicast packets. There is no need to run STP. It is a purely Layer 2 domain. This Cisco FabricPath Layer 2 IS-IS is a separate process than Layer 3 IS-IS.

IS-IS provides following benefits:

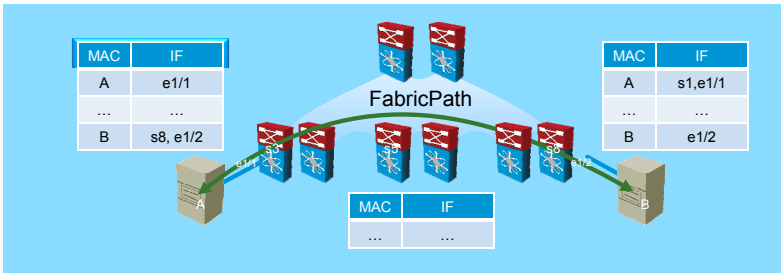
- **Has no IP dependency:** No need for IP reachability to form adjacency between devices
- **Easily extensible:** Using custom type, length, value (TLV) settings, IS-IS devices can exchange information about virtually anything
- **Provides Shortest Path First (SPF) routing:** Excellent topology building and reconvergence characteristics

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Conversational Learning

- Per-port MAC address table only needs to learn the peers that are reached across the fabric.
- A virtually unlimited number of hosts can be attached to the fabric.



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

With Cisco NX-OS Software Release 5.1(1) and the F-Series module, you can use conversational MAC address learning. The MAC address learning is configurable, and can either be conversational or traditional on a VLAN-by-VLAN basis.

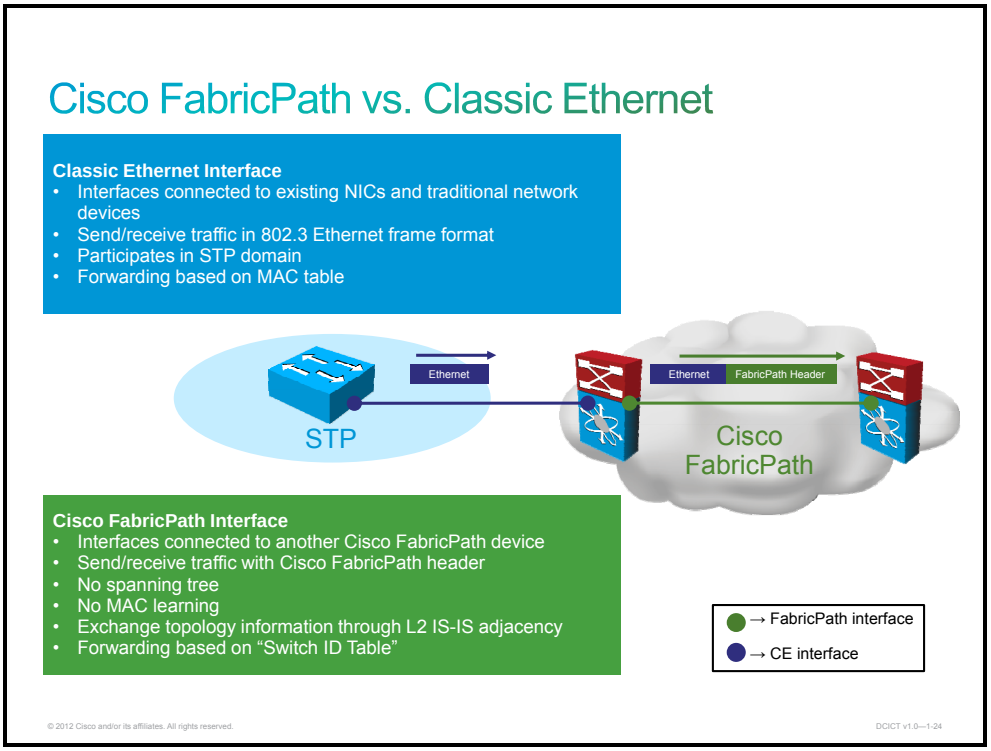
Conversational MAC address learning means that each interface learns only those MAC addresses for interested hosts, rather than all MAC addresses in the domain. Each interface learns only those MAC addresses that are actively speaking with the interface. In effect, conversational MAC address learning consists of a three-way handshake.

Conversational MAC address learning permits the scaling of the network beyond the limits of individual switch MAC address tables.

All Cisco FabricPath VLANs use conversational MAC address learning.

Classical Ethernet VLANs use traditional MAC address learning by default. However, this setting is configurable.

https://t.me/learningnets



To interact with the Classic Ethernet network, you set VLANs to either Classic Ethernet or Cisco FabricPath mode. The Classic Ethernet VLANs carry traffic from the Classic Ethernet hosts to the Cisco FabricPath interfaces. The Cisco Fabric Path VLANs carry traffic throughout the Cisco FabricPath topology. Only the active Cisco FabricPath VLANs that are configured on a switch are advertised as part of the topology in the Layer 2 IS-IS messages.

The following interface modes carry traffic for the following types of VLANs:

- Interfaces on the F Series modules that are configured as Cisco FabricPath interfaces can carry traffic only for Cisco FabricPath VLANs
- Interfaces on the F Series modules that are not configured as Cisco FabricPath interfaces carry traffic for the following:
  - Cisco FabricPath VLANs
  - Classical Ethernet VLANS
- Interfaces on the M Series modules carry traffic only for classic Ethernet VLANs

To have a loop-free topology for the classic Ethernet and Cisco FabricPath hybrid network, the Cisco FabricPath network automatically presents as a single bridge to all connected classic Ethernet devices.

Other than configuring the STP priority on the Cisco FabricPath Layer 2 gateway switches, you do not need to configure anything for the STP to work seamlessly with the Cisco FabricPath network. Only connected classic Ethernet devices form a single STP domain. Those classic Ethernet devices that are not interconnected form separate STP domains.

All classic Ethernet interfaces should be designated ports, which occur automatically, or they will be pruned from the active STP topology. If the system does prune any port, the system returns a syslog message. The system clears the port again only when that port is no longer receiving superior BPDUs.

The Cisco FabricPath Layer 2 gateway switch also propagates the topology change notifications (TCNs) on all its classic Ethernet interfaces.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

The Cisco FabricPath Layer 2 gateway switches terminate STP. The set of Cisco FabricPath Layer 2 gateway switches that are connected by STP forms the STP domain. Because there can be many Cisco FabricPath Layer 2 gateway switches that are attached to a single Cisco FabricPath network, there may also be many separate STP domains.

https://t.me/learningnets

Verifying Cisco FabricPath

This topic explains how to verify correct operation of Cisco FabricPath on the Cisco Nexus 7000 and 5500 Platform switches.

Verifying MAC Address Learning

To verify that local and remote MAC addresses are learned for the Cisco FabricPath VLANs, use the **show mac address table** command.

```
Nexus-7010-PROD-B# show mac address-table dynamic vlan 10
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0000.0000.0001    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0002    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0003    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0004    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0005    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0006    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0007    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0008    dynamic   0          F    F  Eth1/15
* 10       0000.0000.0009    dynamic   0          F    F  Eth1/15
* 10       0000.0000.000a    dynamic   0          F    F  Eth1/15
  10       0000.0000.000b    dynamic   0          F    F  200.0.30
  10       0000.0000.000c    dynamic   0          F    F  200.0.30
  10       0000.0000.000d    dynamic   0          F    F  200.0.30
  10       0000.0000.000e    dynamic   0          F    F  200.0.30
  10       0000.0000.000f    dynamic   0          F    F  200.0.30
```

Local MAC

Remote MAC

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

The **show mac address table** command can be used to verify that MAC addresses are learned on Cisco FabricPath edge devices. The command shows local addresses with a pointer to the interface that the address was learned on. For remote addresses, it provides a pointer to the remote switch from which this address was learned.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying Cisco FabricPath Routing

To examine the Cisco FabricPath routes between the switches in the fabric, use the **show fabricpath route** command.

```
Nexus-7010-PROD-A# show fabricpath route
FabricPath Unicast Route Table
'a/b/c' denotes ftag/switch-id/subswitch-id
'[x/y]' denotes [admin distance/metric]
ftag 0 is local ftag
subswitch-id 0 is default subswitch-id


FabricPath Unicast Route Table for Topology-Default

0/71/0, number of next-hops: 0
        via ----- , [60/0], 0 day/s 00:12:10, local
1/51/0, number of next-hops: 2
        via Eth2/5, [115/40], 0 day/s 00:01:03, isis_fabricpath-default
        via Eth2/6, [115/40], 0 day/s 00:01:03, isis_fabricpath-default
1/72/0, number of next-hops: 2
        via Eth2/11, [115/40], 0 day/s 00:01:13, isis_fabricpath-default
        via Eth2/12, [115/40], 0 day/s 00:01:13, isis_fabricpath-default
```

Topology ID / Switch ID / Subswitch ID

Multipathing

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

The **show fabricpath route** command can be used to view the Cisco FabricPath routing table that results from the Cisco FabricPath IS-IS SPF calculations. The Cisco FabricPath routing table shows the best paths to all the switches in the fabric. If multiple equal paths are available between two switches, all paths will be installed in the Cisco FabricPath routing table to provide ECMP.

https://t.me/learningnets

Summary

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

Summary

- vPCs allow downstream devices to be dual-homed to two upstream switches running vPCs so that the connectivity appears to be one downstream switch connected to one upstream switch.
- Verifying the creation of vPCs involves using **show** commands such as **show vpc.**
- Cisco FabricPath is Layer 2 routing. This technology provides multiple equal cost paths and negates the use of the Spanning Tree Protocol.
- Verifying the configuration of Cisco FabricPath involves using **show** commands such as **show fabricpath route**.

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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 6

Using OTV on Cisco Nexus 7000 Series Switches

Overview

Cisco Overlay Transport Virtualization (OTV) is a “MAC in IP” technique for supporting the extension of Layer 2 VLANs over any transport that can forward IP packets. This lesson describes OTV as a method of data center interconnect

Objectives

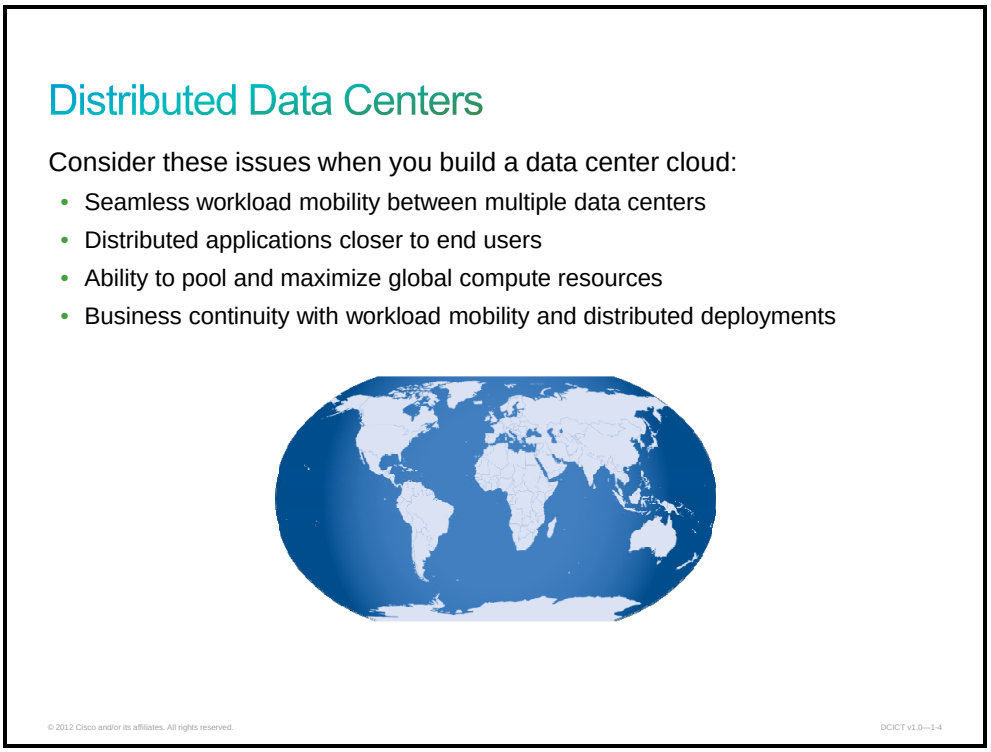
Upon completing this lesson, you will be able to describe OTV as a method of data center interconnect (DCI). You will be able to meet these objectives:

- Describe OTV on the Cisco Nexus 7000 Series Switches
- Verify OTV on the Cisco Nexus 7000 Series Switches

https://t.me/learningnets

OTV on the Cisco Nexus 7000 Series Switches

OTV is a data center interconnect (DCI) technology providing Layer 2 connectivity between multiple data centers over an IP infrastructure. This topic describes OTV on the Cisco Nexus 7000 Series Switches.

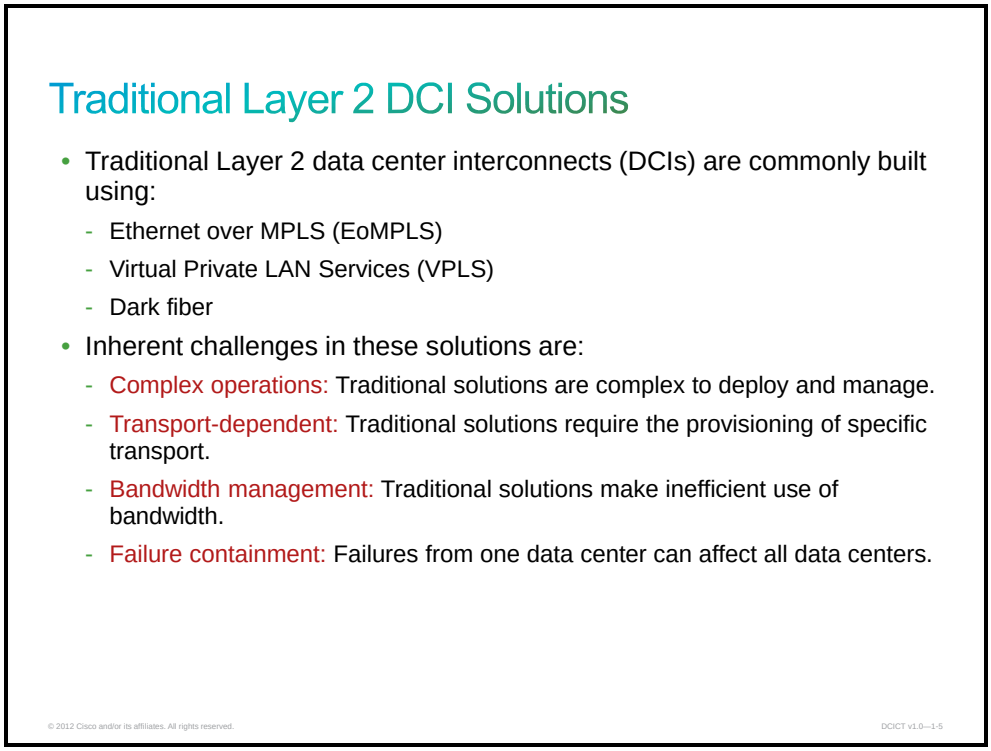


Businesses face the challenge of providing very high availability for applications while keeping operating expenses (OpEx) low. Applications must be available anytime and anywhere with optimal response times.

The deployment of geographically dispersed data centers allows IT architects to put in place effective disaster-avoidance and disaster-recovery mechanisms that increase the availability of the applications. Geographic dispersion also enables optimization of application response time through improved facility placement. Flexible mobility of workloads across data centers helps avoid demand hotspots and more fully utilizes available capacity.

To enable all the benefits of geographically dispersed data centers, the network must extend Layer 2 connectivity across the diverse locations. This connectivity must be provided without compromising the autonomy of data centers or the stability of the overall network.

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



Enabling key technologies such as server clustering and workload mobility across data centers requires Layer 2 connectivity between data centers.

Traditionally, several different technologies have been used to provide Layer 2 DCIs:

- **Ethernet over Multiprotocol Label Switching (EoMPLS):** EoMPLS can be used to provision point-to-point Layer 2 Ethernet connections between two sites using MPLS as the transport.
- **Virtual Private LAN Services (VPLS):** Similar to EoMPLS, VPLS uses an MPLS network as the underlying transport network. However, instead of point-to-point Ethernet pseudowires, VPLS delivers a virtual multiaccess Ethernet network.
- **Dark fiber:** In some cases, dark fiber may be available to build private optical connections between data centers. Dense wavelength-division multiplexing (DWDM) or coarse wavelength-division multiplexing (CWDM) can increase the number of Layer 2 connections that can be run through the fibers. These technologies increase the total bandwidth over the same number of fibers.

Although it is possible to build Layer 2 DCIs based on these technologies, they present a number of challenges:

- **Complex operations:** Traditional Layer 2 VPN technologies can provide extended Layer 2 connectivity across data centers. But these technologies usually involve a mix of complex protocols, distributed provisioning, and an operationally intensive hierarchical scaling model. The provisioning of many point-to-point connections or the complexity of the underlying transport technologies can add significantly to the operational cost of these types of DCIs. A simple overlay protocol with built-in capabilities and point-to-cloud provisioning is crucial to reducing the cost of providing this connectivity.
- **Transport dependent:** Traditional Layer 2 DCI technologies depend on the availability of a specific transport, such as dark fiber or an MPLS network. This dependency limits the options that are available to interconnect geographically dispersed data centers, because it requires the underlying transport to be available at all data center sites. A cost-effective

https://t.me/learningnets

solution for the interconnection of data centers must be transport agnostic. Such independence gives the network architect the flexibility to choose any transport between data centers that is based on business and operational preferences.

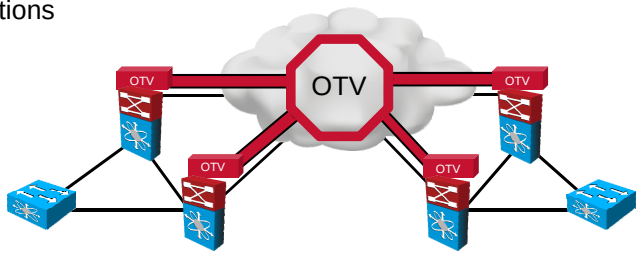
- **Bandwidth management:** Traditional Layer 2 DCI technologies often do not allow the concurrent use of redundant connections, because of the risk of Layer 2 loops. Balancing the load across all available paths while providing resilient connectivity between the data center and the transport network requires added intelligence. Traditional Ethernet switching and Layer 2 VPN do not meet that level of intelligence.
- **Failure containment:** The extension of Layer 2 domains across multiple data centers can cause problems. Traditional Layer 2 extension technologies between multiple data centers often extend the failure domain between the data centers that are causing issues. These failures propagate freely over the open Layer 2 flood domain. A solution that provides Layer 2 connectivity yet restricts the reach of the flood domain is needed. Such a solution would contain failures and thus preserve the resiliency that is achieved by using multiple data centers.

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

Overlay Transport Virtualization

OTV delivers a virtual Layer 2 transport over any Layer 3 infrastructure.

- Overlay: A solution that is independent of the infrastructure technology and services; flexible over various interconnect facilities
- Transport: Technology that transports services for Layer 2 Ethernet and IP traffic
- Virtualization: Technology that provides virtual stateless multiaccess connections



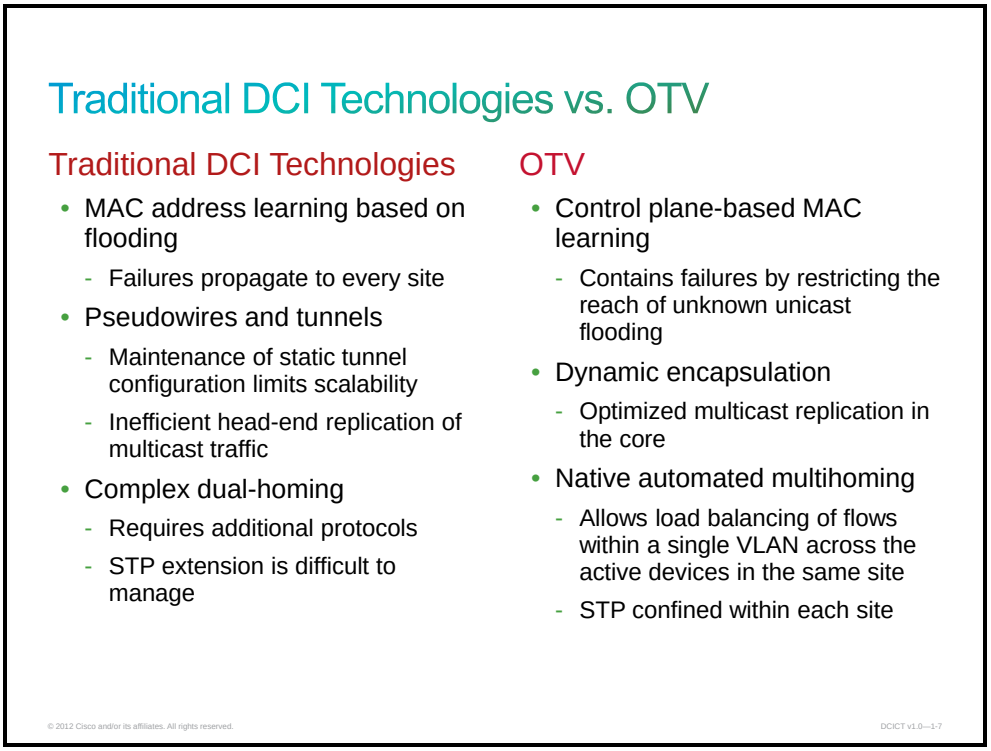
© 2012 Cisco and/or its affiliates. All rights reserved.

OTV overcomes the challenges that are inherent to traditional Layer 2 DCI technologies.

The name of the technology describes its key characteristics:

- **Overlay:** OTV provides an overlay VPN in addition to an IP network. It is independent of the underlying infrastructure technologies and services. It does not impose any restrictions on the underlying infrastructure as long as it is capable of transporting IP packets.
- **Transport:** OTV provides a Layer 2 transport across a Layer 3 network. It can leverage all the underlying capabilities of the underlying transport network, such as fast convergence, load balancing, and multicast replication.
- **Virtualization:** OTV provides a virtual multiaccess Layer 2 network that supports efficient transport of unicast, multicast, and broadcast traffic. Sites can be added to an OTV overlay without a need to provision additional point-to-point connections to the other sites. There are no virtual circuits, pseudowires or other point-to-point connections to maintain between the sites. Packets are routed independently from site to site without a need to establish stateful connections between the sites.

https://t.me/learningnets

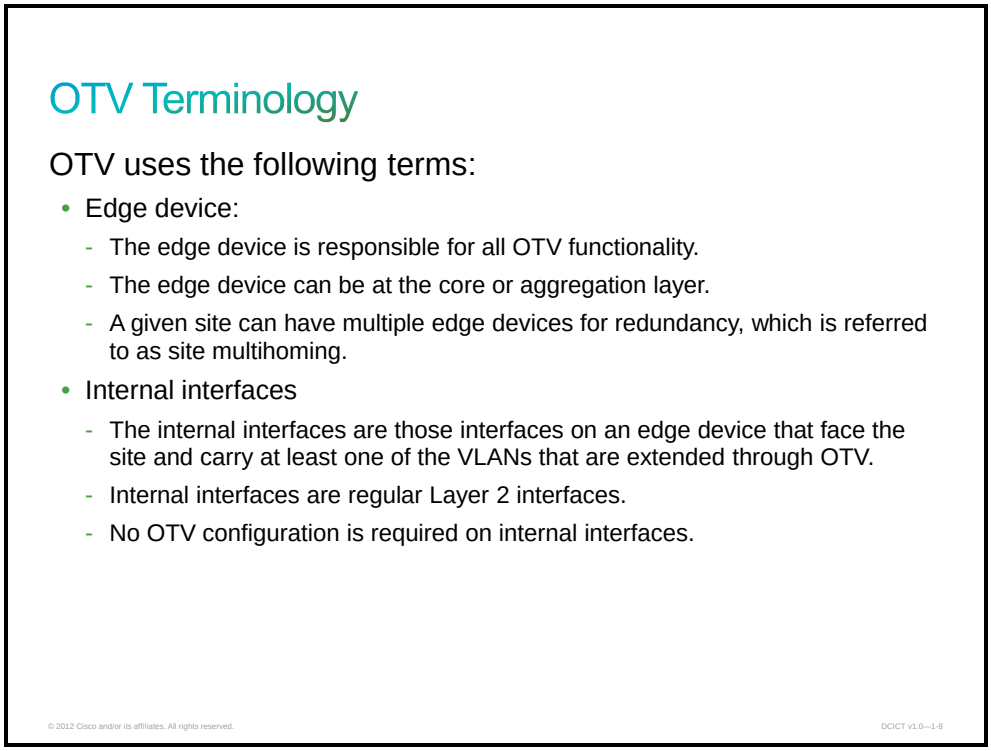


Traditional Layer 2 DCI technologies depend on flooding of unknown unicast, multicast, and broadcast frames to learn the location of the MAC addresses on the network. OTV replaces this mechanism with protocol-driven control plane learning. Flooding of unknown unicasts is no longer necessary. Eliminating unknown unicast flooding helps in containing failures by limiting the scope of the flooding domain.

Traditional Layer 2 DCI technologies use a mesh of point-to-point connections to connect multiple sites. The configuration and maintenance of these tunnels limits the scalability of Layer 2 DCI solutions. In addition, a point-to-point connection model leads to inefficient multicast forwarding. Multicasts need to be replicated at the headend to be sent to the different sites. OTV is a multipoint technology. It leverages the multipoint connectivity model of Layer 3 networks to easily add more sites to the OTV overlay without a need to provision point-to-point connections between the sites. If the Layer 3 network is multicast-enabled, OTV can leverage the inherent multicast capabilities of the Layer 3 network to replicate multicast in the core network, avoiding headend replication.

With traditional Layer 2 DCI technologies dual-homing, a site to the transport network is often complex. Redundant connections introduce the possibility of Layer 2 loops. Managing these loops requires use of protocols that ensure that the Layer 2 topology remains loop-free. Extending Spanning Tree Protocol (STP) between the sites is one example. OTV has a native dual-homing mechanism that allows traffic to be load-balanced to the transport network. By not extending STP across the overlay, each site is allowed to remain an independent spanning-tree domain.

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



To understand the operation of OTV, it is important to establish some of the key terms and concepts.

OTV is an edge function. Layer 2 traffic is received from the switched network. For all VLANs that need to be extended to remote locations, the Ethernet frames are dynamically encapsulated into IP packets that are then sent across the transport infrastructure. A device that performs the OTV encapsulation and decapsulation functions between the Layer 2 network and the transport network is an OTV edge device. OTV edge devices are responsible for all OTV functionality. The OTV edge device can be in either the core or aggregation layer of the data center. A site can have multiple edge devices to provide additional resiliency. This scenario is commonly referred to as multihoming.

By definition, an OTV edge device has interfaces that connect to the transport network and interfaces that connect to the Layer 2 switched network. The Layer 2 interfaces that receive the traffic from the VLANs that are to be extended across OTV are named the internal interfaces. These interfaces are regular Layer 2 interfaces, usually IEEE 802.1Q trunks. No OTV-specific configuration is required on the internal interfaces.

https://t.me/learningnets

OTV Terminology (Cont.)

- Join interface:
  - The join interface is one of the uplinks of the edge device.
  - The join interface is a routed point-to-point link.
    - Can be a single routed port, a routed port channel, or a subinterface of a routed port or port channel
  - The join interface is used to join the overlay network.
- Overlay interface:
  - The overlay interface is a new virtual interface that contains all the OTV configuration.
  - The overlay interface is a logical multiaccess, multicast-capable interface.
  - The overlay interface encapsulates the site Layer 2 frames in IP unicast or multicast packets.

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

The join interface is used to source the OTV encapsulated traffic and send it to the Layer 3 domain of the data center network. The join interface is a Layer 3 entity.

With the current release of the Cisco NX-OS Software, the join interface can only be defined as a routed point-to-point physical or logical interface. It can be a physical routed port or subinterface of a routed port. For additional resiliency, it can also be a Layer 3 port channel or subinterface of a Layer 3 port channel. Currently the join interface cannot be any other type of interface, such as a switch virtual interface (SVI) or loopback interface.

An OTV overlay can only have a single join interface per edge device. Multiple overlays can share the same join interface.

The edge device uses the join interface for different purposes. The join interface is used to join the overlay network and discover the other remote OTV edge devices. When multicast is enabled in the transport infrastructure, an edge device joins specific multicast groups in order to join the overlay network. These groups are available in the transport network and are dedicated to carry control and data plane traffic. The join interface is used to form OTV adjacencies with the other OTV edge devices belonging to the same overlay VPN. The join interface is used to send and receive MAC reachability information and the join interface is also used to send and receive the encapsulated Layer 2 traffic.

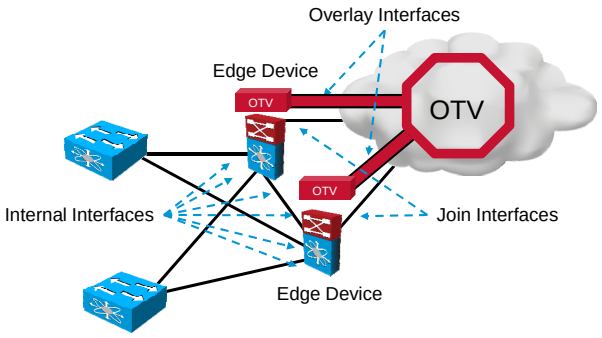
The Overlay interface is a logical multiaccess and multicast-capable interface that must be explicitly defined by the user. The entire OTV overlay configuration is applied on this logical interface. Every time the OTV edge device receives a Layer 2 frame that is destined for a remote data center site, the frame is logically forwarded to the overlay interface. This behavior causes the edge device to perform the dynamic OTV encapsulation on the Layer 2 frame and send it to the join interface toward the routed domain.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

OTV Terminology (Cont.)

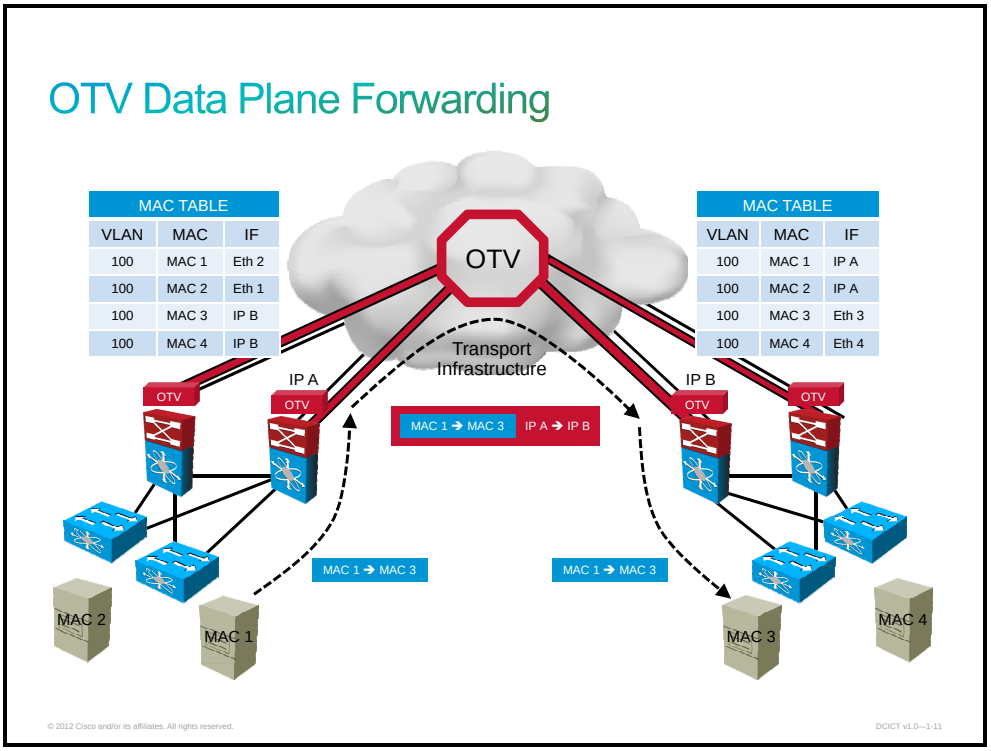
The figure shows the OTV components:



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

The figure illustrates the OTV components, including two edge devices with internal and join interfaces labeled. It also shows a graphical representation of the OTV overlay and the overlay interfaces, which are depicted as tunnels.

https://t.me/learningnets



In OTV data plane forwarding operations, control plane adjacencies are established between the OTV edge devices in different sites and MAC address reachability information is exchanged. Then the MAC address tables in the OTV edge devices contain two different types of entries. The first type is the normal MAC address entry that points to a Layer 2 switch port. The second type in the MAC address table includes entries that point to IP addresses of adjacent OTV neighbors.

To forward frames that are based on the MAC entries that are installed by OTV the following procedure is followed, which is illustrated in the figure:

Step 1 The Layer 2 frame is received at the OTV edge device. A traditional Layer 2 lookup is performed. However, this time the information in the MAC address table for MAC address MAC 3 does not point to a local Ethernet interface. Rather, it points to the IP address of the remote OTV edge device that advertised the MAC reachability information for MAC 3.

Step 2 The OTV edge device encapsulates the original Layer 2 frame. The source IP address of the outer header is the IP address of its join interface. The destination IP address is the IP address of the join interface of the remote edge device.

Step 3 The OTV encapsulated frame, which is a regular unicast IP packet, is carried across the transport infrastructure and delivered to the remote OTV edge device.

Step 4 The remote OTV edge device de-encapsulates the IP packet exposing the original Layer 2 frame.

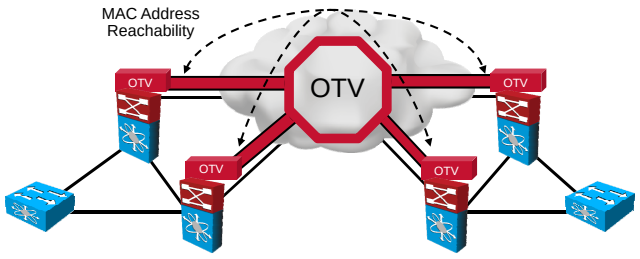
Step 5 The edge device performs a Layer 2 lookup on the original Ethernet frame. It discovers that it is reachable through a physical interface, which means it is a MAC address local to the site.

Step 6 The frame is delivered to the destination host that MAC 3 belongs to through regular Layer 2 Ethernet switching.

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

OTV Control Plane

- The OTV control plane proactively advertises MAC address reachability.
  - The MAC addresses are advertised in the background once OTV has been configured.
  - No specific configuration is required.
- IS-IS is used by OTV as the control protocol between the edge devices.
  - There is no need to configure or understand the operation of IS-IS.



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

Before any encapsulation and de-encapsulation of frames across the overlay can be performed, it is necessary to exchange MAC address reachability information between the sites. This step is required to create corresponding OTV MAC address table entries throughout the overlay.

OTV does not depend on flooding to propagate MAC address reachability information. Instead, OTV uses a control plane protocol to distribute MAC address reachability information to remote OTV edge devices. This protocol runs as an overlay control plane between OTV edge devices. As a result, there is no dependency with the routing protocol used in the Layer 3 domain of the data center or in the transport infrastructure.

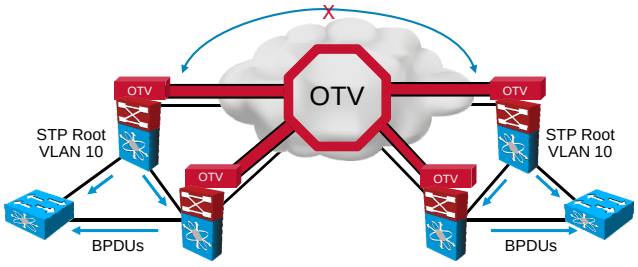
The OTV control plane is transparently enabled in the background after creating the OTV overlay interface and does not require explicit configuration. While it is possible to tune parameters, such as timers, for the OTV protocol, such tuning is more of an exception than a common requirement.

Note The routing protocol that is used to implement the OTV control plane is Intermediate System-to-Intermediate System (IS-IS). It was selected because it is a standards-based protocol, originally designed with the capability of carrying MAC address information in type, length, value (TLV) triplets. Despite the fact that IS-IS is used, the control plane protocol will be generically called “OTV protocol.” This use will differentiate it from IS-IS that is used as an interior gateway protocol (IGP) for IP version 4 (IPv4) or IP version 6 (IPv6). It is not necessary to have a working knowledge of IS-IS configuration to implement OTV. However, some background in IS-IS can be helpful when you troubleshoot OTV.

https://t.me/learningnets

Spanning Tree and OTV

- OTV is site-transparent for STP:
  - Each site maintains its own STP topology with its own root bridges even if the Layer 2 domain is extended across the sites.
  - An OTV edge device only sends and receives BPDUs on internal interfaces.
  - This mechanism is built into OTV and requires no additional configuration.



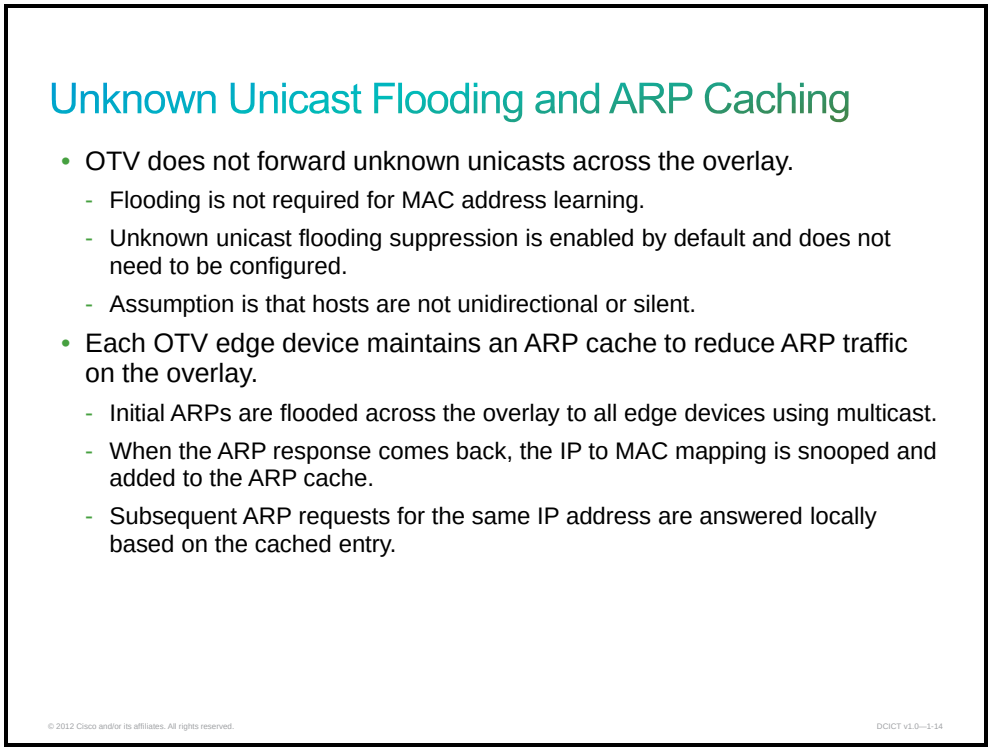
© 2012 Cisco and/or its affiliates. All rights reserved.

OTV by default does not transmit STP bridge protocol data units (BPDUs) across the overlay. This native OTV function does not require the use of any explicit configuration, such as BPDU filtering. This feature allows every site to remain an independent spanning-tree domain: Spanning-tree root configuration, parameters, and the spanning-tree protocol flavor can be decided on a per-site basis.

The separation of spanning-tree domains fundamentally limits the fate sharing between data center sites. A spanning-tree problem in the control plane of a given site would not produce any effect on the remote data centers.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Traditional Layer 2 switching relies on flooding of unknown unicasts for MAC address learning. However, in OTV, this function is performed by the OTV control protocol. By default OTV suppresses the flooding of unknown unicast frames across the overlay. An OTV edge device behaves more like a router than a Layer 2 bridge. It forwards Layer 2 traffic across the overlay only if it has previously received information on how to reach that remote MAC destination. Unknown unicast suppression is enabled by default and does not need to be configured.

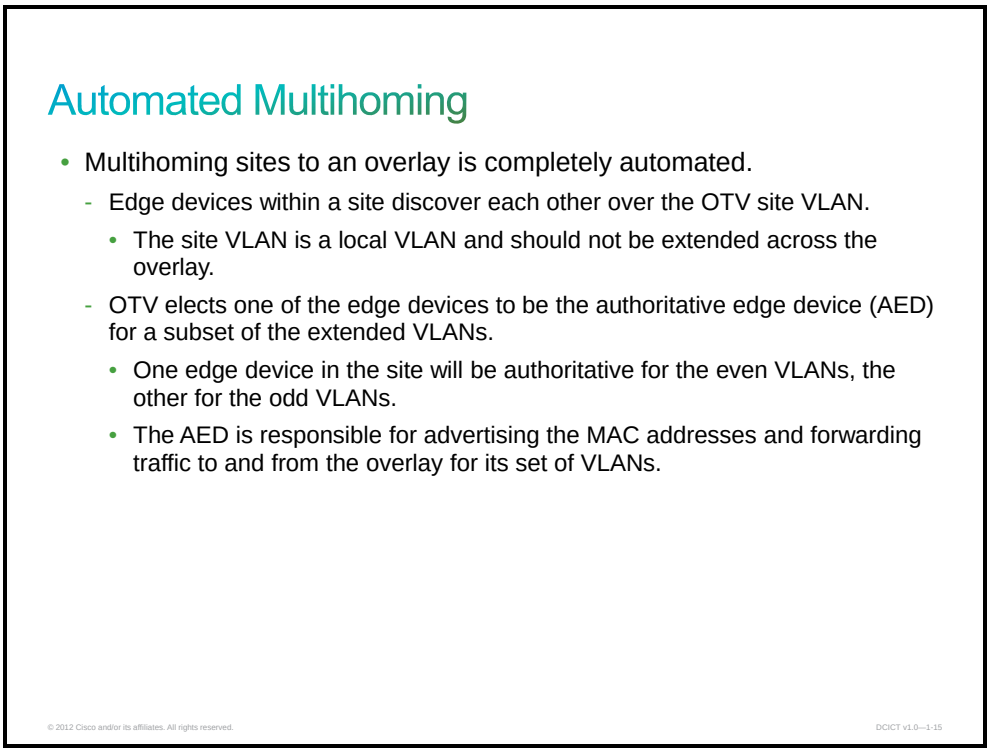
Note	This property of OTV is important to minimize the effects of a server misbehaving and generating streams that are directed to random MAC addresses. This type of behavior could also occur as a result of a denial of service (DoS) attack.

The assumption in the behavior of OTV is that there are no silent or unidirectional devices in the network. It is assumed that sooner or later an OTV edge device will learn the address of a host and communicate it to the other edge devices through the OTV protocol. To support specific applications, like the Microsoft Network Load Balancing (NLB) service, which requires the flooding of Layer 2 traffic to function, a configuration knob will be provided in a future Cisco NX-OS Software release to enable selective flooding. Statically defined MAC addresses allow Layer 2 traffic that is destined to those MAC addresses to be flooded across the overlay, or broadcast to all remote OTV edge devices, instead of being dropped. The expectation is that this configuration will be required only in very specific corner cases, so that the default behavior of dropping unknown unicast would be the usual operation model.

Another function that reduces the amount of traffic that is sent across the transport infrastructure is Address Resolution Protocol (ARP) optimization. Each OTV edge device inspects ARP traffic as it is passed between the internal interfaces and the overlay. Initially, ARP requests received on internal interfaces are broadcast on the overlay to all other edge devices. However, when the ARP response is received on the overlay, the edge router enters the IP and MAC combination into an ARP cache. This is important when subsequent ARP requests are received on an internal interface for an IP address that is present in the cache. In such cases,

https://t.me/learningnets

the OTV edge device will respond to the ARP request on the internal interface. It will not broadcast the request to the overlay again.



One key function included in the OTV protocol is multihoming, where two or more OTV edge devices provide LAN extension services to a given site. Consider this redundant node deployment along with the fact that STP BPDUs are not sent across the OTV overlay. Such a condition could potentially lead to the creation of a bridging loop between sites. To prevent this loop, OTV has a built-in mechanism that ensures only one of the edge devices will forward traffic for a given VLAN. The edge device that has the active forwarding role for the VLAN is called an authoritative edge device (AED) for that VLAN.

The AED has two main tasks:

- Forwarding Layer 2 unicast, multicast, and broadcast traffic between the site and the overlay and vice versa
- Advertising MAC address reachability information to the remote edge devices

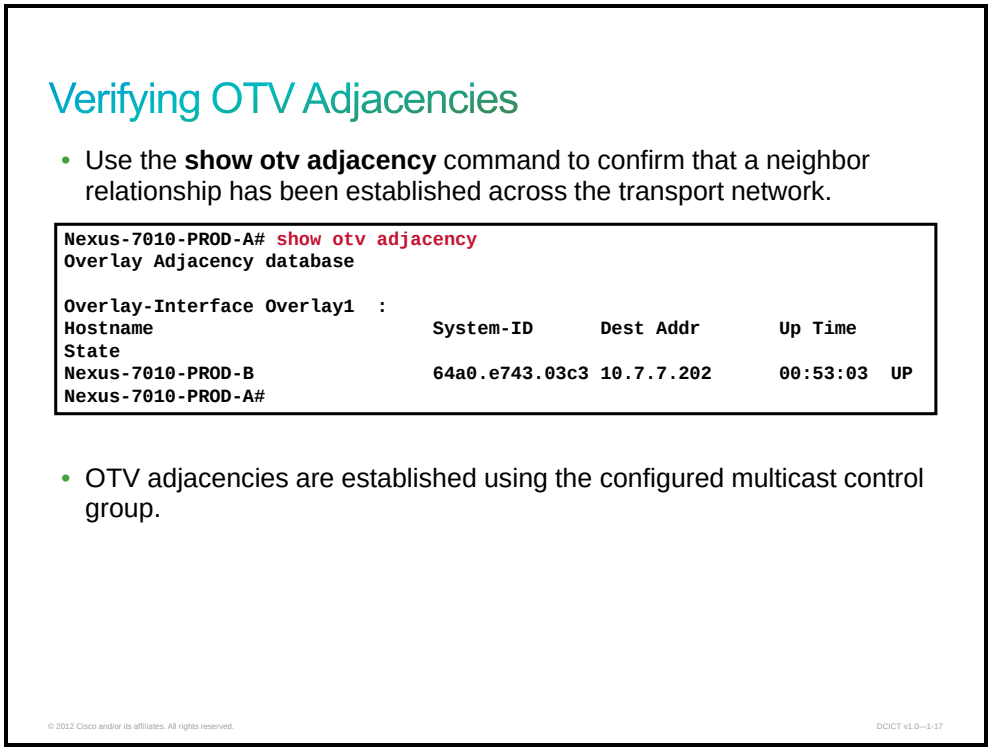
The AED role is negotiated, on a per-VLAN basis, between all the OTV edge devices belonging to the same site. To decide which device should be elected as an AED for a given site, the OTV edge devices establish an internal OTV control protocol peering.

The internal adjacency is established on a dedicated VLAN, named the site VLAN. The Site VLAN should be carried on multiple Layer 2 paths internal to a given site, to increase the resiliency of this internal adjacency. The internal adjacency is used to negotiate the AED role. A deterministic algorithm is implemented to split the AED role for odd and even VLANs between two OTV edge devices. More specifically, the edge device that is identified by a lower system ID will become authoritative for all the even extended VLANs. The device with the higher system ID will be authoritative for the odd extended VLANs. This behavior is hardware-enforced and cannot be tuned in the current Cisco NX-OS Software release.

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

Verifying OTV on the Cisco Nexus 7000 Series Switches

OTV relies on being able to create neighbor adjacencies and establish a routing table showing availability of devices at various sites. This topic explains using **show** commands to verify OTV on the Cisco Nexus 7000 Series Switches.



The **show otv adjacency** command can be used to verify whether the OTV control protocol adjacencies have been properly established. The OTV adjacencies are formed using the OTV control group for the overlay.

If adjacencies fail to establish, verify that the configured control group is the same on all the overlay interfaces and that the overlay interface itself is operational.

https://t.me/learningnets

Verifying OTV Adjacencies (Cont.)

- Use the **show otv overlay** command to verify that the overlay interface is enabled.

```
Nexus-7010-PROD-A# show otv overlay 1

OTV Overlay Information
Site Identifier 0000.0000.2010

Overlay interface Overlay1

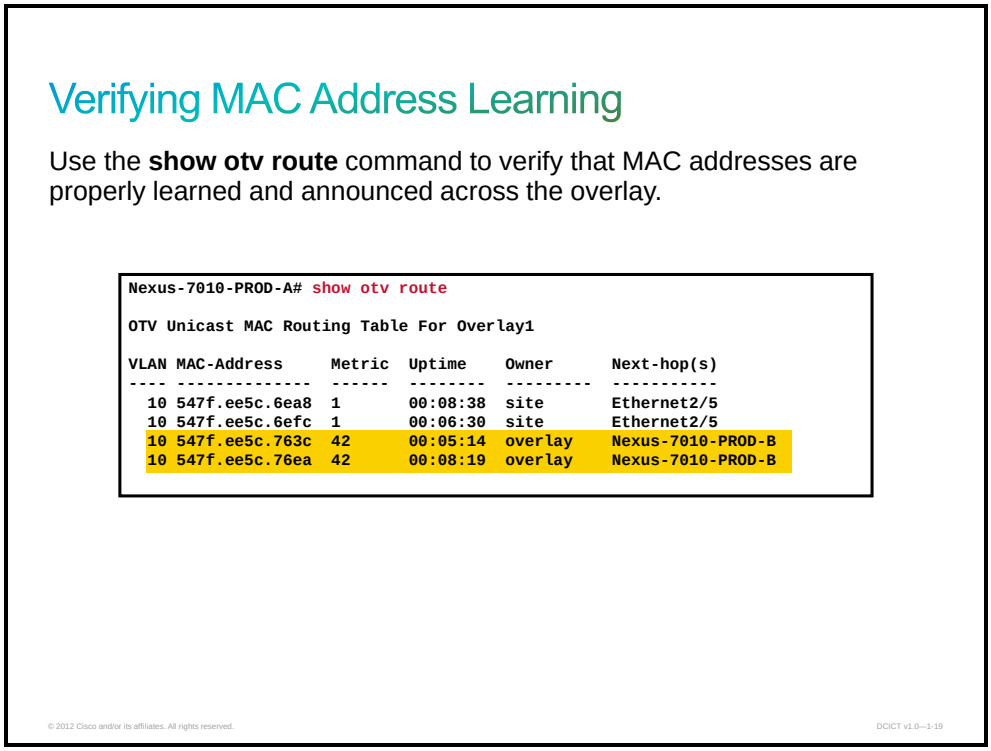
 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-12 (Total:3)
 Control group       : 239.7.7.7
 Data group range(s) : 232.7.7.0/24
 Join interface(s)   : Eth3/1 (10.7.7.201)
 Site vlan           : 13 (up)
 AED-Capable         : Yes
 Capability          : Multicast-Reachable
```

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

To verify that the overlay interface is operational and that the correct parameters have been configured for the overlay interface, you can use the **show otv overlay** command. The primary fields to observe in the output of this command are the “VPN state” and “Control group” fields. The VPN state should be UP and the control group address should match on the local and remote edge device.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Once you have verified that the adjacencies were established, you should verify that the MAC addresses of the target hosts are properly advertised across the overlay. Use the **show otv route** command to see all the MAC addresses that are learned for the extended VLANs. This command shows both local addresses that were learned on the internal interfaces and remote addresses that were learned through the overlay.

Tip	The **show otv route** command only shows MAC addresses for VLANs that are extended on the overlay. Sometimes a MAC address is displayed in the **show mac address-table** command for a VLAN on an internal interface, but the **show otv route** command does not show the address. If this occurs, you should verify that the VLAN was added to the list of extended VLANs for the overlay.

The OTV IS-IS database can be examined using the **show otv isis database** command.

https://t.me/learningnets

Verifying OTV Forwarding

To verify that packets are sent and received on the overlay, use the **show interface overlay** command.

```
Nexus-7010-PROD-A# show interface overlay 1
Overlay1 is up
    MTU 1400 bytes, BW 1000000 Kbit
    Encapsulation OTV
    Last link flapped 00:58:36
    Last clearing of "show interface" counters never
    Load-Interval is 5 minute (300 seconds)
    RX
      0 unicast packets  1473 multicast packets
      1668206 bytes  4709 bits/sec  0 packets/sec
    TX
      0 unicast packets  0 multicast packets
      0 bytes  0 bits/sec  0 packets/sec
```

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

One of the OTV-specific commands that can be helpful in establishing if packets are sent or received on the overlay is the **show interface overlay** command.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

Summary

- OTV is a DCI technology that provides a MAC-in-IP solution with less complexity than traditional DCI solutions.
- OTV configuration is straightforward, and verification that neighbor relationships and routes are in place can be confirmed using various **show otv** commands.

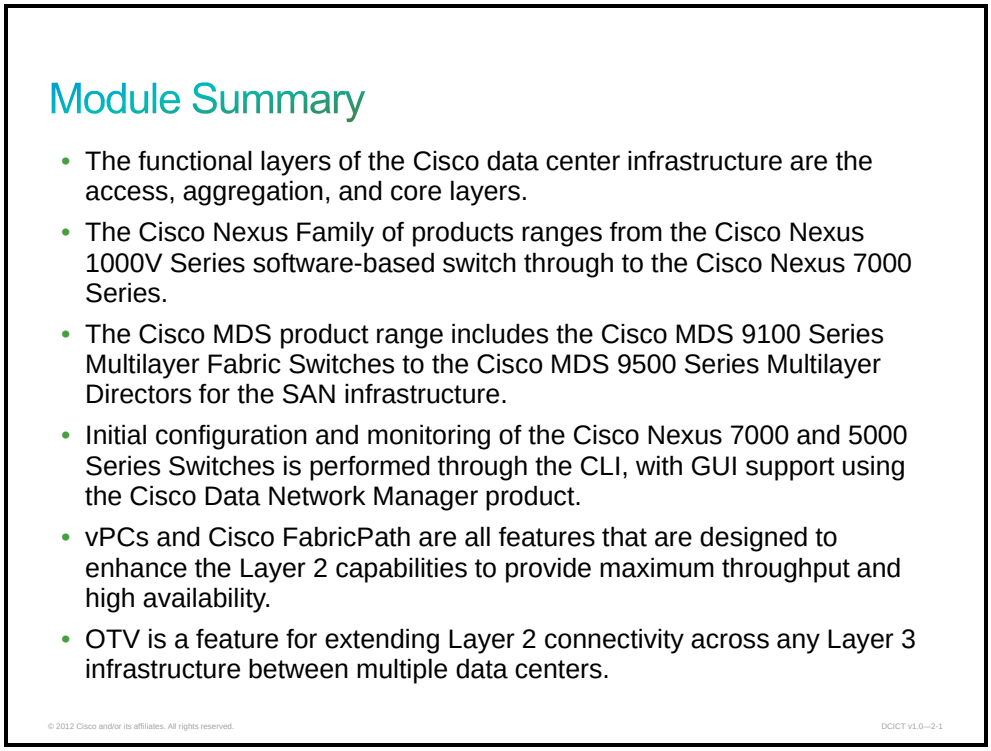
© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

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

Module Summary

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



The Cisco data center is broken down into functional layers. These layers consist of the access, aggregation, and core. Inside the data center, there are often multiple networks including the LAN and SAN networks. Each of these networks uses these functional layers to break down the network logically. The access layer usually provides connectivity to devices such as servers and users. The aggregation layer is usually the Layer 3 demarcation point and where the policies and services are set. The core layer usually provides high-speed connectivity between multiple aggregation layers.

The Cisco Nexus Series product range has been designed for the data center with the Cisco NX-OS Software providing the full range of features that are required to ensure high availability, performance, reliability, security, and manageability for these devices. The Cisco Nexus Series includes the Cisco Nexus 1000V Series (a software-based switch for the virtualized access layer), and the Cisco Nexus 5000 and 7000 Series Switches. To provide additional ports without increasing management points, there is also the Cisco Nexus 2000 Fabric Extender. The Cisco Nexus 2000 Fabric Extender can be connected to either a Cisco Nexus 5000 or 7000 Series switch as an external I/O module.

The Cisco MDS product range is a Fibre Channel switch that is designed for the SAN infrastructure inside the data center. The Cisco MDS Series runs the Cisco NX-OS Software, the same as the Cisco Nexus product range. This process provides consistency within the data center for switch operating systems, making it easier from a management perspective because there are fewer operating systems for administrators to learn. The Cisco MDS Series includes the Cisco MDS 9100 Series through to the Cisco MDS 9500 Series. The Cisco MDS 9500 Series is a modular product and supports a Fibre Channel over Ethernet (FCoE) I/O module. The FCoE module provides customers with the ability to consolidate I/O across 10-Gb/s Ethernet ports.

https://t.me/learningnets

The primary method of monitoring the Cisco Nexus 7000 and 5000 Series switches is through the CLI. Initially the switch has no configuration, only a default administrative user called admin. The first configuration that must be performed is to provide the admin user with an administrative password. Once the administrative password has been set, the administrator would then configure the basic parameters that are required to provide network connectivity. To assist the administrator, there is a setup dialog script to walk the administrator through the basic requirements.

Inside the data center, there will be Layer 2 and Layer 3 connectivity. At the access layer there is often Layer 2 connectivity only, with Layer 3 being provided at the aggregation layer. With any Layer 2 domain, normally the Spanning Tree Protocol (STP) is running. STP is designed to create a loop-free tree-like topology. This principle means that inside the data center, although multiple paths are provided, certain paths will be forwarding and others blocking. Features such as virtual port channels (vPCs) and Cisco FabricPath help avoid this issue where full utilization of the available bandwidth is not possible. vPCs allow a downstream device to be dual-homed to two upstream switches that are connected through a vPC link using a multichassis port channel. This multichassis port channel is known as a vPC on Cisco Nexus Series switch devices. Cisco FabricPath provides a Layer 2 routing capability to replace STP.

As the virtualization layer grows, the requirement for extending Layer 2 between geographically dispersed data centers is growing. When extending Layer 2 between data centers, there are usually issues that need to be resolved such as the size of the Layer 2 fault domain. There are several technologies available to provide this functionality, but they often have specific requirements and can be difficult to manage. Overlay Transport Virtualization (OTV) is a technology that allows the extension of Layer 2 but without the overheads and issues of large Layer 2 fault domains.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Cisco MDS 9000 Series Configuration Guides: http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html
- Cisco Systems, Inc. Cisco Nexus 7000 Series NX-OS Software configuration guides: http://www.cisco.com/en/US/products/ps9402/products_installation_and_configuration_guides_list.html
- Cisco Systems, Inc. Cisco Nexus 7000 Series Switches Release Notes: http://www.cisco.com/en/US/products/ps9402/prod_release_notes_list.html
- Cisco Systems, Inc. Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide: *http://www.cisco.com/en/US/docs/switches/datacenter/nexus2000/sw/configuration/guide/rel_6_0/b_Configuring_the_Cisco_Nexus_2000_Series_Fabric_Extender_rel_6_0.html*
- Cisco Systems, Inc. Cisco Nexus 5000 Series NX-OS Software configuration guides: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html
- Cisco Systems, Inc. Cisco Nexus 5000 Series Switches Release Notes: http://www.cisco.com/en/US/products/ps9670/prod_release_notes_list.html
- Cisco Systems, Inc. Cisco Nexus 7000 Series Switches Data Sheets: http://www.cisco.com/en/US/products/ps9402/products_data_sheets_list.html
- Cisco Systems, Inc. Cisco Nexus 5000 Series Switches Data Sheets: http://www.cisco.com/en/US/products/ps9670/products_data_sheets_list.html

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

- Cisco Systems, Inc. Cisco MDS 9500 Series Multilayer Directors Data Sheets: http://www.cisco.com/en/US/products/ps5990/products_data_sheets_list.html
- Cisco Systems, Inc. Cisco MDS 9100 Series Multilayer Fabric Switches Data Sheets: http://www.cisco.com/en/US/products/ps5987/products_data_sheets_list.html
- Cisco Systems, Inc. Cisco MDS 9200 Series Multilayer Switches Data Sheets: http://www.cisco.com/en/US/products/ps5988/products_data_sheets_list.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 of these data center-related components typically would not be used in the SAN infrastructure design? (Source: Examining Cisco Data Center Functional Layers)

A) access
B) aggregation
C) core
D) collapsed core

Q2) Which type of flow control is used in the SAN environment? (Source: Examining Cisco Data Center Functional Layers)

A) Credit-based flow control is used.
B) The sender controls when frames are sent.
C) Flow control is not required.
D) Sliding window-based flow control is used.

Q3) FCoE is part of which main component of the data center architecture? (Source: Examining Cisco Data Center Functional Layers)

A) virtualization
B) unified fabric
C) unified computing
D) unified data center

Q4) Which of these provides the most efficient use of ports in a SAN infrastructure because fewer or no ports are consumed for ISLs? (Source: Examining Cisco Data Center Functional Layers)

A) core-edge
B) multitier
C) collapsed aggregation
D) collapsed core

Q5) Which Cisco Nexus product supports unified ports? (Source: Reviewing the Cisco Nexus Product Family)

A) Cisco Nexus 7000 Series Switches
B) Cisco Nexus 5000 Series Switches
C) Cisco Nexus 5500 Platform switches
D) Cisco Nexus 2000 Series Fabric Extenders

Q6) Which product supports a front-to-back airflow to address the requirements for hot-aisle and cold-aisle deployments without additional complexity? (Source: Reviewing the Cisco Nexus Product Family)

A) Cisco Nexus 9-Slot Switch
B) Cisco Nexus 10-Slot Switch
C) Cisco Nexus 18-Slot Switch
D) Cisco Nexus 3000 Series Switch

https://t.me/learningnets

Q7) Which license is required to support VDCs on the Cisco Nexus 7000 Series Switches? (Source: Reviewing the Cisco Nexus Product Family)

A) Base license
B) Enterprise license
C) Advanced LAN Enterprise license
D) Scalable Services license

Q8) How many ports are there on the Cisco Nexus 7000 40 Gigabit Ethernet module? (Source: Reviewing the Cisco Nexus Product Family)

A) 2
B) 4
C) 6
D) 8

Q9) Which power supply redundancy mode provides grid redundancy only? (Source: Reviewing the Cisco Nexus Product Family)

A) combined
B) power supply and input source redundancy
C) power supply redundancy
D) input source redundancy

Q10) How many expansion slots are there for the Cisco Nexus 5548 Switch? (Source: Reviewing the Cisco Nexus Product Family)

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

Q11) How many ports on the Cisco Nexus 5010 Switch support 1 Gb/s connectivity on the base module? (Source: Reviewing the Cisco Nexus Product Family)

A) 4
B) 8
C) 12
D) 16

Q12) How many Cisco Nexus 2000 Series Fabric Extenders can be connected to a Cisco Nexus 5500 Platform switch with no Layer 3 daughter card or expansion module? (Source: Reviewing the Cisco Nexus Product Family)

A) 8
B) 12
C) 24
D) 32

Q13) Which product has a fixed module and one slot for an I/O module? (Source: Reviewing the Cisco MDS Product Family)

A) Cisco MDS 9148 Multilayer Fabric Switch
B) Cisco MDS 9506 Multilayer Director Switch
C) Cisco MDS 9124 Multilayer Fabric Switch
D) Cisco MDS 9222i Multiservice Modular Switch

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

Q14) Which supervisor module meets the minimum requirements that are required to support FCoE on the Cisco MDS 9000 Series Switches? (Source: Reviewing the Cisco MDS Product Family)

A) Supervisor-3
B) Supervisor-2A
C) Supervisor-2
D) Supervisor-1

Q15) What is the total aggregate bandwidth that is provided on the Cisco MDS 9513 Multilayer Director switch? (Source: Reviewing the Cisco MDS Product Family)

A) 760 Gb/s
B) 1.1 Tb/s
C) 2.2 Tb/s
D) 4.1 Tb/s

Q16) For how many days may licensed features be evaluated on the Cisco MDS Series switch? (Source: Reviewing the Cisco MDS Product Family)

A) 120
B) 90
C) 60
D) 30

Q17) How many ports are enabled in the port-based license on a Cisco MDS 9124 Multilayer Fabric Switch? (Source: Reviewing Cisco MDS Product Family)

A) 4
B) 8
C) 16
D) 24

Q18) Which management port on the Cisco Nexus 7000 Series Switches provides lights-out RMON and management without the need for separate terminal servers? (Source: Monitoring the Cisco Nexus 7000 and 5000 Series Switches)

A) console port
B) management 0 port
C) any physical interface
D) CMP

Q19) Which command would be used to connect to the CMP from the control processor on a Cisco Nexus 7000 Series switch? (Source: Monitoring the Cisco Nexus 7000 and 5000 Series Switches)

A) switch# **attach cmp**
B) switch# **connect cmp**
C) switch# **cmp attach**
D) switch# **cmp connect**

Q20) Which command would be used to identify the impact of an ISSU event? (Source: Monitoring the Cisco Nexus 7000 and 5000 Series Switches)

A) switch# **show all install impact**
B) switch# **show install all**
C) switch# **show impact install all**
D) switch# **show install all impact**

https://t.me/learningnets

Q21) Which command is a valid command for viewing the details of a 10-Gigabit Ethernet interface? (Source: Monitoring the Cisco Nexus 7000 and 5000 Series Switches)

A) switch# **show interface Gigabit Ethernet 1/3**
B) switch# **show interface 10Gigabit Ethernet 1/3**
C) switch# **show interface Ethernet 1/3**
D) switch# **show interface**

Q22) Which of these best describes the function of Cisco Fabric Services in a vPC domain? (Source: Describing vPCs and Cisco FabricPath in the Data Center)

A) synchronizes vPC control plane information
B) synchronizes vPC data plane information
C) synchronizes vPC control and data plane information
D) assists in the election of the vpc primary switch

Q23) Which command would you use to verify the status of configured virtual port channels on the switch? (Source: Describing vPCs and Cisco FabricPath in the Data Center)

A) switch# **show vpc summary**
B) switch# **show vpc brief**
C) switch# **show vpc database**
D) switch# **show vpc global**

Q24) Which link is used in a virtual port channel to create the illusion of a single control plane? (Source: Describing vPCs and Cisco FabricPath in the Data Center)

A) vPC peer
B) vPC peer keepalive
C) vPC peer link
D) vPC link

Q25) Which command would be used to find specific parameters that caused a consistency check to fail during a vPC configuration? (Source: vPCs and Cisco FabricPath in the Data Center)

A) switch# **show vpc interface**
B) switch# **show vpc global**
C) switch# **show vpc brief**
D) switch# **show vpc consistency-parameters global**

Q26) Cisco FabricPath provides ECMP capabilities. How many ECMP paths does it support? (Source: Describing vPCs and Cisco FabricPath in the Data Center)

A) 4
B) 6
C) 8
D) 16

Q27) Which protocol is used as the control protocol for Cisco FabricPath? (Source: Describing vPCs and Cisco FabricPath in the Data Center)

A) Layer 2 IS-IS
B) Layer 3 IS-IS
C) Layers 2 and 3 IS-IS
D) IS-IS

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

Module Self-Check Answer Key

Q1) B
Q2) A
Q3) B
Q4) D
Q5) C
Q6) B
Q7) C
Q8) C
Q9) D
Q10) A
Q11) B
Q12) C
Q13) D
Q14) B
Q15) C
Q16) A
Q17) B
Q18) D
Q19) A
Q20) D
Q21) C
Q22) B
Q23) A
Q24) C
Q25) D
Q26) D
Q27) A

https://t.me/learningnets

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

Module 2

Cisco Data Center Virtualization

Overview

This module describes the function and benefits of Cisco data center virtualization.

Module Objectives

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

- Describe network device virtualization
- Describe how RAID groups and LUNs virtualize storage for high availability and configuration flexibility
- Describe the benefits of server virtualization in the data center
- Describe the problems that Cisco Nexus 1000V Series switches solve and how the Cisco Nexus 1000V VSM and VEM integrate with VMware ESX
- Validate connectivity of the Cisco Nexus 1000V VSM to VEMs and VMware vCenter, by using the VMware ESX and Cisco Nexus 1000V CLIs

https://t.me/learningnets

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

Lesson 1

Virtualizing Network Devices

Overview

The purpose of this lesson is to describe the virtualization capabilities of Cisco Nexus 7000 and 5000 Series Switches. Using virtual device contexts (VDCs) and Network Interface Virtualization (NIV), Cisco Nexus 5000 and 7000 Series Switches allow you to lower complexity and cost of ownership of data center network and storage resources.

Objectives

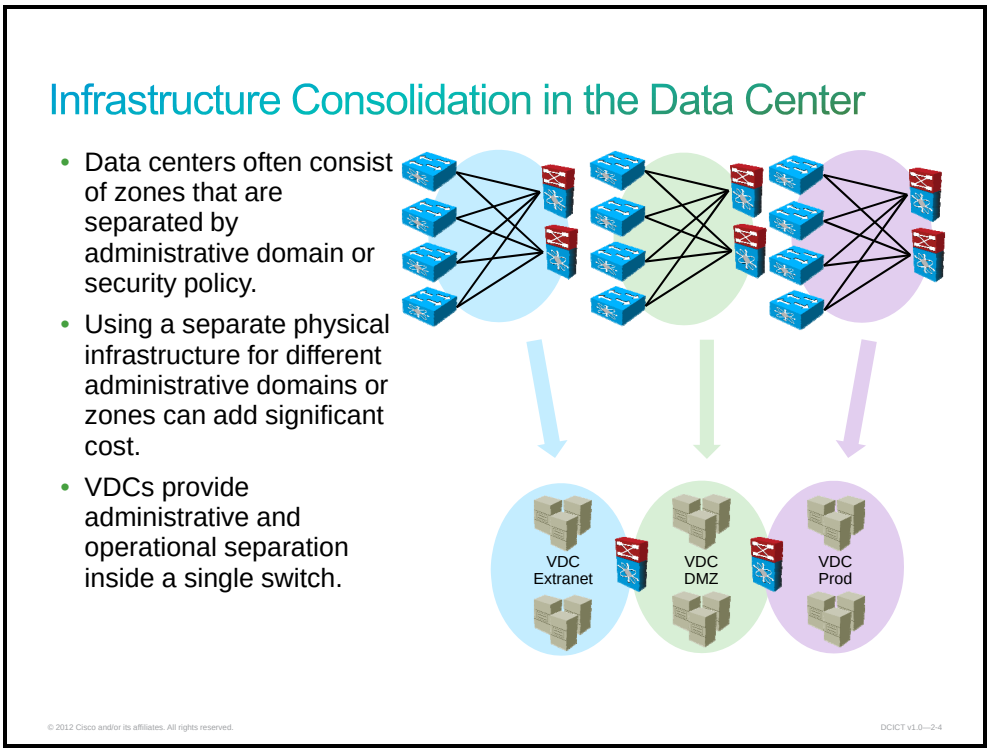
Upon completing this lesson, you will be able to describe the virtualization capabilities of Cisco Nexus 7000 and 5000 Series Switches. You will be able to meet these objectives:

- Describe VDCs on the Cisco Nexus 7000 Series switch
- Verify VDCs on the Cisco Nexus 7000 Series switch
- Navigate between VDCs on the Cisco Nexus 7000 Series switch
- Describe NIV on Cisco Nexus 7000 and 5000 Series Switches

https://t.me/learningnets

Describing VDCs on the Cisco Nexus 7000 Series Switch

This topic describes the requirements and functionality of VDCs on the Cisco Nexus 7000 Series switch.



Data centers are often partitioned into separate domains or zones that are implemented on separate physical infrastructures. The creation of separate physical infrastructures is commonly driven by a need to separate administrative domains for security and policy reasons.

VLANs and virtual routing and forwarding (VRF) instances can be used to separate user traffic on the data plane. However, these technologies do not provide separation of administration and management functions or isolation of fault domains.

Building separate physical infrastructures to separate zones by administrative or security policy can add significant cost to the infrastructure. Depending on the port counts and the functions that are needed in the separate domains, the physical switches in each domain might be underutilized. Consolidation of multiple logical switches on one physical switch can improve hardware utilization. Consolidation can also add flexibility to the data center design.

VDCs allow one physical Cisco Nexus 7000 Series switch to be partitioned into multiple logical switches. This partitioning enables the consolidation of different logical zones onto one physical infrastructure.

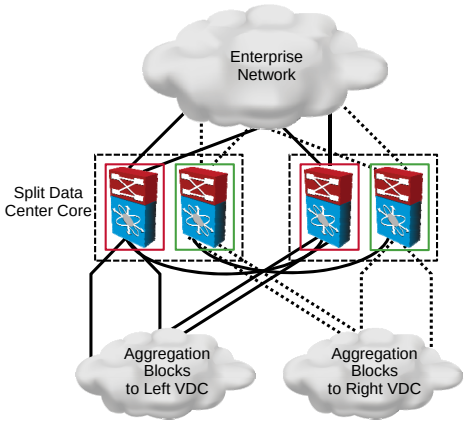
 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

VDC Deployment Scenarios

- VDCs allow multiple logical switches to be implemented on one physical switch.
- Ports can be reallocated between VDCs in a flexible manner.
- Common scenarios:
  - Dual core
  - Multiple aggregation blocks
  - Service insertion

Dual core: Useful in migrations, mergers, and acquisitions



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

A VDC has the same functional characteristics as a physical switch. The VDC can be used in many places in the overall data center network design.

A major advantage of using VDCs instead of separate physical switches is that physical ports can easily be reallocated between VDCs. This capability allows for ease of changes and additions to the network design as the network grows and evolves.

The following scenarios can benefit from the use of VDCs. Because a VDC has characteristics and capabilities like those of a separate physical switch, these scenarios are not VDC-specific topologies; they could be built using separate dedicated switches in roles that are occupied by VDCs. However, VDCs can provide additional design flexibility and efficiency in these scenarios.

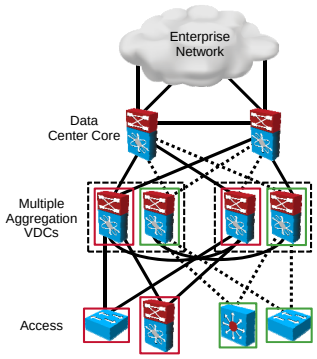
Dual-Core Topology

VDCs can be used to build two redundant data center cores using only a pair of Cisco Nexus 7000 Series Switches. This technique can be useful to facilitate migration when the enterprise network needs to expand to support mergers and acquisitions. If sufficient ports are available on the existing data center core switches, then two additional VDCs can be created for a separate data center core. This approach allows a second data center network to be built alongside the original one. This second network can be built without any impact on the existing network. Eventually, aggregation blocks can be migrated from one core to the other by reallocating interfaces from one VDC to the other.

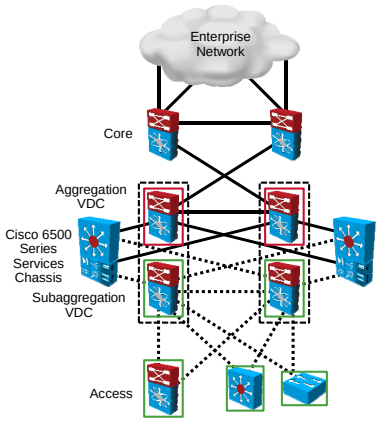
https://t.me/learningnets

VDC Deployment Scenarios (Cont.)

Multiple aggregation blocks: Separation by business unit or function



Service insertion: Separated management and control for access and aggregation layers



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

Multiple Aggregation Blocks

At the aggregation layer of the data center network, a single aggregation block consists of a pair of aggregation switches, for redundancy, and their associated access layer switches. Often, an enterprise has a business requirement to deploy separate aggregation blocks for different business units or functions. The use of VDCs might accomplish this logical segregation without needing to deploy separate physical switches. Administration and management can be delegated to different groups. Configuration changes in the VDC of one aggregation block cannot affect the VDCs for the other aggregation blocks; for example, a separate production and development aggregation block can be built using one pair of aggregation switches.

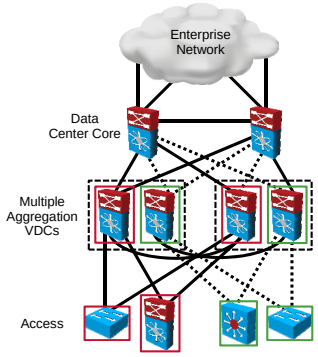
Service Insertion

VRFs are often used to create a Layer 3 hop that separates the servers in the access network from the services in the service chain and the aggregation layer. This approach creates a “services sandwich” consisting of two VRFs with the services chain in between. Instead of VRFs, two VDCs can be used to create this services sandwich. In addition to the control plane and data plane separation that VRFs provide, a VDC provides management plane separation and fault isolation. The VDC services sandwich design increases security by logically separating the switches on the inside and outside of the services chain.

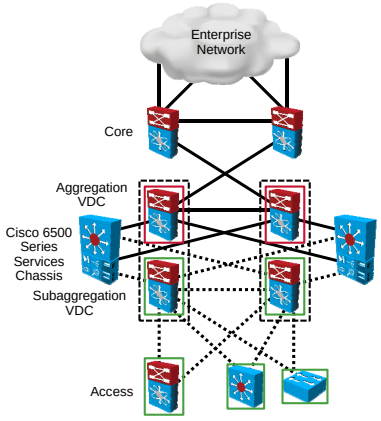
https://t.me/learningnets

VDC Deployment Scenarios (Cont.)

Multiple aggregation blocks: Separation by business unit or function



Service insertion: Separated management and control for access and aggregation layers



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

The Cisco Nexus 7000 Series switch uses several virtualization technologies that are already present in Cisco IOS Software. At Layer 2, you have VLANs and at Layer 3, you have VRFs. These two features are used to virtualize the Layer 3 forwarding and routing tables. The Cisco Nexus 7000 Series switch extends this virtualization concept to VDCs that virtualize the device itself, by presenting the physical switch as multiple, independent logical devices.

Within each VDC is a set of unique and independent VLANs and VRFs, with physical ports assigned to each VDC. This independence also allows virtualization of both the hardware data plane and the separate management domain.

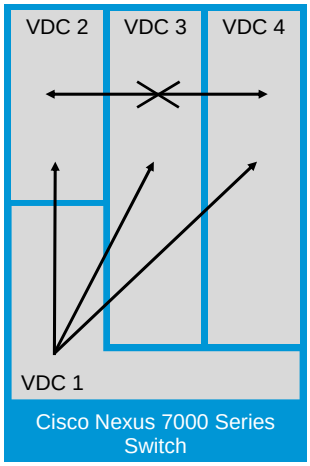
In its default state, the switch control plane runs as a single device context, called VDC 1, which runs approximately 80 processes. Some processes have other spawned threads, which can result in as many as 250 processes actively running on the system at any given time. This collection of processes constitutes what is seen as the control and management plane for a single physical device without any other VDCs enabled. The default VDC 1 is always active, always enabled, and can never be deleted. Even if no other VDC is created, support for virtualization through VRFs and VLANs within VDC 1 is available.

The Cisco Nexus 7000 Series switch can support multiple VDCs. The creation of additional VDCs replicates these processes for each device context that is created. The hardware resources on the supervisor and I/O modules are shared between the VDCs. The processes of the different VDCs share the kernel and infrastructure modules of the Cisco Nexus Operating System (NX-OS) Software, but the processes within each VDC are entirely independent.

https://t.me/learningnets

Virtual Device Contexts

- Cisco Nexus 7000 Series Switches support as many as four possible VDCs per switch.
- VDC 1 is the default VDC.
- The default VDC can create and manage other VDCs:
  - Nondefault VDCs are strictly separated.
  - Ports are allocated to VDCs from the default VDC.
- Default VDC cannot be deleted.
- Default VDC controls shared switch resources.



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

The use of VDCs currently allows one Cisco Nexus 7000 Series switch to be partitioned into as many as four logical switches: the default VDC and three additional VDCs. Initially, all hardware resources of the switch belong to the default VDC. When you first configure a Cisco Nexus 7000 Series switch, you are effectively configuring the default VDC (VDC 1). The default VDC has a special role: It controls all hardware resources and has access to all other VDCs. VDCs are always created from the default VDC. Hardware resources, such as interfaces and memory, are also allocated to the other VDCs from the default VDC. The default VDC can access and manage all other VDCs. However, the additional VDCs have access only to the resources that are allocated to them and cannot access any other VDCs.

VDCs are truly separate virtual switches (vSwitches). They do not share any processes or data structures, and traffic can never be forwarded from one VDC to another VDC inside the chassis. Any traffic that needs to pass between two VDCs in the same chassis must first leave the originating VDC through a port that is allocated to it. The originating VDC then enters the destination VDC through a port that is allocated to that VDC. VDCs are separated on the data, control, and management planes. The only exception to this separation is the default VDC, which can interact with the other VDCs on the management plane. Control and data plane functions of the default VDC are still separated from the other VDCs.

The default VDC has several other unique and crucial roles in the function of the switch:

- Systemwide parameters such as Control Plane Policing (CoPP), VDC resource allocation, and Network Time Protocol (NTP) may be configured from the default VDC.
- Licensing of the switch for software features is controlled from the default VDC.
- Software installation must be performed from the default VDC. All VDCs run the same version of software.
- Reloads of the entire switch may be issued only from the default VDC. Nondefault VDCs may be reloaded independently of other VDCs.

https://t.me/learningnets

If a switch might be used in a multiple-VDC configuration, then the default VDC should be reserved for administrative functions only. Configure all production network connections in nondefault VDCs. This approach will provide flexibility and higher security. Administrative access can easily be granted into the nondefault VDCs to perform configuration functions, without exposing access to reload the entire switch or change software versions. No Layer 3 interfaces in the default VDC need to be exposed to the production data network. Only the management interface needs to be accessible through an out-of-band (OOB) management path. Unused interfaces may be retained in a shutdown state in the default VDC as a holding area until they are needed in the configuration of a nondefault VDC. In this way, the default VDC may be maintained as an administrative context requiring console access or separate security credentials. Following this guideline effectively allows one Cisco Nexus 7000 Series switch to perform the functional roles of as many as three production switches.

VDC Virtualization

- Within each VDC, VLANs and VRFs can be used to provide additional levels of virtualization.
- VLAN numbers and VRF names can be reused within different VDCs.
- VLANs and VRFs in different VDCs are strictly isolated.
- External connections are necessary to forward traffic between VDCs.



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

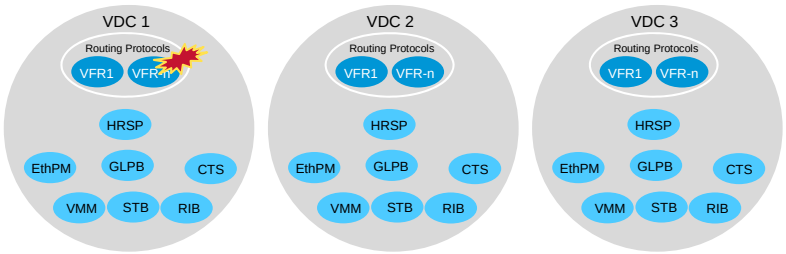
The use of VDCs does not restrict the use of VLANs and VRFs. Within each VDC, you can create VLANs and VRFs. The VLANs and VRFs in a VDC are entirely independent of and isolated from the VLANs and VRFs in any other VDC.

Because VDCs are independent, VLAN numbers and VRF names can be reused in different VDCs. There is no internal connection between the VLANs and VRFs in different VDCs. To connect a VLAN or VRF in one VDC to a VLAN or VRF in a different VDC, an external connection is required. VDCs truly behave as completely separate logical switches.

https://t.me/learningnets

VDC Fault Domains

- Each VDC runs separate processes for control plane and management plane functions, creating a separate fault domain.
- When a process crashes in a VDC, the processes in the other VDCs are unaffected and continue to run unimpeded.



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

When multiple VDCs are created in a physical switch, the architecture of the VDC feature provides a means to prevent failures within any VDC from affecting other VDCs. For example, a spanning-tree recalculation that is started in one VDC does not affect the spanning-tree domains of other VDCs in the same physical chassis. Each recalculation is an entirely independent process. The same applies to other processes, such as the Open Shortest Path First (OSPF) process. Network topology changes in one VDC do not affect other VDCs on the same switch.

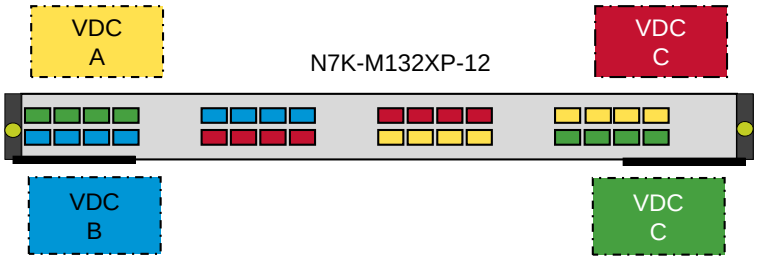
Because the Cisco NX-OS Software uses separate processes in each VDC, the fault isolation extends even to potential software process crashes. If a process crashes in one VDC, then that crash is isolated from other VDCs. The Cisco NX-OS high-availability features, such as stateful process restart, can be applied independently to the processes in each VDC. Process isolation within a VDC is important for fault isolation and is a major benefit for organizations that implement the VDC concept.

In addition, fault isolation is enhanced with the ability to provide per-VDC debug commands and per-VDC logging of messages from syslog. These features give administrators the ability to locate problems within their own VDC.

https://t.me/learningnets

Interface Allocation

- Interfaces on the 32-port N7K-M132XP-12 and F1 and F2 I/O modules must be assigned to VDCs on a per-port group basis.
  - On the N7K-M132XP-12 module, a port group consists of four consecutive odd or even ports.
  - On the F1 and F2 modules, a port group consists of two consecutive ports.



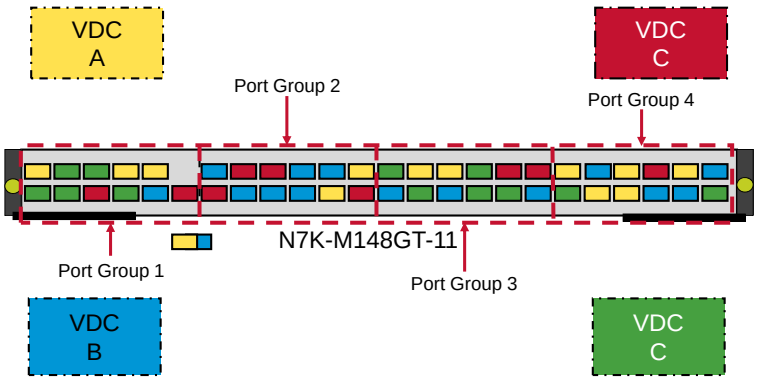
© 2012 Cisco and/or its affiliates. All rights reserved.

Physical ports are allocated to different VDCs from the default VDC. Logical interfaces, such as switch virtual interfaces (SVIs), subinterfaces, or tunnel interfaces, cannot be assigned to a VDC. Logical interfaces are always created in the VDC to which they belong. After a physical port is assigned to a VDC, all subsequent configuration of that port is performed within that VDC. Within a VDC, both physical and logical interfaces can be assigned to VLANs or VRFs.

On most I/O modules, any port can be individually assigned to any VDC. The exceptions to this rule are the N7K-M132XP-12, F1, and F2 I/O modules. On these modules, interfaces can be assigned to a VDC on a per-port-group basis only.

Interface Allocation (Cont.)

- Interfaces on all other I/O modules can be assigned to VDCs on a per-port basis.



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

https://t.me/learningnets

VDC Administration

- By default, the Cisco NX-OS Software has four predefined roles.
- In the default VDC, these roles include the following:
  - **Network-admin:** Has full control of the default VDC and can create, delete, or change nondefault VDCs
  - **Network-operator:** Has read-only rights in the default VDC
- In the nondefault VDCs, these roles include the following:
  - **VDC-admin:** Has full control of a specific VDC, but no rights in the default VDC or other nondefault VDCs
  - **VDC-operator:** Has read-only rights in a specific VDC
- When a network administrator or network operator switches to a nondefault VDC, the same level of rights are inherited:
  - Network administrator has VDC admin rights in nondefault VDCs.
  - Network operator has VDC operator rights in nondefault VDCs.

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

The Cisco NX-OS Software uses role-based access control (RBAC) to control the access rights of users on the switch. By default, Cisco Nexus 7000 Series Switches recognize four roles:

- **Network-admin:** The first user account that is created on a Cisco Nexus 7000 Series switch in the default VDC is admin. The network-admin role is automatically assigned to this user. The network-admin role gives a user complete control over the default VDC of the switch. This role includes the ability to create, delete, or change nondefault VDCs.
- **Network-operator:** The second default role that exists on Cisco Nexus 7000 Series Switches is the network-operator role. This role allows the user read-only rights in the default VDC. The network-operator role includes the right to issue the **switchto** command, which can be used to access a nondefault VDC from the default VDC. By default, no users are assigned to this role. The role must be assigned specifically to a user by a user that has network-admin rights.
- **VDC-admin:** When a new VDC is created, the first user account on that VDC is admin. The VDC-admin role is automatically assigned to this admin user on a nondefault VDC. This role gives a user complete control over the specific nondefault VDC. However, this user does not have any rights in any other VDC and cannot access other VDCs through the **switchto** command.
- **VDC-operator:** The VDC-operator role has read-only rights for a specific VDC. This role has no rights to any other VDC.

When a network-admin or network-operator user accesses a nondefault VDC by using the **switchto** command, that user is mapped to a role of the same level in the nondefault VDC. In other words, a user with the network-admin role is given the VDC-admin role in the nondefault VDC. A user with the network-operator role is given the VDC-operator role in the nondefault VDC.

https://t.me/learningnets

VDC Requirements

- The Cisco NX-OS Advanced Services license is required to create, delete, or modify VDCs.
  - A grace period exists, but when it expires, any VDC configuration is deleted.
- VDCs are created from within the default VDC global configuration context.
  - The network administrator role is required to create, delete, or modify VDCs.
- Physical and logical resources are assigned to VDCs from the default VDC global configuration context.
  - When a physical port is assigned to a VDC, that port can be configured from within that VDC only.

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

Consider the following issues when implementing VDCs. To use VDCs, the Advanced Services license needs to be installed on a Cisco Nexus 7000 Series switch. You can try the feature during a 120-day grace period. However, when the grace period expires, any nondefault VDCs will be removed from the switch configuration. Any existing processes for those VDCs will be terminated.

VDCs can be created, deleted, or changed from the default VDC only. You cannot create VDCs from a nondefault VDC. To create VDCs, a user needs to have network-admin rights in the default VDC.

Physical interfaces and other resources are always assigned to nondefault VDCs from the default VDC. After a physical interface has been assigned to a specific VDC, the configuration for that interface is performed from that nondefault VDC. You cannot configure an interface from any VDC other than the VDC to which the interface is allocated.

https://t.me/learningnets

VDC Requirements (Cont.)

- Initially, all chassis physical resources are part of the default VDC.
- All configuration for the interface is lost when you allocate an interface to another VDC.
- When a VDC is deleted, the operation of the VDC is disrupted and all resources are returned to the default VDC.
- To remove an interface from a nondefault VDC and return it to the default VDC, you must enter VDC configuration mode in the default VDC and allocate the interface to the default VDC.
  - This action cannot be performed from a nondefault VDC.
- **Note:** When performing VDC configuration and modification, it is very important to be in the proper configuration context.

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

Initially, all physical resources are assigned to the default VDC. When interfaces are reallocated to a different VDC, any existing configuration on the interface is removed.

When a VDC is removed, all resources that are associated with that VDC are returned to the default. All processes that belong to the VDC are terminated, and forwarding information for the VDC is removed from the forwarding engines.

You cannot move interfaces from a nondefault VDC to the default VDC from within the nondefault VDC itself. To remove a physical interface from a nondefault VDC, you must enter configuration mode in the default VDC and reallocate the interface to the default VDC.

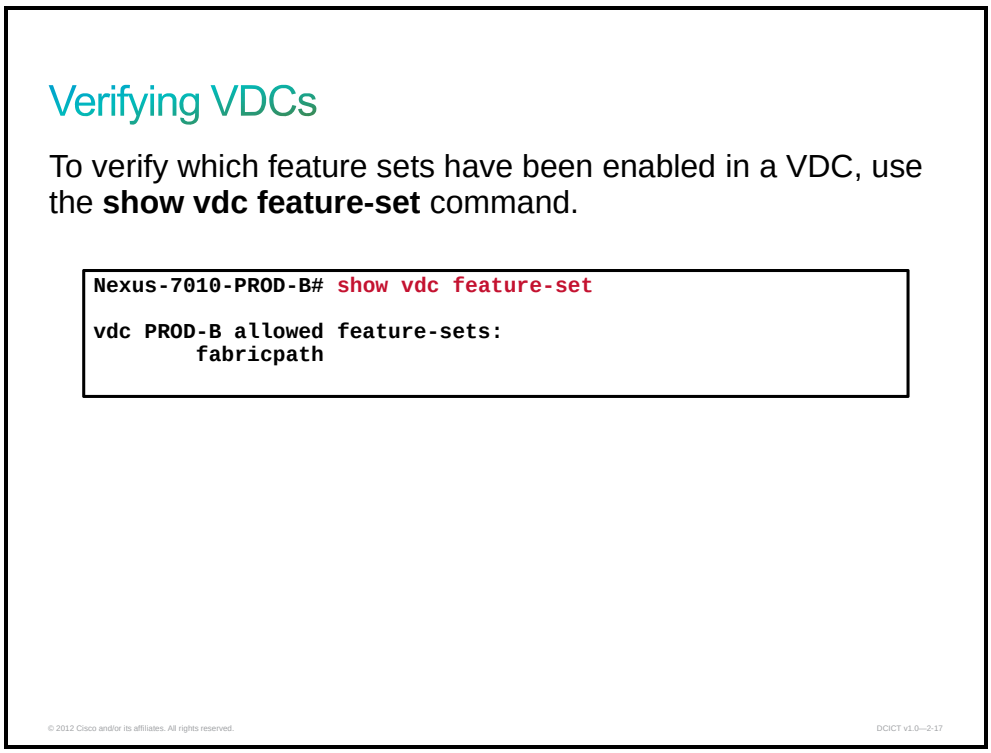
Note	When you configure different VDCs from the default VDC, verify that you are configuring the correct VDC. Accidentally making changes to the wrong VDC can affect switch operations.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying VDCs on the Cisco Nexus 7000 Series Switch

This topic describes how to verify VDCs on the Cisco Nexus 7000 Series switch.



The **show feature** command is used to verify which features have been specifically enabled on the Cisco Nexus switches.

https://t.me/learningnets

Verifying VDCs (Cont.)

- To display VDC summary information, use the **show vdc** command.

```
N7K-1# show vdc

vdc_id              vdc_name  state               mac
------              --------  -----               ---------
1                   N7K-1     active              00:18:ba:d8:3f:fd
2                   RED       active              00:18:ba:d8:3f:fe
3                   BLUE      active              00:18:ba:d8:3f:ff
```

- From a nondefault VDC, only information for that VDC is visible.

```
N7K-1-RED# show vdc

vdc_id              vdc_name  state               mac
------              --------  -----               ----------
2                   RED       active              00:18:ba:d8:3f:fe
```

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

The scope of the **show vdc** commands depends on the VDC in which the commands are executed. When these commands are executed in a nondefault VDC, the displayed information is restricted to that VDC only. If these commands are executed from the default VDC, then they display information on all VDCs, unless a specific VDC is entered as a command option.

Issuing the **show vdc** command from within the default VDC context lists all the active and current VDCs. The default VDC has visibility over all nondefault VDCs.

Issuing the **show vdc** command within a nondefault VDC context provides information about that VDC only. Nondefault VDCs have no visibility to one another or to the default VDC.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying VDCs (Cont.)

- To display detailed VDC information, use the **show vdc detail** command.

```
N7K-1# show vdc detail

vdc id: 1
vdc name: N7K-1
vdc state: active
vdc mac address: 00:18:ba:d8:3f:fd
vdc ha policy: RELOAD
vdc dual-sup ha policy: SWITCHOVER
vdc boot Order: 1
vdc create time: Sun Jan  2 04:02:58 2011
vdc reload count: 0
vdc restart count: 0

vdc id: 2
vdc name: RED
vdc state: active
vdc mac address: 00:18:ba:d8:3f:fe
vdc ha policy: RESTART
vdc dual-sup ha policy: SWITCHOVER
vdc boot Order: 1
vdc create time: Sat Jan 22 22:47:17 2011
vdc reload count: 0
vdc restart count: 0
```

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

The **show vdc detail** command provides more detailed information about the VDCs, including name, state, and high-availability policies.

Verifying VDC Interface Allocation

- This example shows how to verify VDC interface information from within the default VDC.

```
N7K-1# show vdc membership

vdc_id: 1 vdc_name: N7K-1 interfaces:
     Ethernet2/1  Ethernet2/2  Ethernet2/3  Ethernet2/4  Ethernet2/5
     Ethernet2/6  Ethernet2/7  Ethernet2/8  Ethernet2/9  Ethernet2/10
     Ethernet2/11 Ethernet2/12 Ethernet2/13 Ethernet2/14 Ethernet2/15
     Ethernet2/16 Ethernet2/17 Ethernet2/18 Ethernet2/19 Ethernet2/20
     Ethernet2/21 Ethernet2/22 Ethernet2/23 Ethernet2/24 Ethernet2/25
     Ethernet2/26 Ethernet2/27 Ethernet2/28 Ethernet2/29 Ethernet2/30
     Ethernet2/31 Ethernet2/32 Ethernet2/33 Ethernet2/34 Ethernet2/35
     Ethernet2/36 Ethernet2/37 Ethernet2/38 Ethernet2/39 Ethernet2/40
     Ethernet2/41 Ethernet2/42 Ethernet2/43 Ethernet2/44 Ethernet2/45
     Ethernet2/48

vdc_id: 2 vdc_name: RED interfaces:
     Ethernet2/47

vdc_id: 3 vdc_name: BLUE interfaces:
     Ethernet2/46
```

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

The **show vdc membership** command can be used to display the interfaces that are allocated to the VDCs.

https://t.me/learningnets

VDC Running Configuration

- From the default VDC, the running configuration for all VDCs on the device can be saved to the startup configuration by using one command.
- The running configurations for all VDCs can be viewed from the default VDC.

```
N7K-1# copy running-config startup-config vdc-all
```

```
N7K-1# show running-config vdc-all
!Running config for default vdc: N7K-1

version 5.0(3)
license grace-period
no hardware ip verify address identical
<…output omitted…>

!Running config for vdc: RED

switchto vdc RED

version 5.0(3)
feature telnet
<…further output omitted…>
```

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

You can save the running configuration for all VDCs by using the **copy running-config startup-config vdc-all** command. The **show running-config vdc-all** command displays the current configuration files for all VDCs.

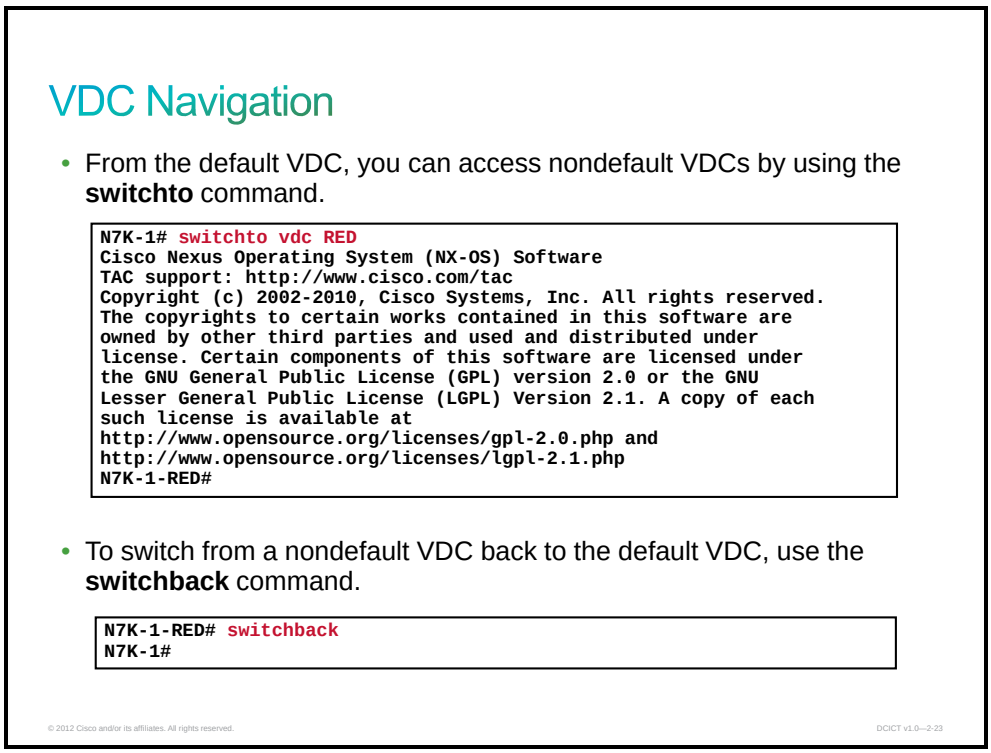
Both commands must be issued from the default VDC, which has visibility for all VDCs. You cannot view the configuration of other VDCs from a nondefault VDC.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Navigating Between VDCs on the Cisco Nexus 7000 Series Switch

This topic describes how to navigate between VDCs on the Cisco Nexus 7000 Series switch.



You can navigate between the default and nondefault VDCs by using the **switchto vdc** command. This action changes the context from the default to the specified nondefault VDC. This command cannot be used to navigate directly between nondefault VDCs. To navigate from one nondefault VDC to another, you must first issue the **switchback** command to return to the default VDC. You can follow that command with a **switchto** command to enter the configuration context for the desired nondefault VDC. This command is necessary to perform the initial setup of the VDCs. When user accounts and IP connectivity are configured properly, the VDC can be accessed over the network by using Secure Shell (SSH) or Telnet.

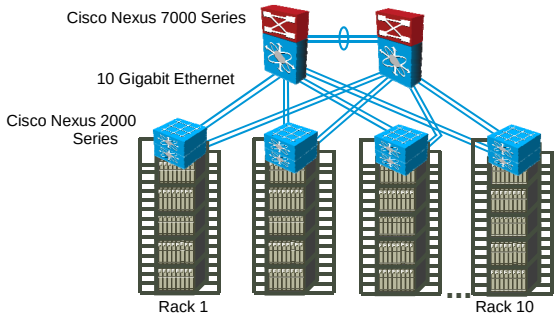
https://t.me/learningnets

Describing NIV on Cisco Nexus 7000 and 5000 Series Switches

This topic describes NIV on Cisco Nexus 7000 and 5000 Series Switches.

Using FEXs with Cisco Nexus Switches

- Cisco Nexus 2000 Series fabric extenders serve as remote line cards of a Cisco Nexus 5000 or 7000 Series switch.
  - FEXs are managed and configured from the Cisco Nexus switch.
- Together, the Cisco Nexus switches and Cisco Nexus 2000 Series fabric extenders combine benefits of ToR cabling with EoR management.



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

Cisco Nexus 2000 Series Fabric Extenders can be deployed with Cisco Nexus 5000 or Cisco Nexus 7000 Series Switches. This deployment can create a data center network that combines the advantages of a top-of-rack (ToR) design with the advantages of an end-of-row (EoR) design.

Dual redundant Cisco Nexus 2000 Series Fabric Extenders are placed at the top of each rack. The uplink ports on the fabric extenders (FEXs) are connected to a Cisco Nexus 5000 or Cisco Nexus 7000 Series switch that is installed in the EoR position. From a cabling standpoint, this design is a ToR design. The cabling between the servers and the Cisco Nexus 2000 Series fabric extender is contained within the rack. Only a limited number of cables need to be run between the racks to support the 10 Gigabit Ethernet connections between the FEXs and the Cisco Nexus switches in the EoR position.

From a network-deployment standpoint, however, this design is an EoR design. The FEXs act as remote line cards for the Cisco Nexus switches, so the ports on the Cisco Nexus 2000 Series Switches act as ports on the associated switch. In the logical network topology, the FEXs disappear and all servers appear as directly connected to the Cisco Nexus switch. From a network-operations perspective, this design has the simplicity that is typically associated with EoR designs. All the configuration tasks for this type of data center design are performed on the EoR switches. No configuration or software maintenance tasks are associated with the FEXs.

https://t.me/learningnets

FEX Operation

- The FEX is not a switch.
  - The FEX is a line card (module) of the parent switch.
- The parent switch:
  - Performs all configuration and management functions
  - Makes all forwarding, security, and QoS decisions
- Cisco Nexus 7000 Series Switches support as many as 32 FEXs.
- Cisco Nexus 5000 Series Switches support as many as 12 FEXs.
- Cisco Nexus 5500 Series Switches support as many as 24 FEXs.

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

All Cisco Nexus 2000 Series platforms are controlled by a parent Cisco Nexus switch. The FEX is not a switch. When connected to the parent switch, the FEX operates as a line card (module) of the parent switch. The parent switch controls all configuration, management, and software updates. In the data plane, all forwarding, security, and quality of service (QoS) decisions are made at the parent switch. Fibre Channel over Ethernet (FCoE) can be supported with the Cisco Nexus 2232 Fabric Extender.

Cisco Nexus 2000 Series Fabric Extenders on Cisco Nexus 7000 Series Switches

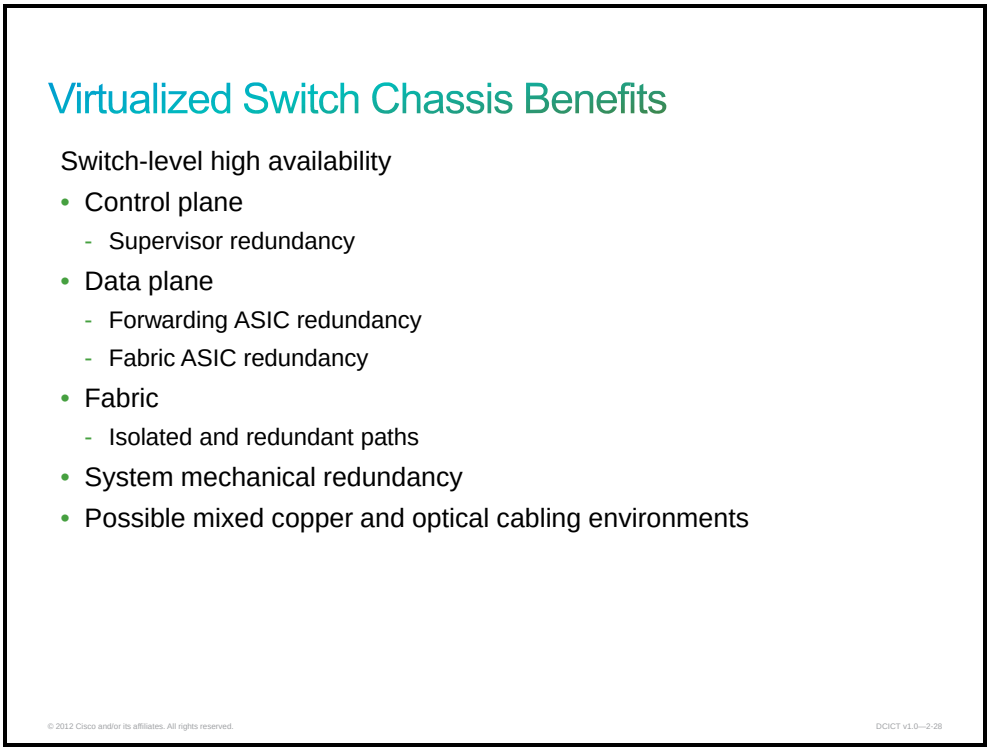


	Cisco Nexus 2224TP GE	Cisco Nexus 2248TP-E	Cisco Nexus 2232PP
FEX host interfaces	24	48	32
FEX host interfaces type	100 BASE-T/1000 BASE-T ports	100 BASE-T/1000 BASE-T ports	1/10 Gigabit Ethernet ports SFP or SFP+
FEX fabric interfaces	2	4	8
Fabric speed	20 Gb/s	40 Gb/s	80 Gb/s
Oversubscription	1.2:1	1.2:1	4:1
Performance	65 mpps	131 mpps	595 mpps
Minimum software	Cisco NX-OS Release 5.2	Cisco NX-OS Release 5.1	Cisco NX-OS Release 5.2

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

The table lists FEXs that are supported on the Cisco Nexus 7000 Series switch.

https://t.me/learningnets



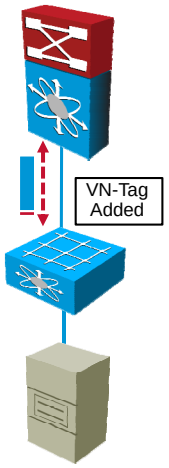
Switch-level high availability addresses these features:

- Control plane
  - Supervisor redundancy
- Data plane
  - Forwarding ASIC redundancy
  - Fabric ASIC redundancy
- Fabric
  - Isolated and redundant paths
- System mechanical redundancy
  - Power supply
  - Fan

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

Cisco Nexus Switch 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) down to the Cisco Nexus 2000 Series (host interface).
  - VNTag is added to the packet between the fabric extender and parent switch.
  - 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.

The Cisco Nexus 2000 Series fabric extender operates as a remote line card to the Cisco Nexus switch. All control and management functions are performed by the parent switch. Forwarding is also performed by the parent switch.

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

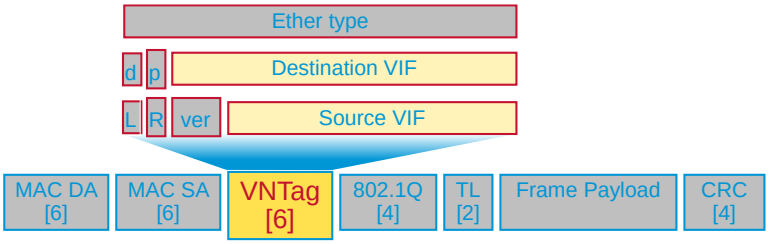
Packets sent to and from the Cisco Nexus 2000 Series fabric extender have a virtual network tag (VNTag) added to them so that the upstream switch knows how to treat the packets and which policies to apply. This tag does not affect the host, which is unaware of any tagging.

The link or links between the Cisco Nexus 2000 Series fabric extender and the upstream switch multiplex the traffic from multiple devices that connect to the FEX.

https://t.me/learningnets

VNTag Overview

- VNTag technology is being standardized under IEEE 802.1Qbh.
  - **d:** Direction bit (0 is host-to-network forwarding; 1 is 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 source Cisco Nexus 2000 Series)
  - **VIF:** Virtual interface



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

VNTag technology is being standardized under IEEE 802.1Qbh.

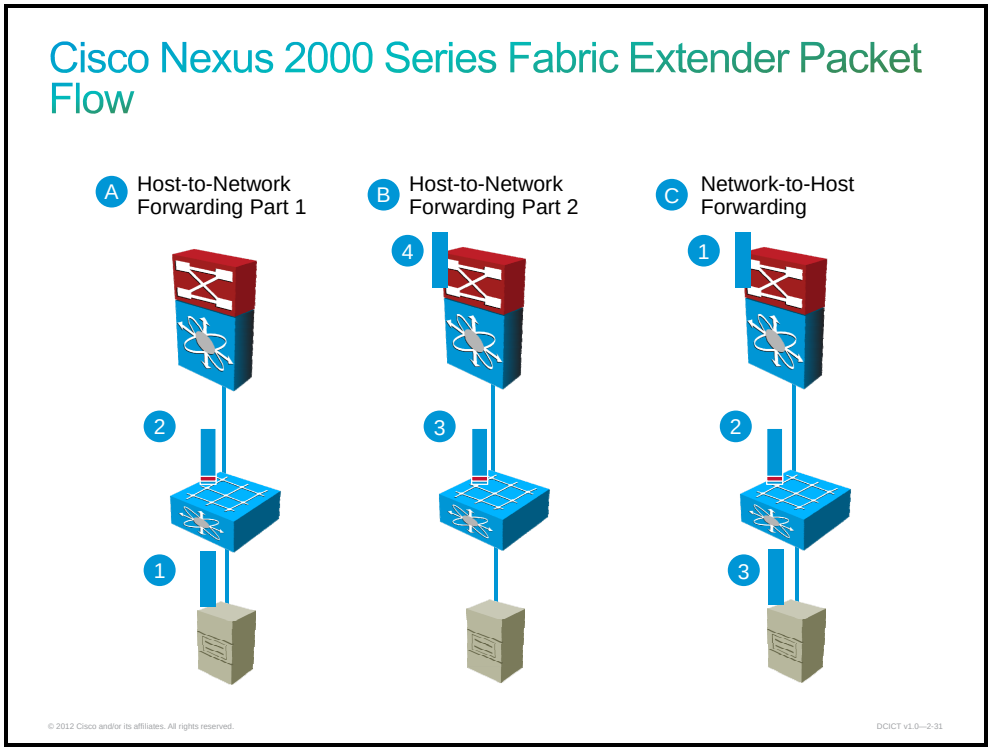
Because a fabric extender does not support local switching, forwarding is on the parent switch. Forwarding from and to a FEX is based on the VNTag. VNTag is an NIV technology. A VNTag allows the FEX to act as a data path of the Cisco Nexus 7000 Series switch, for all policy and forwarding. The VNTag is added to the packet between the fabric extender and Cisco Nexus switch. The VNTag is stripped before the packet is sent to hosts.

The figure describes the VNTag fields in an Ethernet frame.

Note The VNTag is completely transparent to the end host.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



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

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

1. The frame arrives from the host.
2. The Cisco Nexus 2000 Series switch adds a VNTag, and the packet is forwarded over a fabric link, using a specific VNTag. The Cisco Nexus 2000 Series switch adds a unique VNTag for each Cisco Nexus 2000 Series host interface. These are the VNTag 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 switch extracts the VNTag, which identifies the logical interface that corresponds to the physical host interface on the Cisco Nexus 2000 Series. The Cisco Nexus switch applies an ingress policy that is based on the physical Cisco Nexus Series switch port and logical interface:
   - Access control and forwarding are based on frame fields and virtual (logical) interface policy.
   - Physical link-level properties are based on the Cisco Nexus Series switch port.
4. The Cisco Nexus switch strips the VNTag and sends the packet to the network.

https://t.me/learningnets

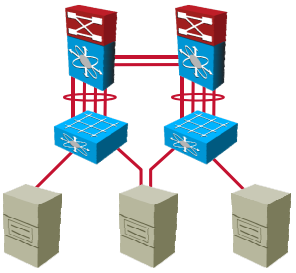
When the network sends a packet to the host (diagram C), these events occur:

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

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

Cisco Nexus 7000 Series Fabric Extender Fabric Ports Connectivity

- You must connect the FEX to its parent Cisco Nexus 7000 Series device by using the 32-port 10-Gb/s M1 module (N7K-M132XP-12) and the 32-port 10-Gb/s M1-XL module (N7K-M132XP-12L).
- N7K-M132XP-12 and N7K-M132XP-12L ports need to be in shared rate mode for FEX connectivity.
- Multiple FEX uplink ports can use the same 10-Gb/s ports group in shared rate mode on Cisco Nexus 7000 Series.
- The failure of FEX uplinks or the Cisco Nexus 7000 Series causes failure of server ports on the FEX.

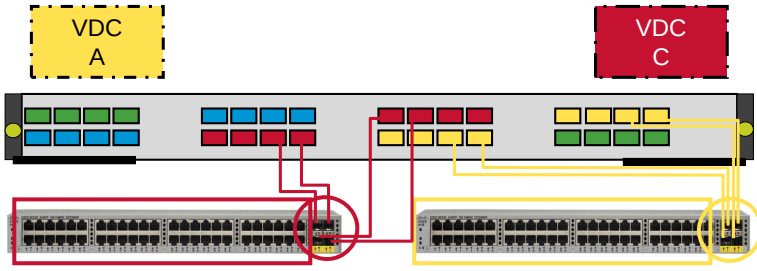


© 2012 Cisco and/or its affiliates. All rights reserved.

When a FEX is connected to a Cisco Nexus 7000 Series switch, it should be connected by using a 32-port 10-Gb/s M1 module.

Cisco Nexus 7000 Series Fabric Extender and VDC

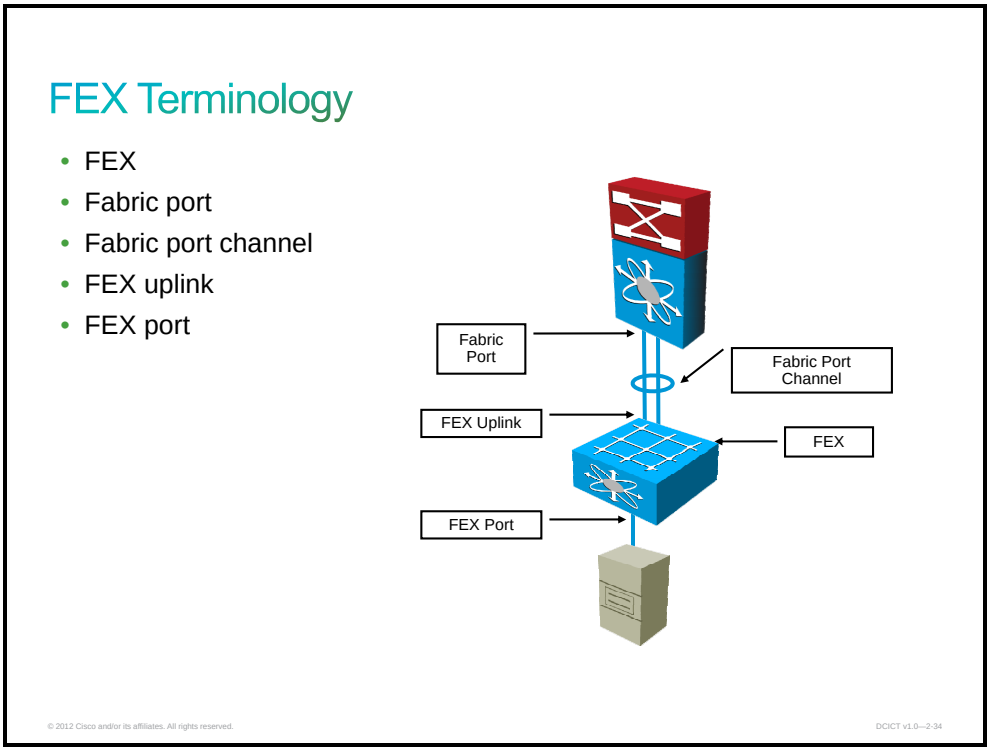
- All FEX fabric links must be in the same VDC.
- All FEX host ports belong to the same VDC.



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

When you use VDCs on a Cisco Nexus 7000 Series switch, all FEX fabric links should be in the same VDC. In addition, all FEX host ports belong to the same VDC to which the FEX fabric links are connected.

https://t.me/learningnets



The figure describes terminology that is used with FEXs:

- **FEX:** The Cisco Nexus 2000 Series fabric extender is a FEX.
- **Fabric port:** The Cisco Nexus switch side of the link that is connected to the FEX is called the fabric port.
- **Fabric port channel:** The fabric port channel is the port channel between the Cisco Nexus switch and FEX.
- **FEX uplink:** The network-facing port on the FEX that is, the FEX side of the link that is connected to the Cisco Nexus switch) is called the FEX uplink. The FEX uplink is also referenced as a network interface.
- **FEX port:** The server-facing port on the FEX is called the FEX port, also referred to as server port or host port. The FEX port is also referenced as a host interface.

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

FEX Configuration on Cisco Nexus 7000 Series Switches

FEX configuration on Cisco Nexus 7000 Series switches differs slightly from configuration on Cisco Nexus 5000 Series Switches:

- The FEX feature set needs to be installed in the default VDC before the feature set can be used in nondefault VDCs.
- Use of the FEX feature set can be allowed or disallowed per VDC. To allow the default:

```
N7K-1(config)# install feature-set fex
```

- Cisco Nexus 7000 Series Switches support only dynamic pinning, so the FEX fabric interfaces must be members of a port channel.

```
N7K-1(config-vdc)# no allow feature-set fex
```

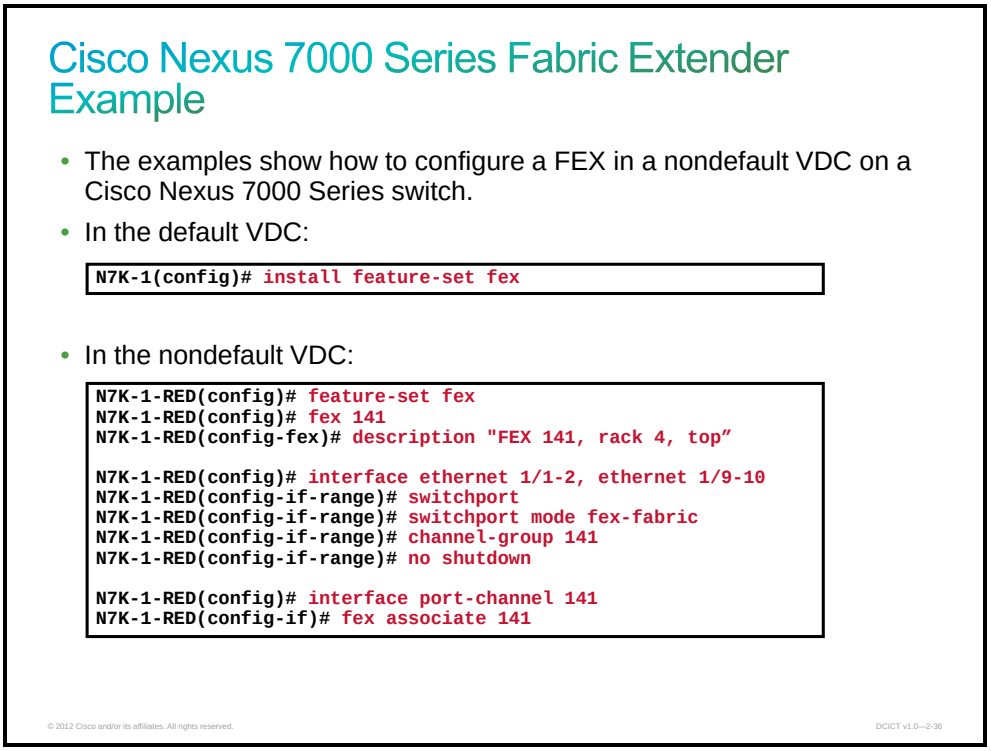
© 2012 Cisco and/or its affiliates. All rights reserved.

Configuration of a FEX on a Cisco Nexus 7000 Series switch is slightly different than the configuration on a Cisco Nexus 5000 Series switch. This difference is partially caused by the VDC-based architecture of Cisco Nexus 7000 Series Switches. Before any FEX can be configured in any VDC, the services that the FEX feature requires must be installed in the default VDC. To enable the use of the FEX feature set, use the **install feature-set fex** command in the default VDC. After the FEX feature set has been installed in the default VDC, you can enable the feature set in any VDC by using the **feature-set fex** command.

You can restrict the use of the FEX feature set to specific VDCs only. By default, all VDCs can enable the FEX feature set after it has been installed in the default VDC. If you want to disallow the use of FEXs in a specific VDC, you can use the **no allow feature-set fex** command in VDC configuration mode for that VDC.

Another difference is that Cisco Nexus 5000 Series Switches do not support dynamic pinning. Dynamic pinning makes it unnecessary to specify the maximum number of pinning interfaces by using the **pinning max-links** command.

https://t.me/learningnets



The example in the figure shows how to configure a Cisco Nexus 2000 Series fabric extender for use with a Cisco Nexus 7000 Series switch. The example shows the configuration to enable the FEX feature set in the default VDC, followed by the configuration in the nondefault VDC to which the FEX is connected.

Note	The use of a port channel to associate the FEX is mandatory on Cisco Nexus 7000 Series Switches.

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

Determine the Status of the FEX

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

```
N5K# 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) –
Eth1/13       1      eth  access up      none                        10G(D) --
Eth1/14       1      eth  access up      none                        10G(D) --
Eth1/15       1      eth  access up      none                        10G(D) --
Eth1/16       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.

Note	The output in the figure is from a Cisco Nexus 5000 Series switch.

https://t.me/learningnets

Verifying FEX

```
N7K-1-RED(config-if)# show fex
  FEX         FEX           FEX                              FEX
Number    Description      State        Model              Serial
------------------------------------------------------------------
141        FEX0141          Online     N2K-C2248TP-1GE  JAF1432CKHC
```

```
N7K-1-RED(config-if)# show inventory fex 141
NAME: "FEX 110 CHASSIS",  DESCR: "N2K-C2248TP-1GE  CHASSIS"
PID: N2K-C2248TP-1GE     ,  VID: V03 ,  SN: JAF1432CKHC

NAME: "FEX 141 Module 1",  DESCR: "Fabric Extender Module: 48x1GE, 4x10GE Supervisor"
PID: N2K-C2248TP-1GE     ,  VID: V03 ,  SN: SSI141308T5

NAME: "FEX 141 Fan 1",  DESCR: "Fabric Extender Fan module"
PID: N2K-C2248-FAN       ,  VID: N/A ,  SN: N/A

NAME: "FEX 141 Power Supply 1",  DESCR: "Fabric Extender AC power supply"
PID: N2200-PAC-400W      ,  VID: V02 ,  SN: LIT14291RYF

NAME: "FEX 141 Power Supply 2",  DESCR: "Fabric Extender AC power supply"
PID: N2200-PAC-400W      ,  VID: V02 ,  SN: LIT14291S0B

110        FEX0141          Online     N2K-C2248TP-1GE  JAF1432CKHC
```

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

The figure describes commands that are used to verify the FEX:

- The **show fex** *FEX-number* command displays module information about a FEX.
- The **show inventory fex** *FEX-number* displays inventory information for a FEX.

Note The output in the figure is from a Cisco Nexus 7000 Series switch.

Verifying FEX (Cont.)

```
N7K-1-RED(config-if)# sh fex detail
FEX: 141 Description: FEX0141   state: Online
  FEX version: 5.1(1) [Switch version: 5.1(1)]
  FEX Interim version: 5.1(1.3)
  Switch Interim version: 5.1(1)
  Extender Model: N2K-C2248TP-1GE,  Extender Serial: JAF1441ANJE
  Part No: 73-12748-05
  Card Id: 99, Mac Addr: 58:8d:09:16:7a:42, Num Macs: 64
  Module Sw Gen: 12594  [Switch Sw Gen: 21]
 pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Po141
  Fabric interface state:
    Po141 - Interface Up. State: Active
    Eth1/1  - Interface Up. State: Active
    Eth1/2  - Interface Up. State: Active
    Eth1/9  - Interface Up. State: Active
    Eth1/10 - Interface Up. State: Active
  Fex Port        State  Fabric Port  Primary Fabric
       Eth141/1/1  Down       Po141       Po141
       Eth141/1/2  Down       Po141       Po141
```

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

The **show fex detail** command displays detailed information about the FEX.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

FEX Host Interfaces

The FEX host interfaces appear in the running configuration and as interfaces of the Cisco Nexus switch.

```
N7K-1-RED# show running-config | begin "interface Ethernet141"

interface Ethernet141/1/1

interface Ethernet141/1/2

interface Ethernet141/1/3
-- output omitted --
```

```
N7K-1-RED# show interface ethernet 141/1/1

Ethernet141/1/1 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 switch. The interfaces are listed as FEX number, module, and slot number designation, in that order. 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 are connected to the Cisco Nexus switch through the fabric interfaces and physically resident on the Cisco Nexus 2000 Series switch.

Note The output in the figure is from a Cisco Nexus 7000 Series switch.

https://t.me/learningnets

Verifying Port Pinning

Verify the port pinning distribution by using the **show interface fex-intf** command.

```
N5K# show interface ethernet 1/9 fex-intf
Fabric           FEX
Interface        Interfaces
---------------------------------------------------
 Eth1/9          Eth141/1/1    Eth141/1/2    Eth141/1/3    Eth141/1/4
                 Eth141/1/5    Eth141/1/6    Eth141/1/7    Eth141/1/8
                 Eth141/1/9    Eth141/1/10   Eth141/1/11   Eth141/1/12
                 Eth141/1/13   Eth141/1/14   Eth141/1/15   Eth141/1/16
                 Eth141/1/17   Eth141/1/18   Eth141/1/19   Eth141/1/20
                 Eth141/1/21   Eth141/1/22   Eth141/1/23   Eth141/1/24
                 Eth141/1/25   Eth141/1/26   Eth141/1/27   Eth141/1/28
                 Eth141/1/29   Eth141/1/30   Eth141/1/31   Eth141/1/32
                 Eth141/1/33   Eth141/1/34   Eth141/1/35   Eth141/1/36
                 Eth141/1/37   Eth141/1/38   Eth141/1/39   Eth141/1/40
                 Eth141/1/41   Eth141/1/42   Eth141/1/43   Eth141/1/44
                 Eth141/1/45   Eth141/1/46   Eth141/1/47   Eth141/1/48
N5K# show interface ethernet 1/10 fex-intf
Fabric           FEX
Interface        Interfaces
---------------------------------------------------
 Eth1/10
```

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

The port pinning can be verified by issuing the **show interface fex-intf** command. The output of this command displays each uplink and the set of FEX interfaces that are associated to that uplink.

The output in the figure shows that all FEX interfaces are associated with uplink Ethernet 1 /9 and that no interfaces use interface Ethernet 1 /10.

Note The output in the figure is from a Cisco Nexus 5000 Series switch.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying the FEX Installation

- All FEXs that are connected to the system should now be discovered and have an associated virtual slot number.

```
N7K-1-RED# show interface fex-fabric
     Fabric      Fabric       Fex                FEX
Fex  Port      Port State    Uplink    Model          Serial
--------------------------------------------------------------------
141    Eth1/9        Active     1    N2K-C2248TP-1GE  JAF1419ECAC
141   Eth1/10        Active     2    N2K-C2248TP-1GE  JAF1419ECAC
---   Eth1/11        Discovered 3    N2K-C2248TP-1GE  JAF1420AHHE
---   Eth1/12        Discovered 4    N2K-C2248TP-1GE  JAF1420AHHE
```

- The server-facing ports on the FEX can be configured from the Cisco Nexus switch.

```
N7K-1-RED(config)# interface ethernet 141/1/13
N7K-1-RED(config-if)# description ServerX
N7K-1-RED(config-if)# switchport access vlan 10
N7K-1-RED(config-if)# shutdown
```

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

When the FEX is fully recognized by the Cisco Nexus switch and configured, use the **show interface fex-fabric** command to display the discovered FEXs. You can also see the associated FEX number, which represents the virtual slot number of the FEX in the Cisco Nexus switch virtualized chassis in the active state.

The individual server-facing FEX ports can now be configured from the Cisco Nexus switch.

Note	The output in the figure is from a Cisco Nexus 7000 Series switch.

https://t.me/learningnets

Summary

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

Summary

- VDCs on a Cisco Nexus 7000 Series switch provide organizations the ability to consolidate on fewer physical switches while retaining the separation of administrative domains.
- Various VDC **show** commands can be used to verify which VDCs have been created, the operational status of the VDCs, and which interfaces are members of which VDC.
- The switch administrator can use the **switchto** command to switch to a nondefault VDC from the default VDC. VDC administrators cannot use the **switchto** command to switch to the default VDC.
- Cisco Nexus 7000 and 5000 Series Switches support NIV through the addition of remote line modules (Cisco Nexus 2000 Series Fabric Extenders) that are connected to the switch.

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

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Virtualizing Storage

Overview

The purpose of this lesson is to describe how storage is virtualized for high availability and configuration flexibility.

Objectives

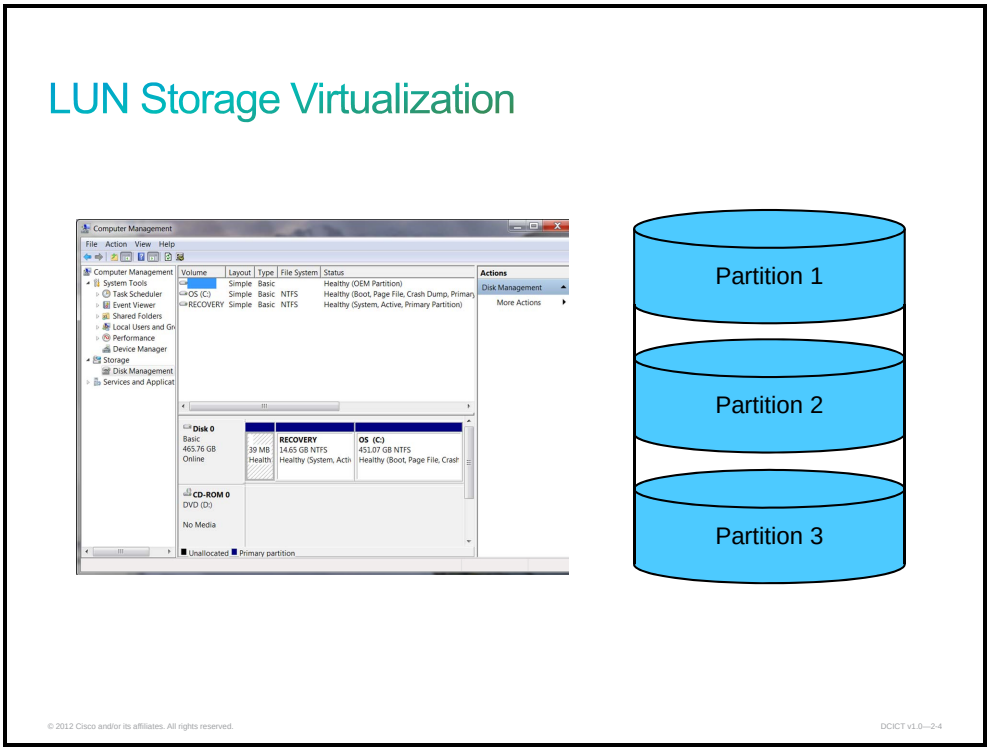
Upon completing this lesson, you will be able to describe storage virtualization. You will be able to meet these objectives:

- Describe LUN storage virtualization
- Describe storage system virtualization

https://t.me/learningnets

LUN Storage Virtualization

This topic describes logical unit number (LUN) storage virtualization.



Virtualization of storage helps to provide location independence by abstracting the physical location of the data. The user is presented with a logical space for data storage by the virtualization system, which manages the process of mapping that logical space to the physical location.

You can have multiple layers of virtualization or mapping, by using the output of one layer as the input for a higher layer of virtualization. Virtualization maps the relationships between the back-end resources and front-end resources. The back-end resource refers to a LUN that is not presented to the computer or host system for direct use. The front-end LUN or volume is presented to the computer or host system for use.

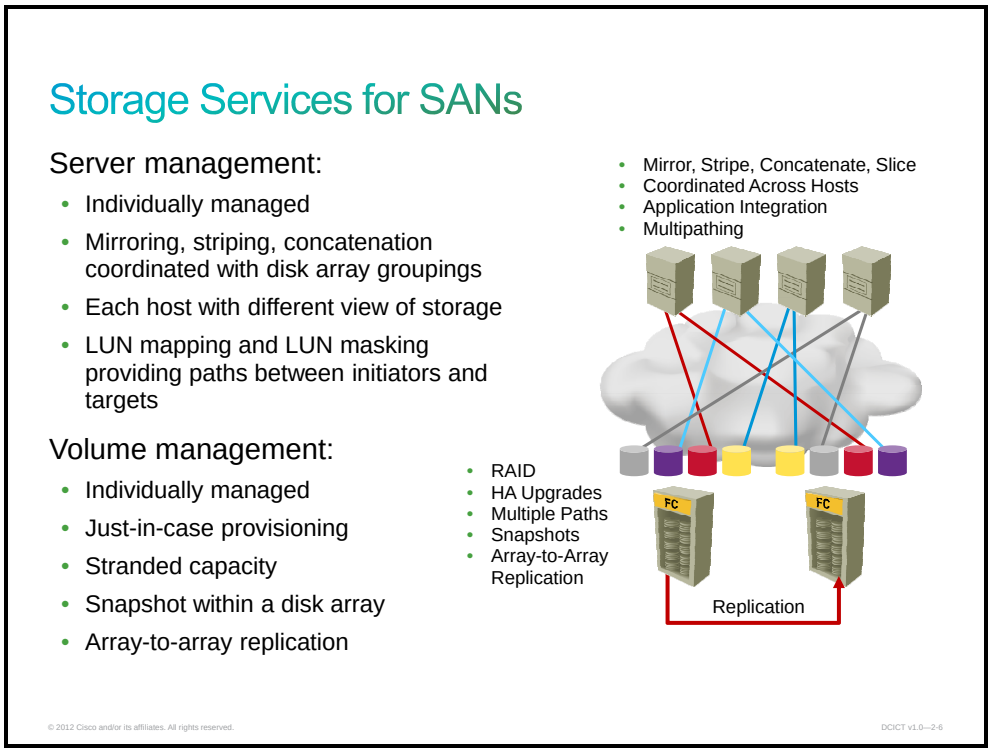
How the mapping is performed depends on the implementation. Typically, one physical disk is broken down into smaller subsets in multiple megabytes or gigabytes of disk space.

In a block-based storage environment, one block of information is addressed by using a LUN and an offset within that LUN, known as a Logical Block Address (LBA).

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

Storage System Virtualization

This topic describes how storage controllers can virtualize the back-end storage from multiple vendors for a single point of management.



In the SAN, a mixture of ad hoc storage services is often provided.

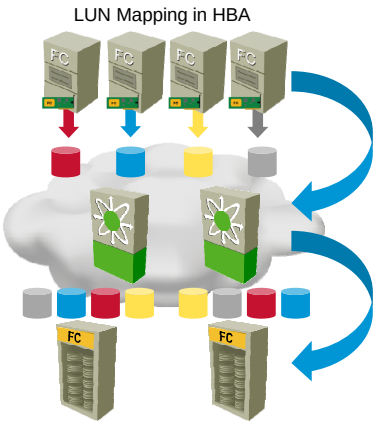
Managing individual volumes and multipathing at the host level adds to the complexity of SAN administration. Each server requires its own investment in management and attention. SAN administrators typically overprovision storage in this scenario, as a strategy to reduce the amount of time that is spent on resource management. This overprovisioning often results in underutilized and wasted storage.

In this scenario, redundancy and replication tasks are often achieved at the array level. This scenario can be created in a same box-to-same box configuration or by using a third-party software utility to replicate across heterogeneous storage. This scenario adds an additional layer of complexity to the information life cycle and overall SAN management. Low-value data can end up residing on expensive storage.

https://t.me/learningnets

LUN Mapping and LUN Masking

- LUN masking:
  - LUN masking is used in the storage array to mask or hide LUNs from servers that are denied access.
  - Storage array makes specific LUNs available to server ports identified by their pWWN.
- LUN mapping:
  - Server maps some or all visible LUNs to volumes.
- Heterogeneous SANs with storage arrays and JBODs from different vendors:
  - Difficult to configure
  - Costly to manage
  - Difficult to replicate and migrate data



RAID Configuration and
LUN Masking in Storage Arrays

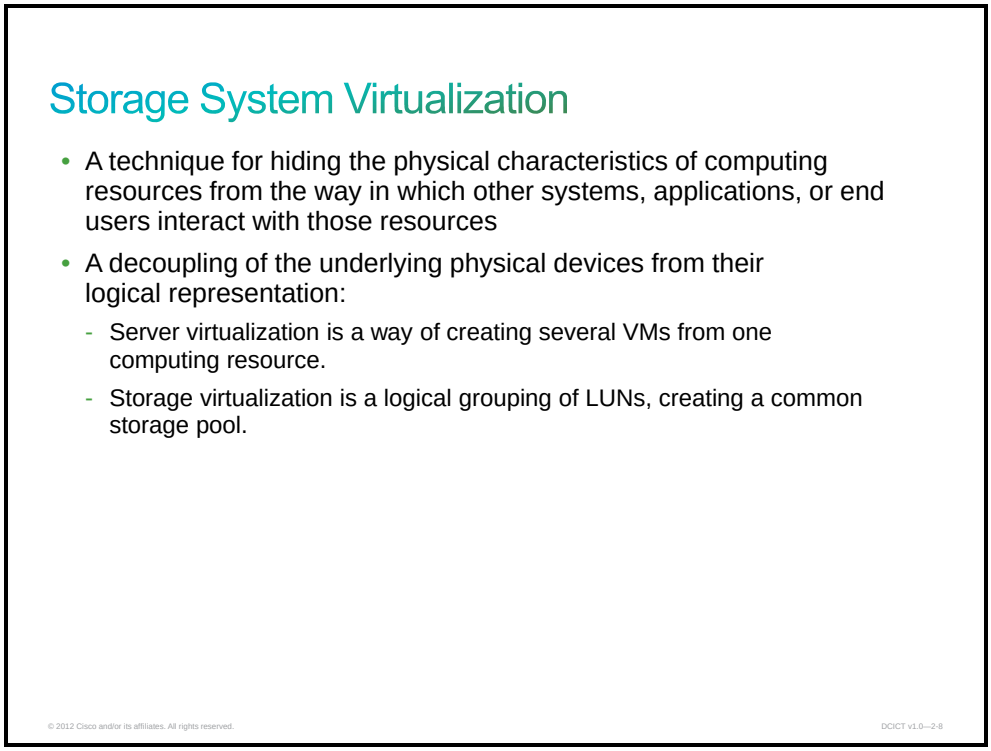
© 2012 Cisco and/or its affiliates. All rights reserved.

In most SAN environments, each individual LUN must be discovered only by one server host bus adapter (HBA). Otherwise, the same volume will be accessed by more than one file system, leading to potential loss of data or security. There are basically three ways to prevent this multiple access:

- **LUN masking:** LUN masking, a feature of enterprise storage arrays, provides basic LUN-level security by allowing LUNs to be seen only by selected servers that are identified by their port world wide name (pWWN). Each storage array vendor has its own management and proprietary techniques for LUN masking in the array. In a heterogeneous environment with arrays from different vendors, LUN management becomes more difficult.
- **LUN mapping:** LUN mapping, a feature of Fibre Channel HBAs, allows the administrator to selectively map some LUNs that have been discovered by the HBA. LUN mapping must be configured on every HBA. In a large SAN, this mapping is a large management task. Most administrators configure the HBA to automatically map all LUNs that the HBA discovers. They then perform LUN management in the array (LUN masking) or in the network (LUN zoning).
- **LUN zoning:** LUN zoning, a proprietary technique that Cisco MDS switches offer, allows LUNs to be selectively zoned to their appropriate host port. LUN zoning can be used instead of, or in combination with, LUN masking in heterogeneous environments or where Just a Bunch of Disks (JBODs) are installed.

Note	JBODs do not have a management function or controller and so do not support LUN masking.

https://t.me/learningnets

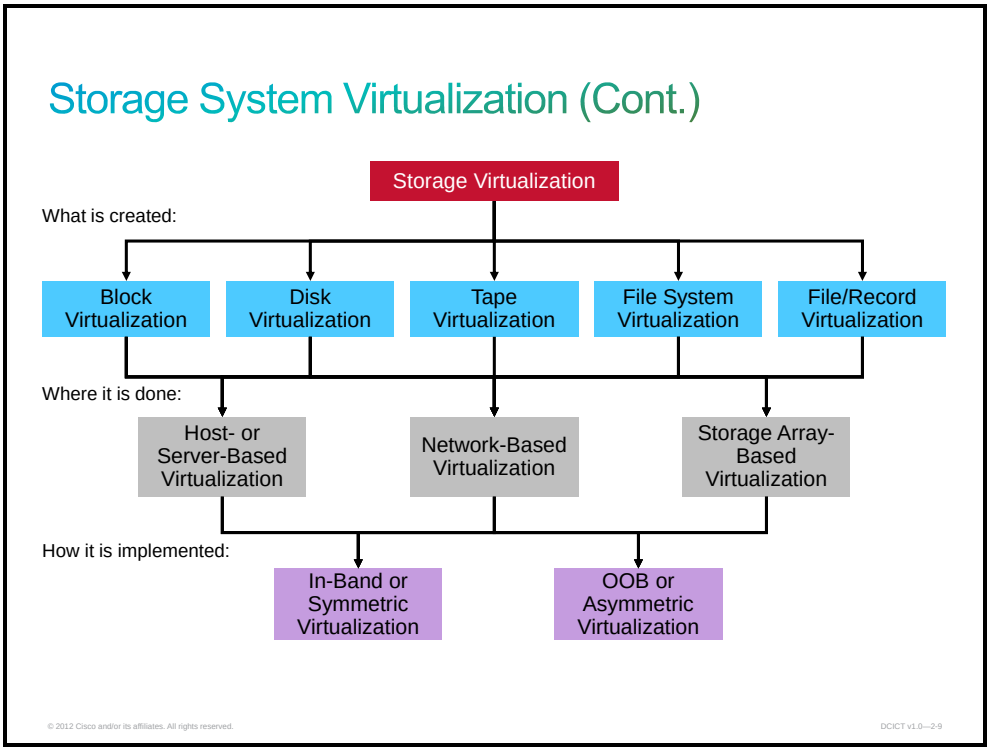


Virtualization provides a process for presenting a logical grouping or subset of computing resources. Virtualization provides a means of hiding the physical characteristics of computing resources. In essence, virtualization decouples the underlying physical device from the logical representation.

In a heterogeneous environment, LUN management can become costly and time consuming. Storage virtualization is often used instead, to create a common pool of all storage resources and perform LUN management within the network:

- **Server virtualization:** A method of creating several virtual machines (VMs) from one computing resource
- **Storage virtualization:** A logical grouping of LUNs, creating a common storage pool

https://t.me/learningnets



The figure shows different types of storage system virtualization:

- **Block virtualization:** RAID groups
- **Disk virtualization:** Disk drives in a storage array
- **Tape virtualization:** Tape drives in a virtual tape library
- **File system virtualization:** Server virtualization
- **File and record virtualization:** Volume virtualization

Virtualization can be performed in these places:

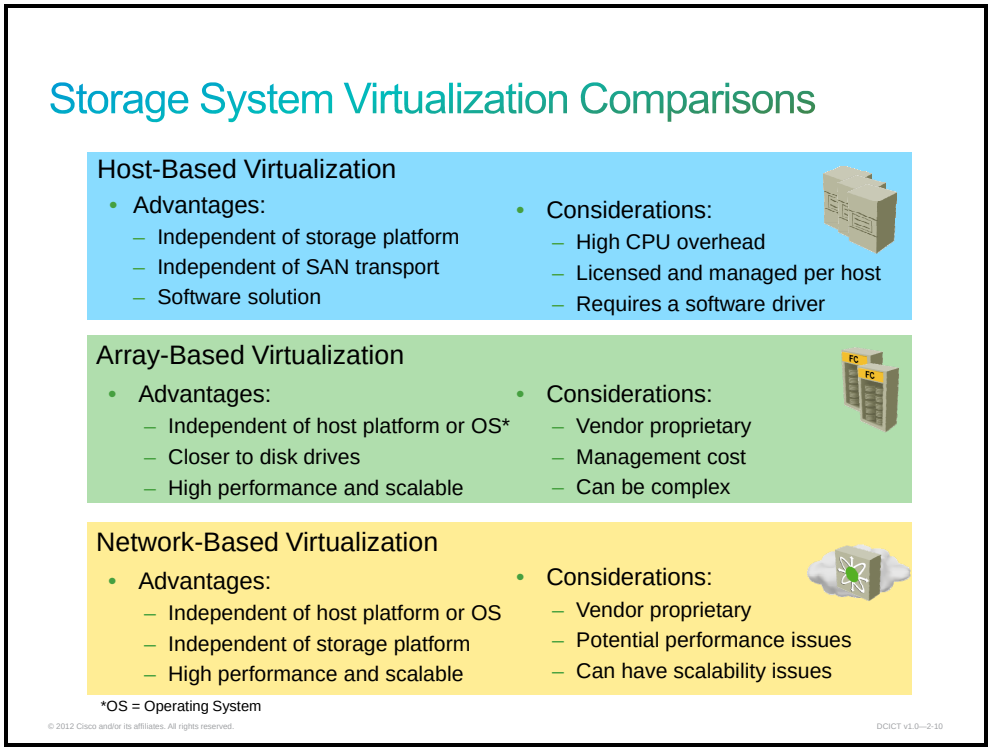
- Host or server
- Network in a SAN appliance or in the Fibre Channel switch
- Storage device or storage array

The Storage Network Industry Association (SNIA) defines two main definitions for storage system virtualization:

- **In-band (symmetric) virtualization:** Data and control take the same path.
- **Out-of-band (OOB, asymmetric) virtualization**: Data and control take different paths.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



Host-Based Virtualization

Host-based virtualization has certain advantages and disadvantages:

- **Advantages:** This software solution is independent of the storage platform or vendor and of the underlying SAN technology, Fibre Channel, Internet Small Computer Systems Interface (iSCSI), and so on. Veritas offers a software-based solution.
- **Considerations:** Host-based virtualization is licensed on a per host basis and can be expensive. The virtualization software is complex and must trap and redirect every frame before it reaches the storage target. Each host is managed individually and requires a software driver to be installed. Any hosts that connect to the SAN without the installed driver can still access the virtual LUNs directly, which can lead to data corruption.

Array-Based Virtualization

- **Advantages:** The array is independent of the host platform or operating system. The virtualization function is in the array controller and is much closer to the physical disk drives. This closeness makes the solution more responsive, creates less SAN data traffic, and provides higher performance and more scalability within the array.
- **Considerations:** This solution is vendor-proprietary, complex, and difficult to manage in a heterogeneous SAN. Reconfiguration of the array, redistribution of LUNs, and RAID configuration often require specialized knowledge. This reconfiguration can be expensive if performed by an engineer onsite.

Network-Based Virtualization

- **Advantages:** Network-based virtualization is provided through a virtualization appliance or within the Fibre Channel switch. This type of virtualization is completely independent of the storage platform, host platform, or operating system. By connecting to the SAN, network-based virtualization has access to all connected hosts and storage devices and can provide high performance and scalability.

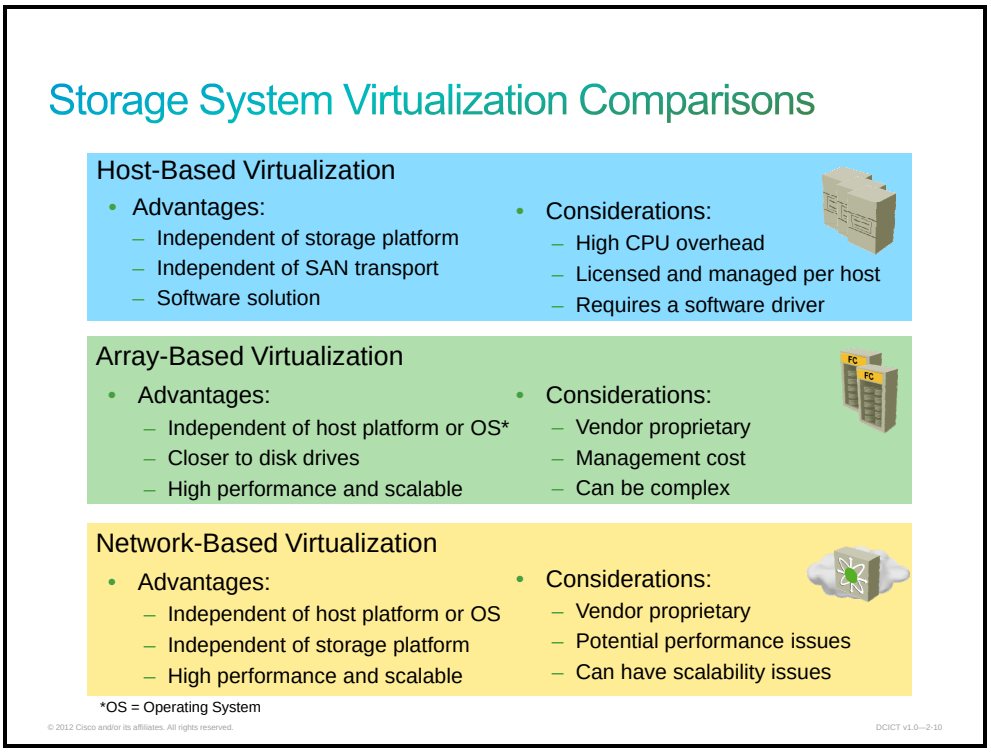
https://t.me/learningnets

- **Considerations:** Network-based virtualization is proprietary and often seen as a bottleneck, particularly if it is in-band and in the data path. Network-based virtualization can have scalability issues and can be a single point of failure if the devices are not clustered. Switch-based virtualization resolves many of these objections by having access to the high-performance backplane and can scale across the SAN. Solutions that are based on the Fabric Application Interface Standard (FAIS) conform to a standard and are vendor independent.

Storage Virtualization

FAIS is an ANSI T11 standards-based effort to create a common application programming interface (API) for fabric applications to run on an underlying hardware platform. FAIS supports storage functions that perform classic enterprise-level storage transformation processes—for example, virtualization and RAID.

Note FAIS is pronounced *face*.



With symmetric virtualization, all I/Os and metadata are routed via a central virtualization storage manager. Data and control messages use the same path. This design is architecturally simpler but can create a bottleneck.

The virtualization engine does not need to reside in a completely separate device. The engine can be embedded in the network as a specialized switch, or it can run on a server. To provide alternate data paths and redundancy, two or more virtual storage management devices are usually used. This redundancy can lead to issues of consistency between the metadata databases that are used to perform the virtualization.

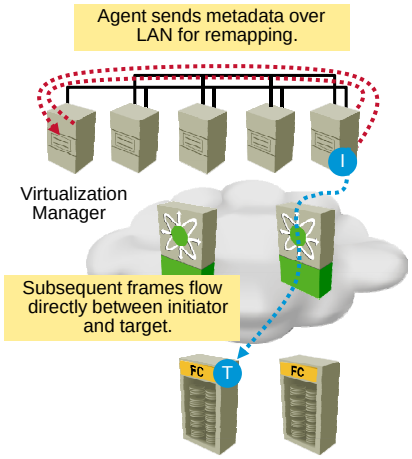
All data I/Os are forced through the virtualization appliance, restricting the SAN topologies that can be used and possibly causing a bottleneck. The bottleneck is often addressed by using caching and other techniques to maximize the performance of the engine. However, this technique again increases complexity and leads to consistency problems between engines.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Asymmetric Virtualization (OOB)

- Each server runs a virtualization agent:
  - Intercepts block I/O requests
  - Sends the meta-data (CDB and LUN) to a virtualization manager on the LAN
- The virtualization manager remaps the CDB and LUN and returns it to the server.
- The server sends the modified control frame to the storage target port.
- All subsequent data and response frames flow directly between initiator and target.
- Advantages:
  - Low latency
- Considerations:
  - Requires agent in host to intercept control frame.
  - Remapping of CDB and LUN adds latency to first frame in exchange.
  - Virtualization manager could be single point of failure.



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

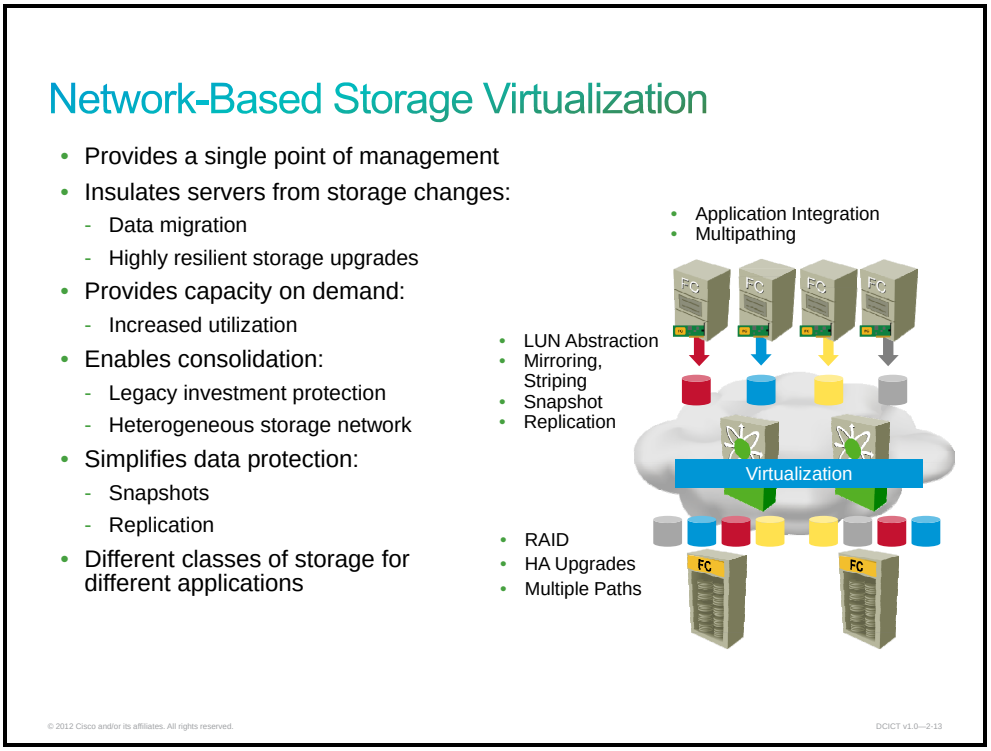
In asymmetric virtualization, the I/O is split into three parts:

- The server intercepts the block I/O requests.
- The server queries the metadata manager to determine the physical location of the data.
- The server stores or retrieves the data directly across the SAN.

The metadata can be transferred in-band over the SAN or OOB over an Ethernet link. The latter approach is more common because it avoids IP metadata traffic slowing the data traffic throughput on the SAN. OOB transfer also does not require Fibre Channel HBAs that support IP.

Each server that uses the virtualized part of the SAN must have a special interface or agent that is installed to communicate with the metadata manager. The metadata manager translates the logical data access to physical access for the server. This special interface might be software or hardware.

https://t.me/learningnets



Network-based virtualization offers substantial benefits that overcome the challenges of traditional SAN management solutions. Network-based virtualization simplifies management and consolidates it into a single point. Hosts and storage are now independent of the various management solutions:

- Servers are no longer responsible for volume management and data migration.
- Network-based virtualization enables real-time provisioning of storage, reducing the waste and overhead of overprovisioning storage.
- Existing and heterogeneous storage assets can be consolidated and fully used.
- Data is better protected by simplified snapshot and replication techniques.
- Different classes of data are easier to assign.

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

Summary

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

Summary

- A LUN is a raw disk that can be partitioned into smaller LUNs. Host systems are provided access to the smaller LUNs, therefore hiding the raw disk from the system itself.
- Host systems must not access the same storage area as other host systems, or corruption of data can occur. To provide access to the storage systems, storage system virtualization is used. This can be in the form of OOB, in-band, or network-based storage virtualization.

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

https://t.me/learningnets

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

Lesson 3

Virtualizing Server Solutions

Overview

The purpose of this lesson is to describe the benefits of server virtualization in the data center.

Objectives

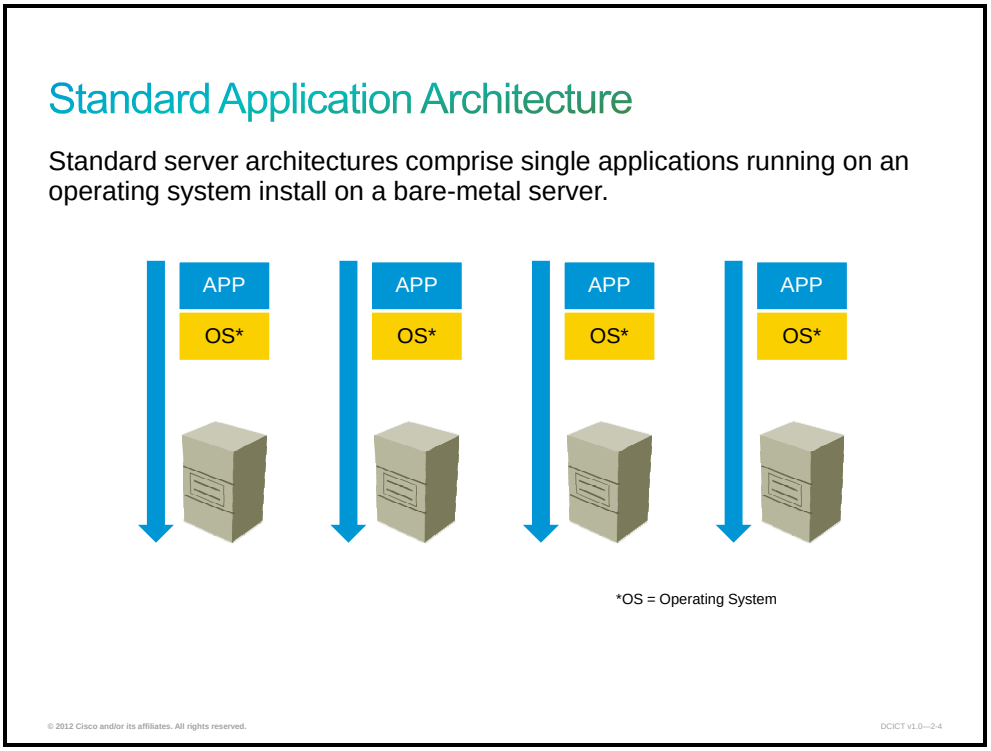
Upon completing this lesson, you will be able to describe the benefits of server virtualization in the data center. You will be able to meet these objectives:

- Describe the benefits of server virtualization in the Cisco data center
- Describe available data center server virtualization solutions

https://t.me/learningnets

Benefits of Server Virtualization

This topic describes the benefits of server virtualization.



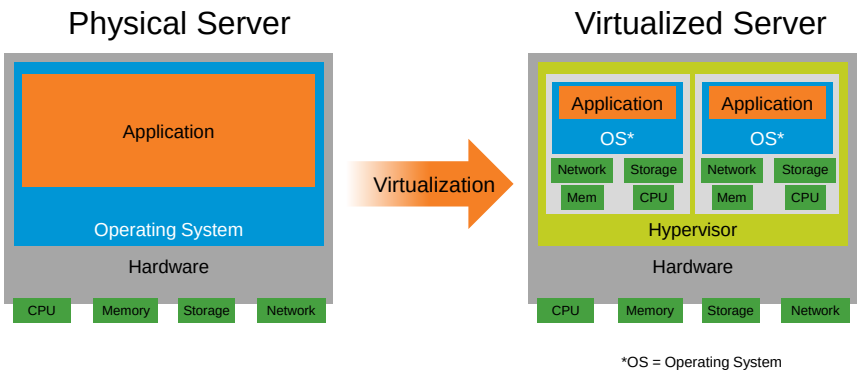
Standard deployments of applications on computer systems consist of one operating system that is installed on a computer. To provide the greatest stability in business environments, good practice dictates that one type of application should be present on an operating system. Otherwise, compatibility issues can cause unforeseen problems.

Although it provides great stability, this approach is not very cost efficient. Over time, growth in the number of applications greatly expands the number of servers in a data center.

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

Server Virtualization

- Abstracting operating system and application from physical hardware
- Hardware independence and flexibility



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

In IT terms, virtualization is a procedure or technique that abstracts physical resources from the services that they provide. In server virtualization, a layer is added between the server hardware and the operating system or systems that are installed in addition to the hardware.

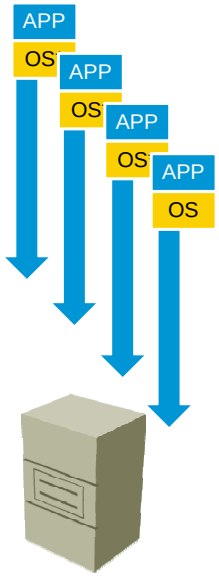
This strategy has several benefits:

- The virtualization layer provides uniformity of hardware to the installed operating system. Suppose that over time, three vendors each provide a server in the data center. Although the three servers are each different, they appear the same to the operating system that will be installed on them. Differences in performance might exist, but not in resource types, such as network cards or graphics adapters.
- The virtualization layer can segment the physical hardware into multiple, separate resource units that all draw from the same physical pool. That means that you can have multiple instances of the same type of operating system (or even different types of operating systems) running in parallel on one server. These separate instances are all completely independent of one another and, unless there is a hardware failure, have similar independence. The instances behave as if they actually were running on individual servers.

https://t.me/learningnets

Basic Virtualized Server Environment

- Virtual server environments allow multiple operating systems to be installed on a single physical server.
- This method of server consolidation provides reduced costs and increased utilization.



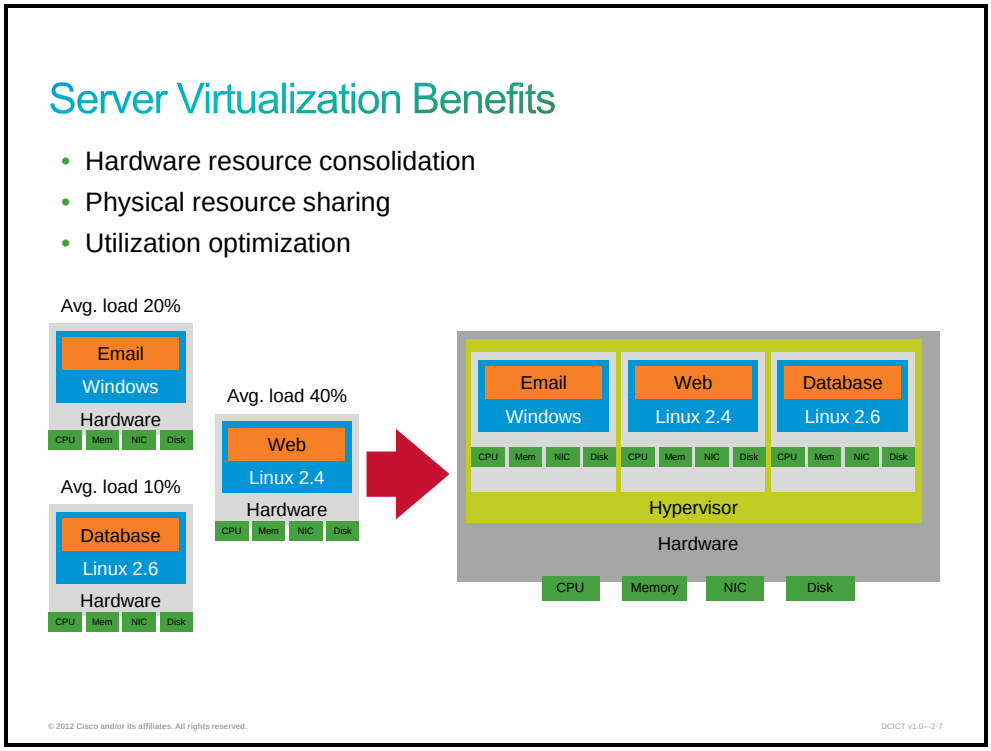
*OS = Operating System

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

Server virtualization allows you to run multiple operating systems on one physical server. This feature can be useful in testing or production environments.

For example, you can copy a production virtual operating system and then practice upgrade procedures on it. This strategy ensures that you can try an upgrade procedure in an environment that is the same as the production environment, without any mistakes affecting users. After you have thoroughly tested the upgrade, evaluated potential consequences, and understood any production impact, a real upgrade can take place.

https://t.me/learningnets



Because of a “one application to one server” mentality, servers in nonvirtualized environments are often underutilized. This condition has a direct impact on both operational and capital expenditure. This underutilization leads to a need for more servers than would otherwise be required. Every server means additional physical space, power, and cooling requirements. As the number grows, management challenges increase as well.

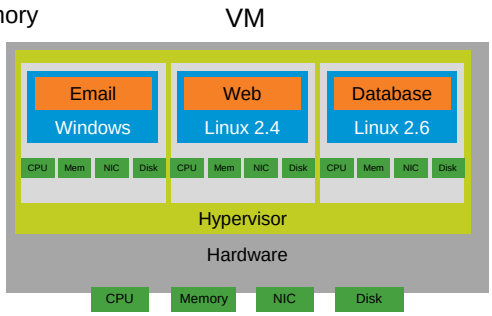
Depending on the average load of existing deployments, three or more servers can often be put onto one piece of hardware. If the data center environment is prepared with virtualization in mind, the ratio can be much greater.

Virtualization also provides more flexibility because the deployment of a new virtual server takes significantly less time than deployment of a physical one.

https://t.me/learningnets

Hypervisor–Abstraction Layer

- Hypervisor or VM Monitor
  - Thin operating system between hardware and VM
  - Controls and manages hardware resources
  - Manages VMs (create, destroy, and so on)
- Virtualizes hardware resources
  - CPU process time-sharing
  - Memory span from physical memory
  - Network
  - Storage



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

The abstraction layer that sits between the operating system and hardware is typically referred to as a hypervisor. This layer is like an operating system but is not intended for installation of applications. Rather, the hypervisor supports installation of multiple types of operating systems.

A hypervisor performs several tasks:

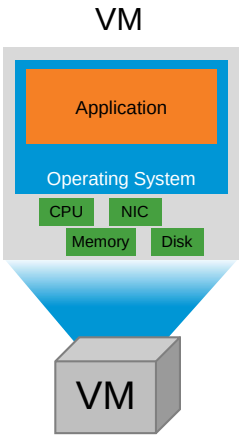
- It provides resources to individual operating systems.
- It uses the resources of the physical server (the host) on which it is installed to provide smaller or greater chunks to the operating systems (virtual machines [VMs]) that are installed on it.
- It provides connectivity between individual VMs as well as between VMs and the outside world.
- It ensures separation between the individual VMs.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

VM Definition

- VM contains operating system and application.
  - Operating system = guest operating system.
  - Guest operating system does not have full control over hardware.
  - Applications are isolated from one another.
- Per VM:
  - vMAC address
  - vIP address
  - Memory, CPU, storage space



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

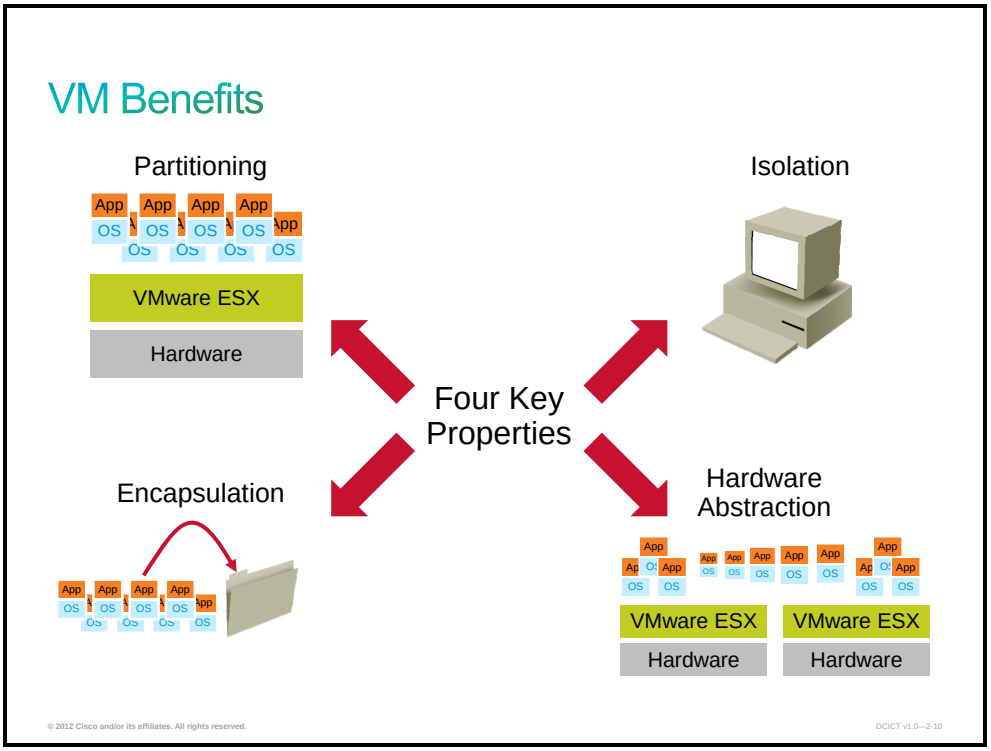
A VM is a logical container that holds all the resources that an operating system requires for normal operation. Components such as a graphics adapter, memory, processor, and networking are present.

As far as the operating system in a VM is concerned, there is no difference between the virtualized components and components in a physical PC or server. The difference is that these components are not physically present, but are rather virtualized representations of host resources. Processor or memory resources are a greater or smaller percentage of actual resources. The virtualized hard disk might be a specially formatted file that is visible as a disk to the virtualized (guest) operating system. The virtualized network interface card (NIC) might simply be a simulated virtual component that the hypervisor manipulates into acting as a physical component.

A VM is supposed to be as close an equivalent as possible of a physical PC or server. The VM must contain the same set of physical identifiers as you would expect from any other such device. These identifiers include the MAC address, IP address, universally unique identifier (UUID) and world wide name (WWN) address.

One of the benefits of a VM is that the administrator can easily set such identifiers to a desired value. This task can be done by simply manipulating the configuration file of the VM.

https://t.me/learningnets



VMs have several benefits as compared to physical devices.

VM Partitioning

VMs allow more efficient use of resources. A single host can serve many VMs, providing that it has sufficient resources to do so.

In practice, the memory capacity of a host and the memory requirements of VMs are the limiting factor.

A hypervisor on a host assigns sufficient resources to every VM that is defined.

VM Isolation

Security and reliability of VMs that share the same host is often a concern for prospective clients. In practice, VMs in a virtualized environment enjoy the same level of separation or security that is present in classic environments.

VMs that share the same host are completely isolated from one another. The only operational hazard is improper design of shared resources, such as network bandwidth or disk access.

Although failure of a crucial hardware component such as a motherboard or a power supply can bring down all the VMs that reside on the affected host, recovery can be much swifter than in a classic environment. Other hosts in the virtual infrastructure can take over VMs from the failed host, and downtime to the affected services is measured in minutes instead of hours.

VM Encapsulation

VMs are a set of files that describe them, define their resource usage, and specify unique identifiers. As such, VMs are extremely simple to back up, modify, or even duplicate.

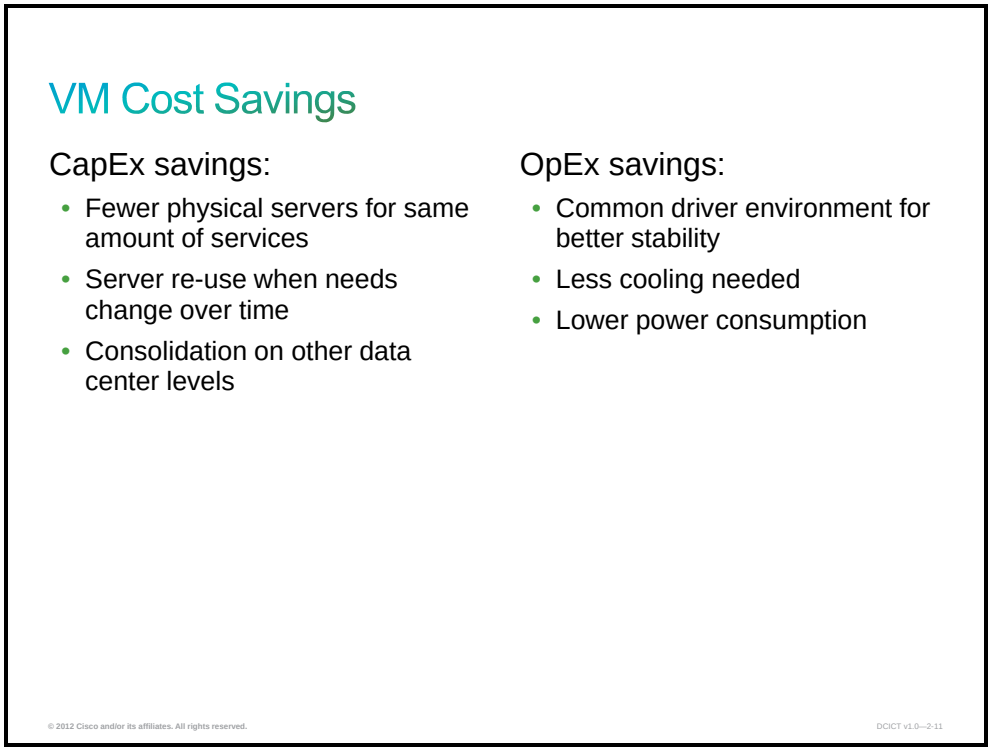
This feature provides benefits for everyday operations such as backup of VMs or deployments in homogeneous environments such as classrooms.

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

VM Hardware Abstraction

VMs are easy to move between hosts. Abstraction can have several benefits:

- **Optimum performance:** If a VM on a given host exceeds the resources of the host, then that VM can be moved to another host that has sufficient resources.
- **Maintenance:** If there is a need to perform maintenance or an upgrade cycle on a host, the VMs from that host can be temporarily redistributed to other hosts. After the maintenance is complete, the process can be reversed, preventing downtime for users.
- **Resource optimization:** If the resource usage of one or more VMs has decreased, one or more hosts might not be needed for a period. In such cases, VMs can be redistributed and emptied hosts can be powered off to save on cooling and power.

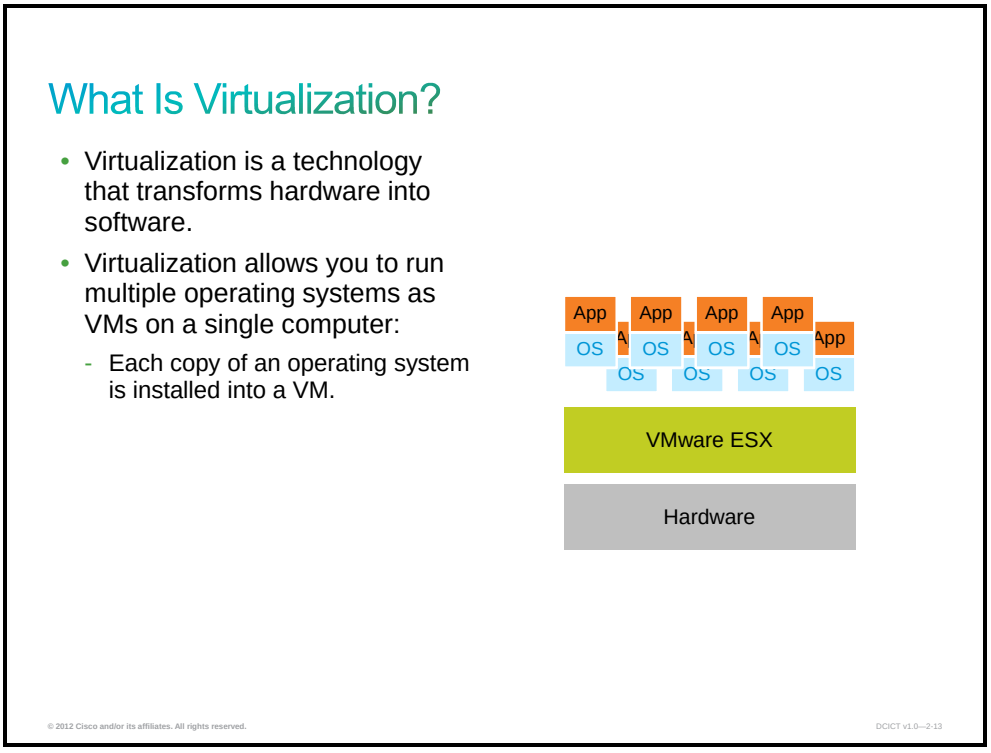


The figure describes cost savings that are associated with a virtualized infrastructure.

https://t.me/learningnets

Available Data Center Server Virtualization Solutions

This topic describes available data center server virtualization solutions.



Virtualization is a technique to take the physical resources of a host machine and turn them into a pool. VMs can tap this pool, depending on what they need.

Whereas the VMs share the resources from the pool, the resources that are assigned to one VM become unavailable to other VMs.

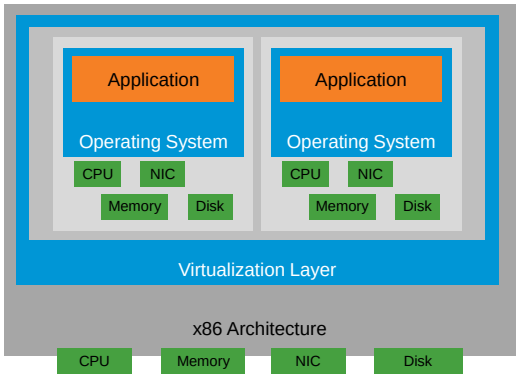
Virtualized operating systems are not simulated. They are complete, proper, off-the-shelf operating systems that run on virtualized hardware.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

How Does Virtualization Work?

A virtualization layer is installed and uses either a hosted or hypervisor architecture.



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

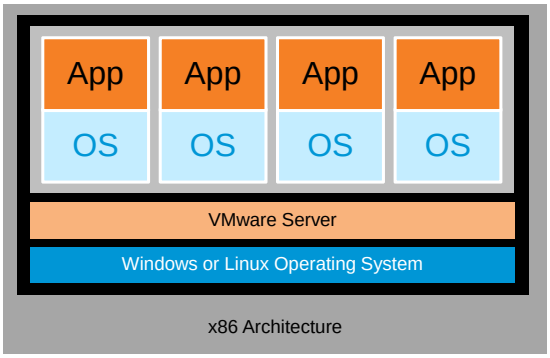
Several types of virtualization are available:

- **Full virtualization:** Complete simulation of the hardware environment allows unmodified operating systems to run within virtualized sandboxes. There are two types of full virtualization:
  - Host operating system-based virtualization
  - Bare-metal virtualization
- **Partial virtualization:** When some (but not all) of the hardware environment is simulated, certain guest software might need to be modified to run in the environment.
- **Paravirtualization:** No simulation of hardware requires substantial modification of guest operating systems so that they can be installed in different virtual domains on one server.

https://t.me/learningnets

Host Operating System-Based Virtualization

A host-based virtualization system requires an operating system (such as Windows or Linux) to be installed on the computer.



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

Host operating system-based virtualization requires a PC or server that already has an installed operating system, such as a Microsoft Windows or a Linux environment. In addition to the operating system, a virtualization application is installed. Within that application, VMs are deployed.

The benefit of such an approach is that there is no need for a separate PC for everyday use and a machine for virtualization. The drawback is that the host operating system uses up some resources that could be assigned to the VMs.

An example of such virtualization is a VMware workstation.

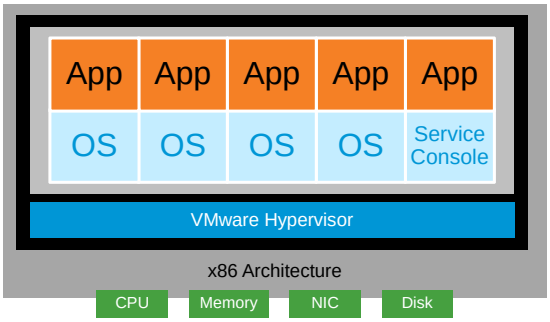
This type of virtualization is most often used for application development and testing.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Virtualization Using a Bare-Metal Hypervisor

A bare-metal hypervisor system does not require an operating system. The hypervisor *is* the operating system.



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

Bare-metal hypervisor virtualization means that the virtualization logic acts as its own mini operating system. VMs are installed in addition to that mini operating system. This method minimizes the amount of resources that are unavailable to the VMs. Also, this method avoids potential bugs and security vulnerabilities of the host operating system.

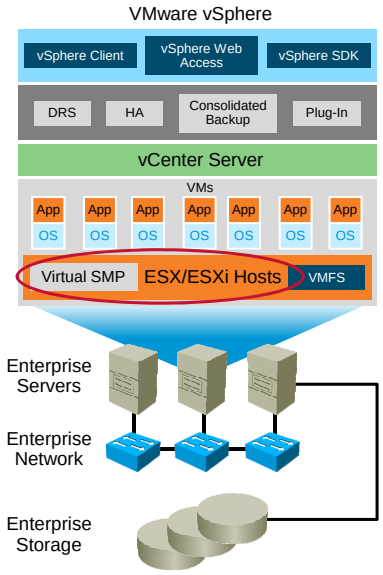
Such an approach is often used for server deployments in which a server is dedicated as a host and all its resources are dedicated for VM deployments.

VMware ESX and ESXi are examples of such an implementation.

https://t.me/learningnets

ESX and ESXi: Virtualization Platform

- ESX and ESXi are bare-metal, efficient, and reliable hypervisors running directly on the server.
- ESX and ESXi abstract CPU, memory, storage, and networking for multiple VMs.



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

VMware ESX and ESXi are two main products that are used for data center virtualization deployments.

ESXi is a small-footprint version of the VMware ESX server, with installed space consumption of approximately 70 MB. As such, ESXi is well positioned for installations to bootable USB sticks or direct integration into server hardware itself.

Both ESX and ESXi servers have similar characteristics in all aspects of VM creation and running.

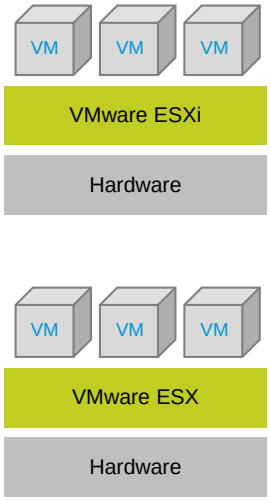
Although ESXi is now the main product, many installations of ESX are still used in data centers.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

ESX and ESXi Features

- Can use standard and distributed vSwitches, NIC teaming, and VLANs
- Can use the VMware vStorage VMFS for storing VMs
- Can be managed by VMware vCenter Server
- Can take advantage of various VMware vSphere features, such as VMware vMotion
- Can be accessed by using the VMware vSphere Client



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

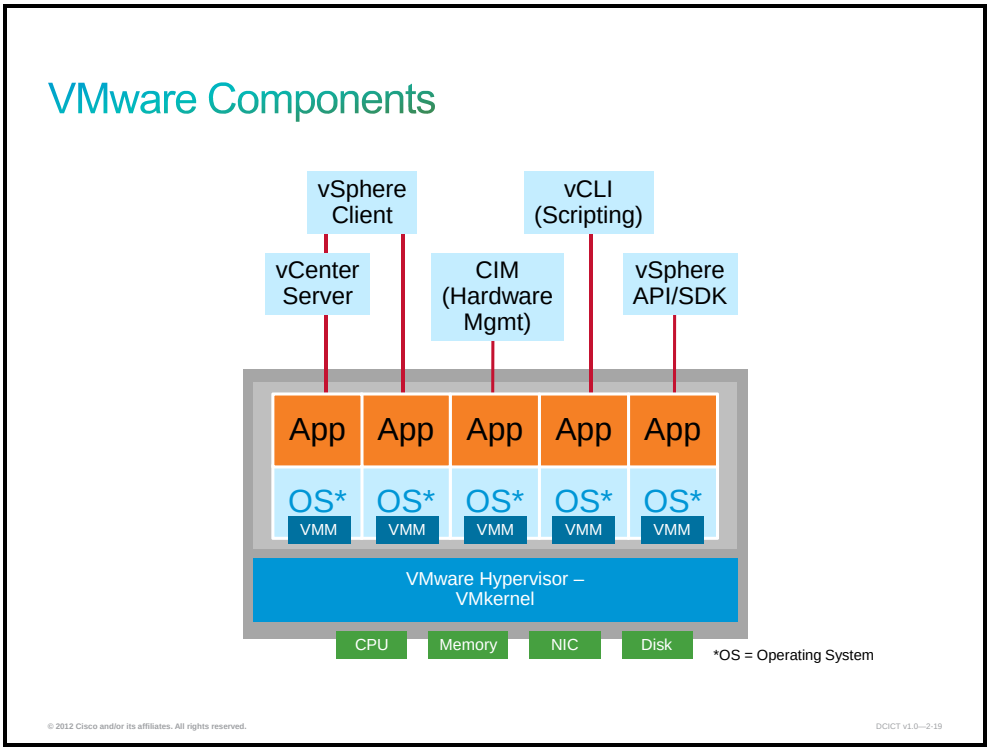
The VMware ESX or ESXi environment can use multiple network connectivity options. A standalone host can use a virtual switch (vSwitch). Multiple hosts, working as a group, can use a distributed version of the vSwitch. To increase bandwidth and reliability, network teaming is supported. For security and traffic separation, there is support for VLANs.

For storage requirements, there are multiple possibilities. Local storage can be used, although doing so is highly unusual because it reduces the mobility options of VMs. Mount points on either SAN or network attached storage (NAS) locations are more common.

With larger deployments, or for advanced functionalities such as VMware vMotion or a distributed vSwitch, the VMware vCenter component can manage an ESX or ESXi environment as a group.

VMware vSphere Client serves as a client-side application through which administrative tasks are carried out.

https://t.me/learningnets

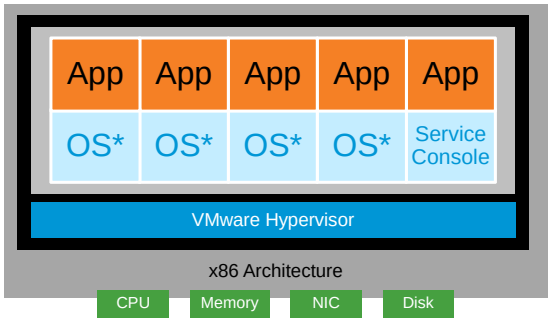


The figure is a visual overview of the VMware infrastructure components.

https://t.me/learningnets

VMware Server

- Also referred to as a host
- Provides physical resources for VMs



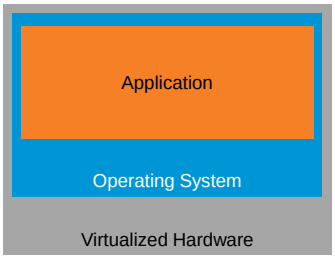
*OS = Operating System

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

A VMware server is typically referred to as a host because it hosts a multitude of VMs. Depending on the amount of physical resources in the host, the number of VMs can be greater or smaller.

Virtual Machine

- Refered to as a VM
- Behaves as a real PC or server
- Gets resources from a host
- Can be migrated between hosts



VM

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

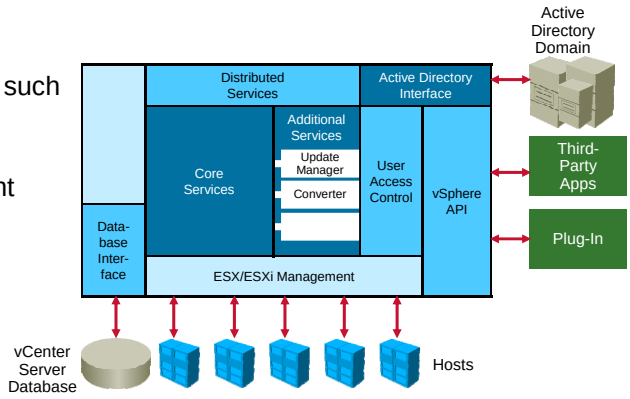
A VM is a logical representation of a real PC or server. A VM consists of virtualized hardware resources that serve as a basis for installation of a standard operating system.

Because a VM is a logical representation that resides on a standardized hypervisor layer, VMs are also easy to migrate between hosts.

https://t.me/learningnets

vCenter Server

- Management point for multiple hosts
- Can be a physical server or a VM
- Required for some advanced functions such as vMotion and the distributed vSwitch
- Should be redundant



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

VMware vCenter Server is a central management component. vCenter Server requires an additional license but serves as a focal point for the management of multiple hosts.

vCenter Server correlates traffic between hosts for functionalities that span more than one host. Features such as vMotion, VMware vNetwork Distributed Switch (vDS), and disaster tolerance require the vCenter Server.

As such, the vCenter Server is considered a crucial component of advanced ESX and ESXi deployments. Failure of a vCenter Server will not stop production traffic, but functionalities that depend on the server will be unavailable until it is restored. For that reason, having a redundant vCenter Server deployment is highly recommended.

vCenter Server can exist as a physical device or can be installed as a VM on one of the hosts that it manages. In a redundant installation, the two instances should be installed on two different hosts for greater availability.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

vSphere Client

- GUI for vCenter Server
- Can manage vCenter environment or individual hosts
- Has a web alternative



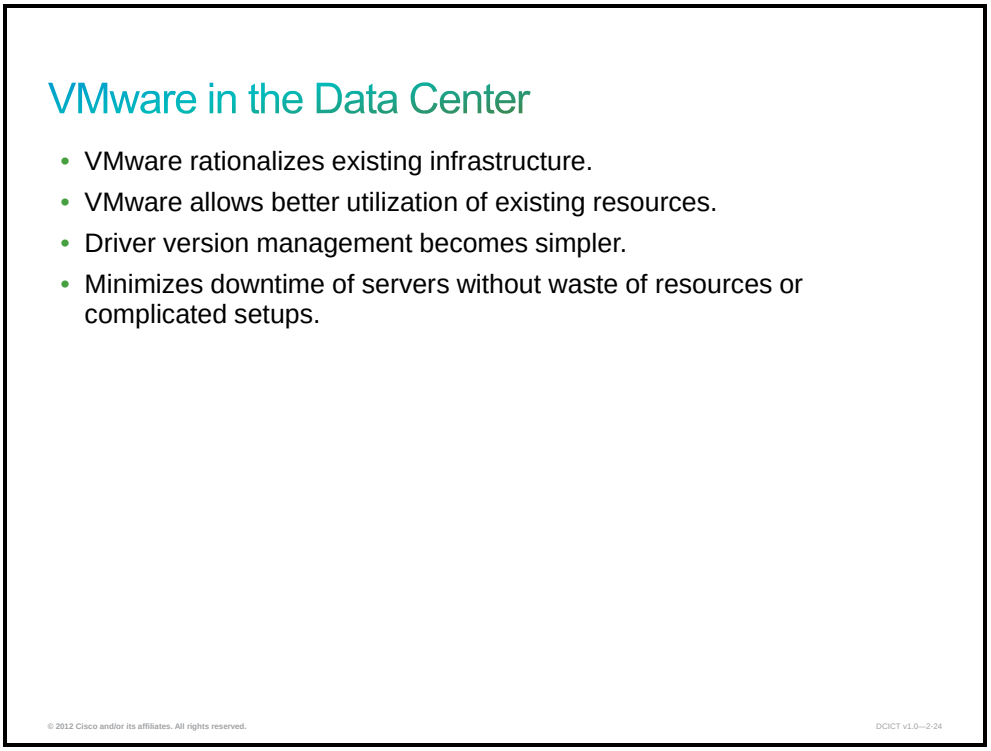
© 2012 Cisco and/or its affiliates. All rights reserved.

vSphere Client is a GUI component of the VMware environment. This component connects to a vCenter Server and allows interaction with the ESX and ESXi servers and VMs.

Depending on access permissions, users can access parts or all of the virtualized environment. For example, users can connect to the consoles of the VMs but be unable to start or stop the VMs or change the parameters of the host on which the VMs run.

To simplify user management, the VMware environment supports integration with the Microsoft Active Directory user database. Therefore, when users log on to the vSphere Client, they can use their domain credentials. This feature creates a more secure environment by avoiding several sets of credentials. (The use of multiple credentials increases the risk of people writing them down so that they do not forget or confuse them.)

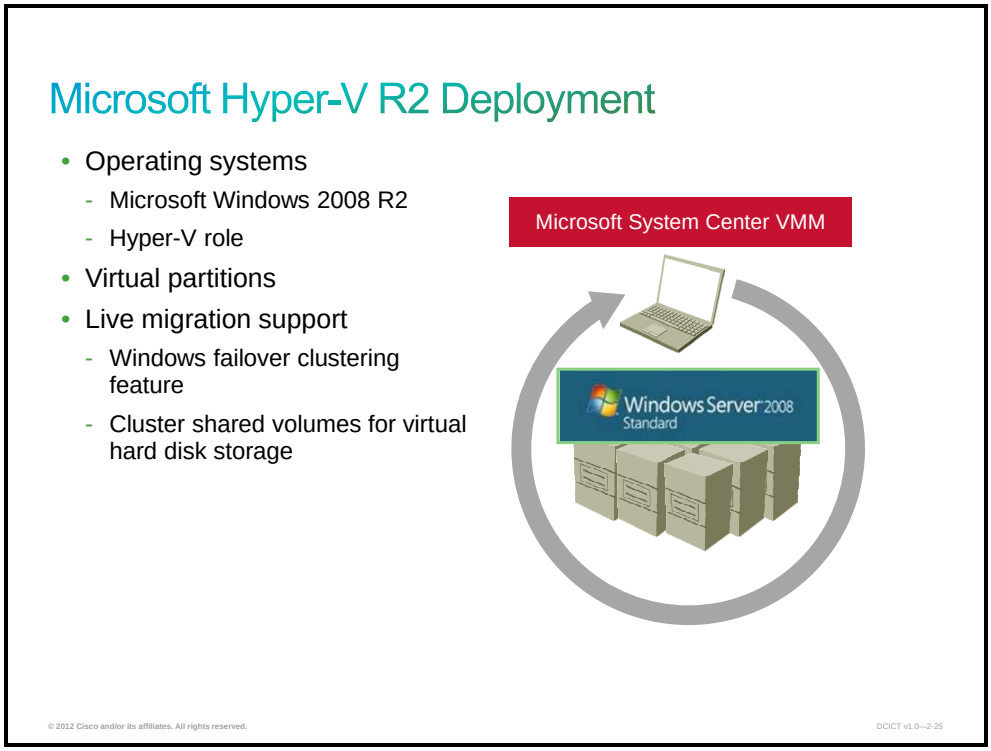
https://t.me/learningnets



VMware has many tangible benefits for a data center environment:

- Existing infrastructure is often underutilized. VMware allows higher utilization of existing resources without the purchase of additional equipment.
- Change management is simplified, and the complexity that is usually associated with large environments is minimized. This result is mainly because all VMs share the same hardware components and have the same drivers. This fact eliminates the driver hunt that can be associated with large environments with varying server types. This benefit also simplifies troubleshooting.
- Even with basic functionality, service downtime can be decreased by quickly restarting failed VMs on an alternative host. VMware boot times are generally shorter than the boot times for a physical server.

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



Microsoft Hyper-V is an alternative to VMware virtualization products.

Overview

Hyper-V is offered as a server role that is packaged into the Microsoft Windows Server 2008 R2 installation or as a standalone server. In either case, Hyper-V is a hypervisor-based virtualization technology for x64 versions of Windows Server 2008. The hypervisor is a processor-specific virtualization platform.

Hyper-V isolates operating systems that run on the VMs from one another through partitioning or logical isolation by the hypervisor. Each hypervisor instance has at least one parent partition that runs Windows Server 2008. The parent partition houses the virtualization stack, which has direct access to hardware devices such as NICs. This partition is responsible for creating the child partitions that host the guest operating systems. The parent partition creates these child partitions by using the hypercall application programming interface (API), which is exposed to Hyper-V.

Virtual Partitions

A virtualized partition does not have access to the physical processor, nor does it manage its real interrupts. Instead, the partition has a virtual view of the processor and runs in a guest virtual address space. Depending on the configuration of the hypervisor, this space might not be the entire virtual address space. A hypervisor can choose to expose only a subset of the processors to each partition. The hypervisor, using a logical synthetic interrupt controller, intercepts the interrupts to the processor and redirects them to the respective partition. Hyper-V can hardware-accelerate the address translation between various guest virtual address spaces by using an I/O memory management unit (IOMMU). An IOMMU operates independently of the memory management hardware that the CPU uses.

Child partitions do not have direct access to hardware resources; instead, they have a virtual view of the resources, in terms of virtual devices. Any request to the virtual device is redirected, via the virtual machine bus (VMBus), to the devices in the parent partition, which manages the requests. The VMBus is a logical channel that enables interpartition

https://t.me/learningnets

communication. The response is also redirected via the VMBus. If the devices in the parent partition are also virtual devices, then the response is redirected further until it reaches the parent partition, where it gains access to the physical devices.

Parent partitions run a virtualization server provider, which connects to the VMBus and processes device-access requests from child partitions. Child partition virtual devices internally run a virtualization server client, which redirects the requests, via the VMBus, to virtualization server providers in the parent partition. This entire process is transparent to the guest operating system.

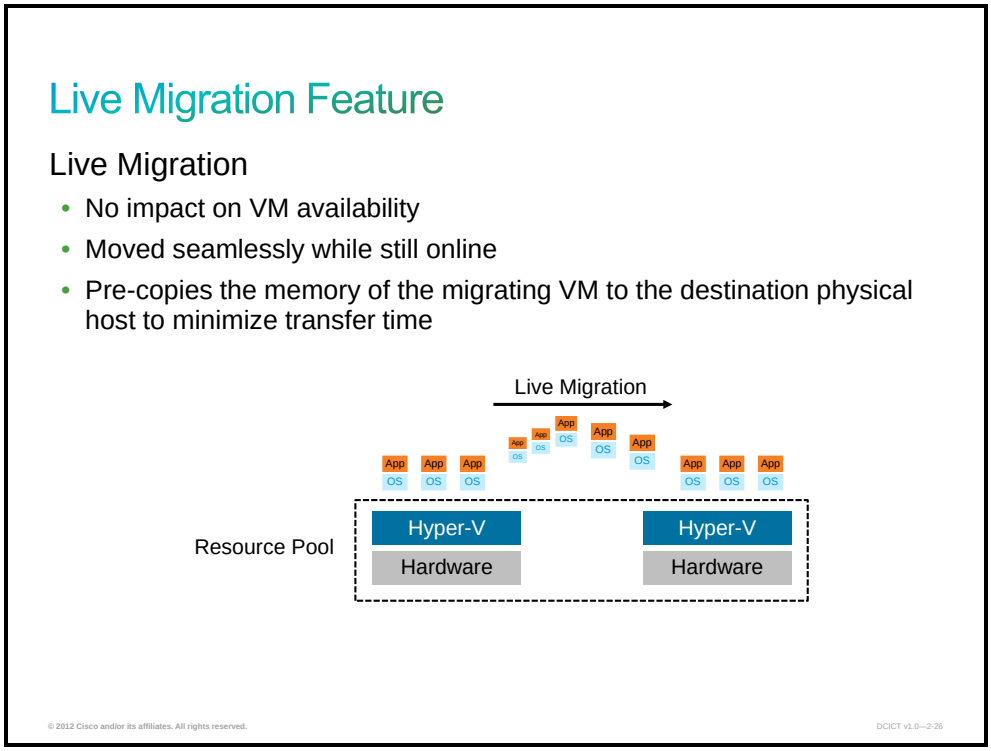
Cluster Shared Volumes

With Windows Server 2008 R2, Hyper-V uses cluster shared volumes storage to support live migration of Hyper-V VMs from one Hyper-V server to another. Cluster shared volumes enable multiple Windows servers to access SAN storage by using one consistent namespace for all volumes on all hosts.

Multiple hosts can access the same logical unit number (LUN) on SAN storage so that virtual hard disk files can be migrated from one hosting Hyper-V server to another. Cluster shared volumes are available as part of the failover clustering feature of Windows Server 2008 R2.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Live Migration Feature

Hyper-V in Windows Server 2008 R2 supports live migration through the use of cluster shared volumes. This architecture allows an individual VM that is still online and actively supporting user sessions to be moved seamlessly to a different physical (Hyper-V) server without disruption.

Hyper-V live migration moves running VMs without affecting availability to users. By precopying the memory of the migrating VM to the destination physical host, live migration minimizes the transfer time of the VM. A live migration is deterministic, meaning that the administrator or script that initiates the live migration can control which computer is the destination for the live migration. The guest operating system of the migrating VM is unaware that the migration is happening, so no special configuration for the guest operating system is needed.

https://t.me/learningnets

Summary

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

Summary

- Server consolidation in the data center provides many benefits, such as cost savings, smaller footprint, and increased utilization of available physical equipment.
- Vendors such as Microsoft offer virtualization solutions. These solutions permit multiple operating systems to run in separate VMs on the same physical host.

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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

Using the Cisco Nexus 1000V Series Switch

Overview

The purpose of this lesson is to describe the problems that are solved by Cisco Nexus 1000V Series Switches and how the Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM) integrate with VMware ESX.

Objectives

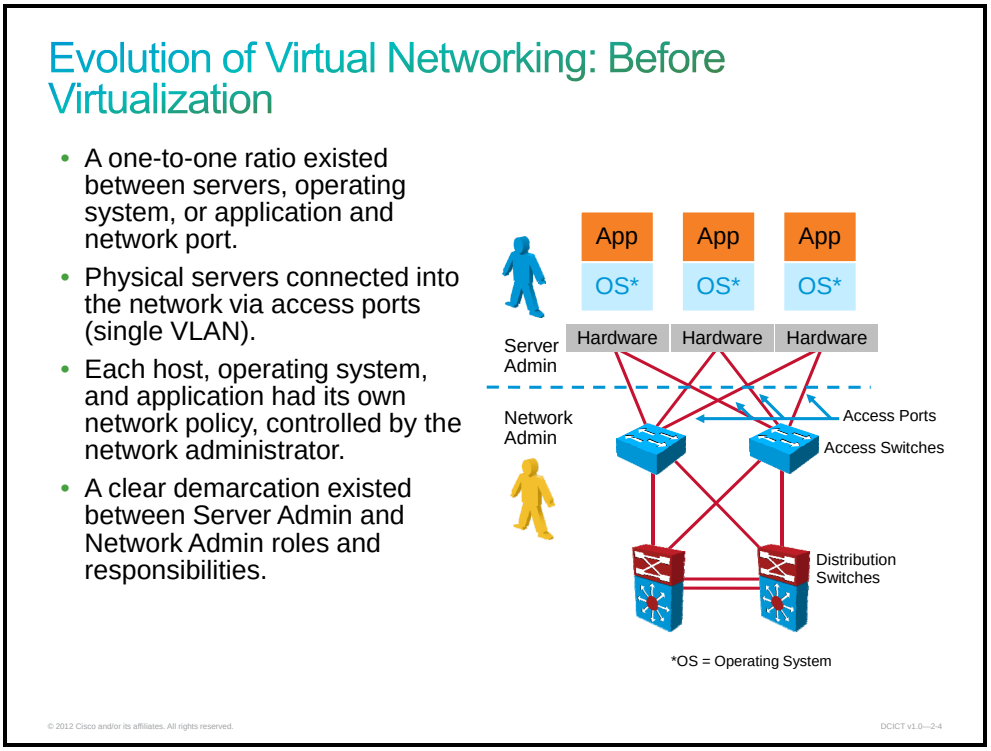
Upon completing this lesson, you will be able to describe the problems that Cisco Nexus 1000V Series switches solve. You will be able to meet these objectives:

- Describe the limitations of VMware vSwitch
- Describe the advantages of VMware vDS
- Describe how the Cisco Nexus 1000V Series switch takes network visibility to the VM level
- Describe how the VSM and VEM integrate with VMware ESX or ESXi and vCenter

https://t.me/learningnets

Limitations of VMware vSwitch

This topic describes the limitations of VMware vSwitch.



Before virtualization, each server ran its own operating system, usually with a single application. The network interface cards (NICs) were connected to access layer switches to provide redundancy. Network security, quality of service (QoS), and management policies were created on these access layer switches and applied to the access ports that corresponded to the appropriate server.

If a server needed maintenance or service, it was disconnected from the network, during which time any crucial applications would need to be offloaded manually to another physical server. Connectivity and policy enforcement were static and seldom required any modifications.

Server virtualization has made networking, connectivity, and policy enforcement much more challenging. A feature such as VMware vMotion, in which devices running applications can move from one physical host to another, is one example.

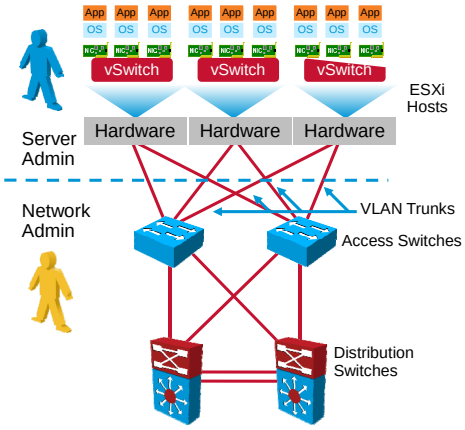
The challenges include the following:

- Providing network visibility from the virtual machine (VM) virtual NICs (vNICs) to the physical access switch
- Creating and applying policies to the vNICs
- Providing consistent mobility of the policies that are applied to the VMs during a vMotion event

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

Evolution of Virtual Networking: vSwitches

- Layer 2 switches become embedded within the ESX hypervisor to switch packets between VMs and the outside world.
- Multiple VMs are required to share the same physical uplinks to the network (VMNICs) as well as the same network policy:
  - No longer a one-to-one relationship between server and network port.
  - Network visibility ends at the physical access port.
- Segmentation between VMs is provided by 802.1Q VLANs:
  - Requires VLAN trunks for server connectivity into the network
- Server or virtualization administrator owns the virtual network configuration and manages it through vCenter Server.

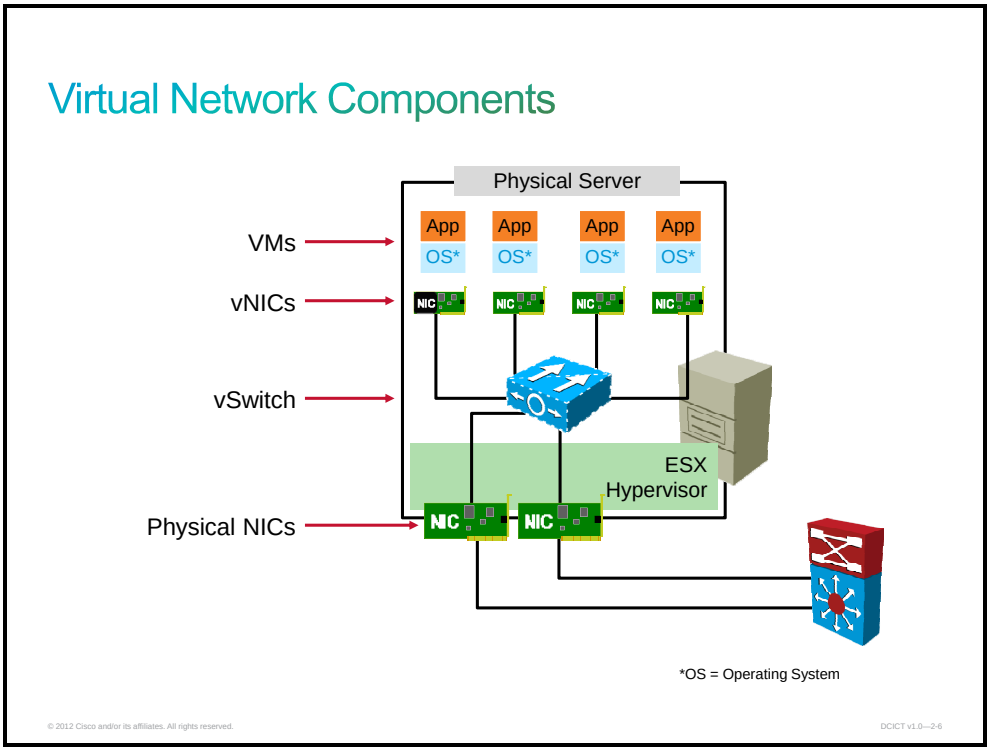


© 2012 Cisco and/or its affiliates. All rights reserved.

The VMware server virtualization solution extends the access layer into the VMware ESX or ESXi server with the VM networking layer. These components are used to implement server virtualization networking:

- **Physical networks:** Physical devices connect VMware ESX hosts for resource sharing. Physical Ethernet switches are used to manage traffic between ESX hosts, the same as in a regular LAN environment.
- **Virtual networks:** Virtual devices run on the same system for resource sharing.
- **Virtual Ethernet switch (vSwitch):** Similar to a physical switch, a vSwitch maintains a table of connected devices, which is used for frame forwarding. Can be connected via an uplink to a physical switch via a physical NIC. Does not provide the advanced features of a physical switch.
- **Port group:** A port group is a subset of ports on a vSwitch for VM connectivity.
- **Physical NIC:** A physical NIC is used to create the uplink from the ESX or ESXi host to the external network. (This NIC is represented by an interface known as a VMNIC.)

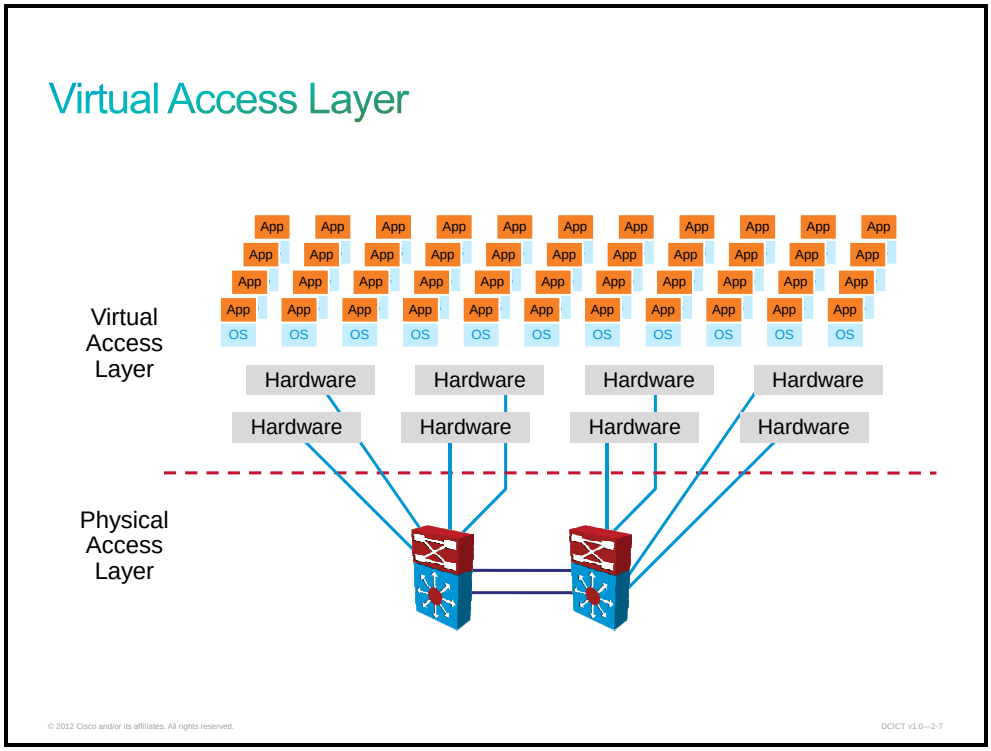
https://t.me/learningnets



In the VMware environment, the network is one component that can be virtualized. The network component includes most aspects of networking, apart from the NIC in the host. Even then, the NICs most often play the role of uplink ports in the vSwitches that are created on the VMware hypervisor level.

Furthermore, VMs have a selection of types of vNICs that act the same as a physical NIC would in a physical PC or server.

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

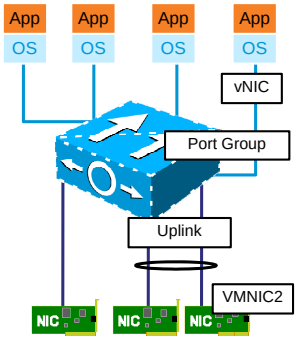


The switching component in the ESX or ESXi host creates the need to manage a virtual access layer as well as the physical access layer. The virtual switching infrastructure creates this virtual access layer, which the VMware administrators manage.

https://t.me/learningnets

vNetwork Standard Switch Overview

- Associates physical (VMNIC) and vNICs
- Connectivity for
  - VM communication within and between ESX hosts
  - Service console for ESX management
  - VMKernel for vMotion, iSCSI, fault tolerance (FT) logging
- Service console port: Assigned to a VLAN
- VMKernel port (or ports): Assigned to a VLAN
- VM port group (or ports):
  - Assigned to a VLAN
  - VMs assigned to a port group
- Uplink (or uplinks):
  - External connectivity
  - VMNIC associated with single vSwitch only



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

A VMware vSwitch is a virtual construct that performs network switching between the VMs on a particular VMware host and the external network.

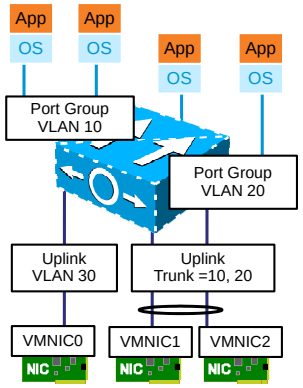
Zero, one, or multiple (as many as 32) network ports can be assigned to a vSwitch. If more ports exist than are assigned, the result is both greater bandwidth and reliability. If the vSwitch does not have a port assigned to it, then it will switch the traffic between VMs within the host only.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

vNetwork Standard Switch Operation

- Multiple switches on a single ESX or ESXi host:
  - No internal communication between vSwitches
- Operates as physical Layer 2 switch:
  - Forwards frames per MAC addresses
  - Maintains MAC address table
  - Internal switching for VMs
- Supports:
  - Trunk ports with 802.1q for VLANs
  - Port channel: NIC teaming
- Cisco Discovery Protocol



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

A single VMware host can have multiple vSwitches configured, if needed. Those switches will be separate from one another, as the VMs are.

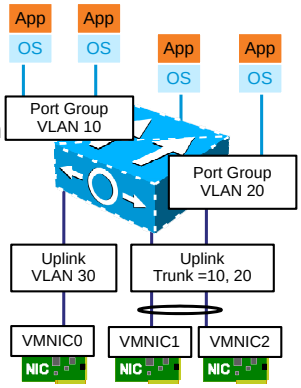
vSwitches are Layer 2 devices. They do not support routing of Layer 3 traffic. The vSwitch performs switching of traffic between VMs that are present on the same host. Other traffic is forwarded to the uplink port.

vSwitches support the trunking functionality, port channels, and the Cisco Discovery Protocol for discovering and responding to the neighbor device network queries.

https://t.me/learningnets

vNetwork Standard Switch Operation (Cont.)

- No participation in Spanning Tree Protocol (STP), Dynamic Trunking Protocol (DTP), or Port Aggregation Protocol (PAgP)
  - Internal vSwitch (or vSwitches)
  - Testing and traffic isolation
  - VMs accessible via VMware vCenter
- Virtual guest tagging
  - VLAN 4095
  - Tagged traffic passed up to guest operating system
- NIC teaming
  - Connect multiple VMNICs to a single vSwitch
- Outbound load balancing only
  - vSwitch port based
  - Source MAC
  - IP hash

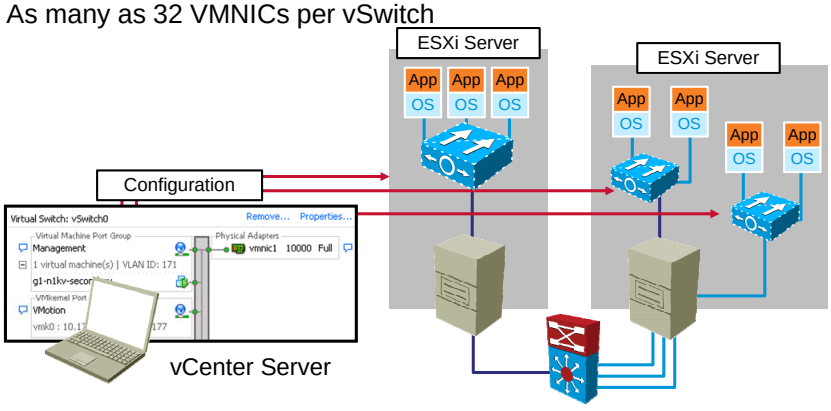


© 2012 Cisco and/or its affiliates. All rights reserved.

The figure continues to describe the standard switch operation in a virtual network.

vNetwork Standard Switch Configuration

- Configured and managed in software by VMware administrator
- Created in the ESX or ESXi host
- Managed as separate virtual networks
- As many as1016 useable ports per vSwitch
- As many as 32 VMNICs per vSwitch

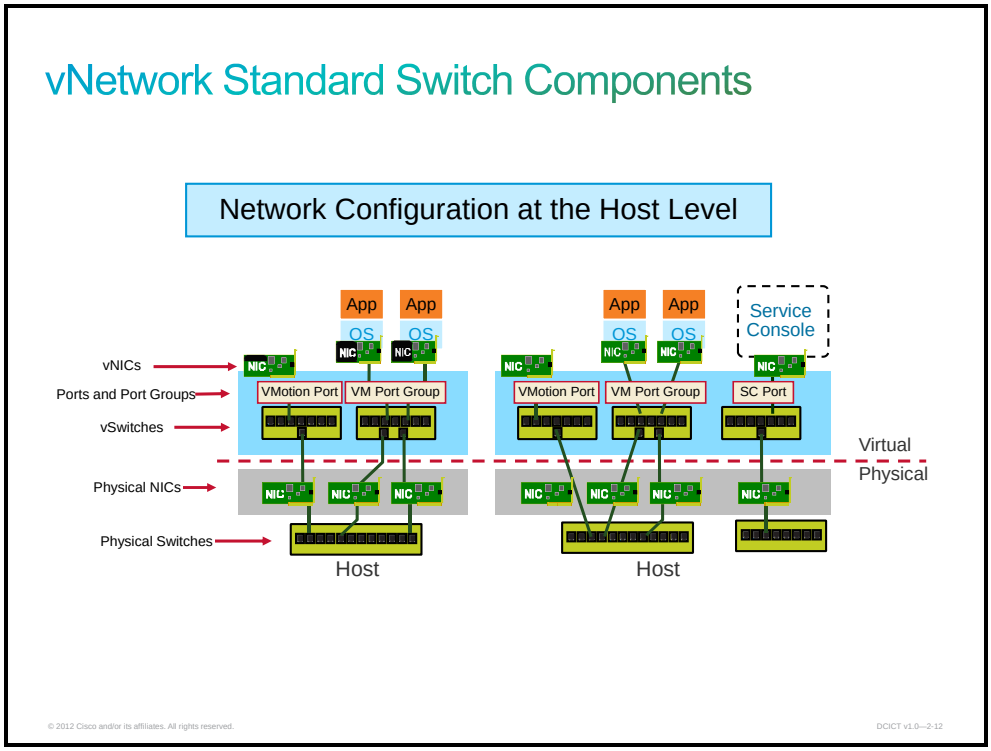


© 2012 Cisco and/or its affiliates. All rights reserved.

A vSwitch is created on the ESX or ESXi host where it is needed. A single host can have multiple vSwitches, and every vSwitch acts as an independent entity.

A vSwitch can have as many as 1016 ports, and as many as 32 NIC ports can be assigned to it.

https://t.me/learningnets



The vSwitch consists of several components, some virtual and some physical:

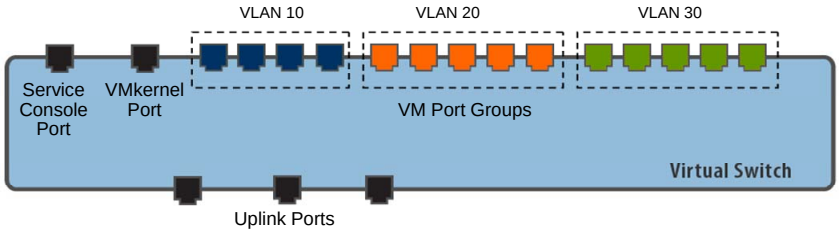
- **vNICs:** Present on VMs and connected to the vSwitch through its virtual ports
- **Port groups:** Logical groupings of switch ports that have the same configuration
- **Physical NICs:** Act as switch uplinks and provide connectivity to the external network.

https://t.me/learningnets

vNetwork Standard Switch Ports

A vSwitch allows these connection types:

- VMkernel port
- Service console port (ESX only)
- VM port group



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

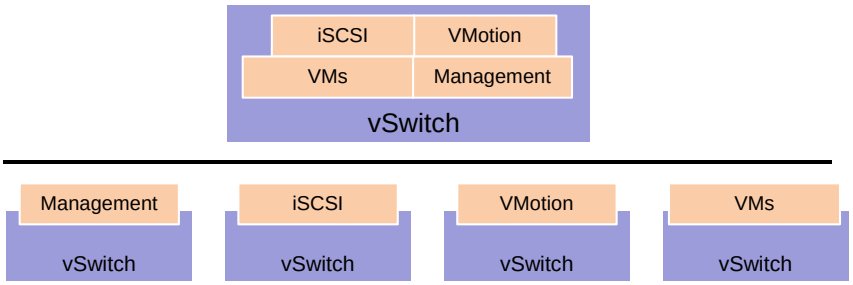
The most common types of switch ports are the VM ports and port groups. These ports represent connection points for the VMs.

Two other port types might be present on the VMware host. First is the VMkernel port. This port type is used for advanced functions such as vMotion or access to Network Files System (NFS) and Internet Small Computer Systems Interface (iSCSI) storage. Management traffic also flows over this port. Second is the service console port. This port type, which is available only on the ESX (not the ESXi) version of the host, is used to gain console-type CLI access to the host management.

https://t.me/learningnets

vNetwork Standard Switch Examples

Different networks can coexist on the same vSwitch or on separate vSwitches.



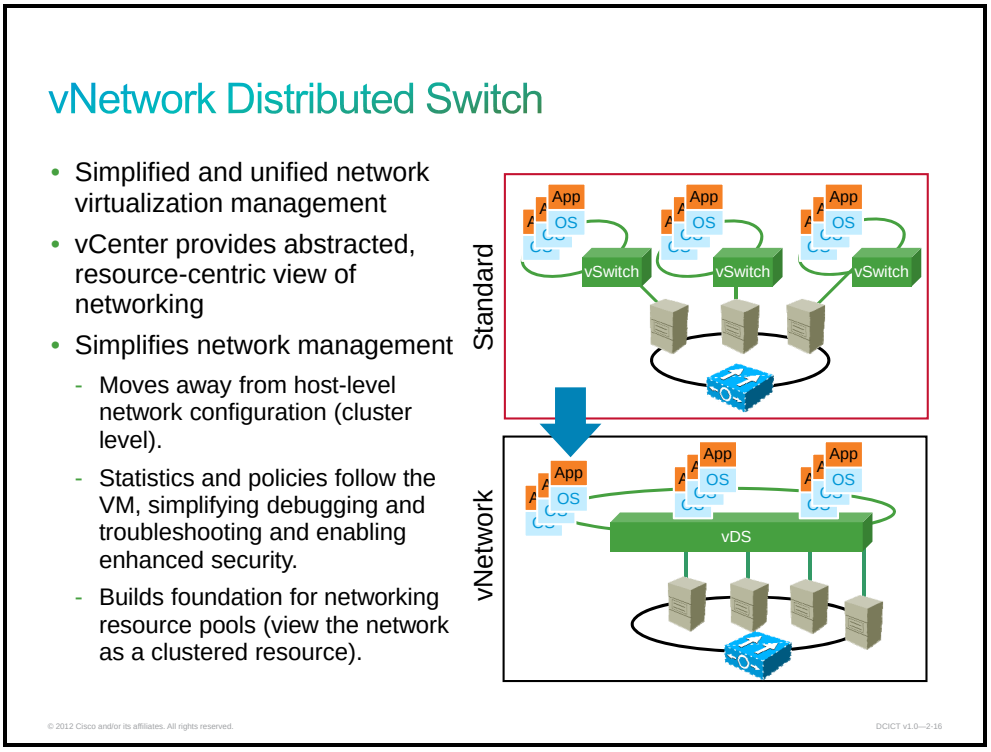
© 2012 Cisco and/or its affiliates. All rights reserved.

The figure shows an example of network segregation, achieved by the use of either VLANs or different vSwitches. Either solution might be appropriate, depending on the circumstances.

https://t.me/learningnets

Advantages of VMware vDS

This topic describes the benefits of VMware vNetwork Distributed Switch (vDS).



VMware vSphere 4 introduced vDS, a distributed virtual switch (vSwitch). With vDS, multiple vSwitches within an ESX or ESXi cluster can be configured from a central point. The vDS automatically applies changes to the individual vSwitches on each ESX or ESXi host.

The feature is licensed and relies on VMware vCenter Server. The feature cannot be used for individually managed hosts.

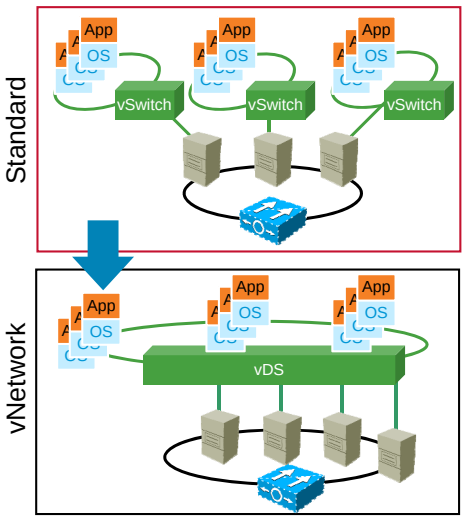
The VMware vDS and vSwitch are not mutually exclusive. Both devices can run in tandem on the same ESX or ESXi host. An example of this type of configuration would be running the Cisco Nexus 1000V VSM on a host that it is controlling. In this scenario, the VSM runs on a vSwitch that is configured for VSM connectivity, while controlling a vDS that runs a VEM on the same host.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

vNetwork Distributed Switch

- Simplified and unified network virtualization management
- vCenter provides abstracted, resource-centric view of networking
- Simplifies network management
  - Moves away from host-level network configuration (cluster level).
  - Statistics and policies follow the VM, simplifying debugging and troubleshooting and enabling enhanced security.
  - Builds foundation for networking resource pools (view the network as a clustered resource).



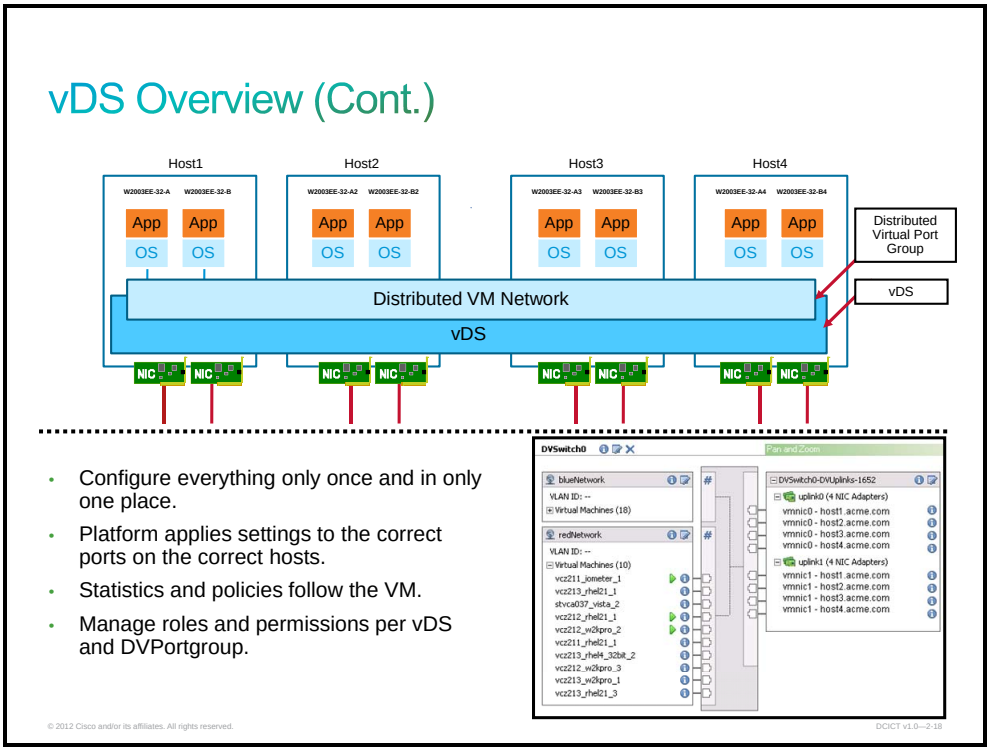
© 2012 Cisco and/or its affiliates. All rights reserved.

The vDS is the next step to a vSwitch. Whereas the vSwitch is managed individually on the host on which it was created, the vDS is managed globally across all hosts.

The vDS provides greater management uniformity because segmentation can be avoided.

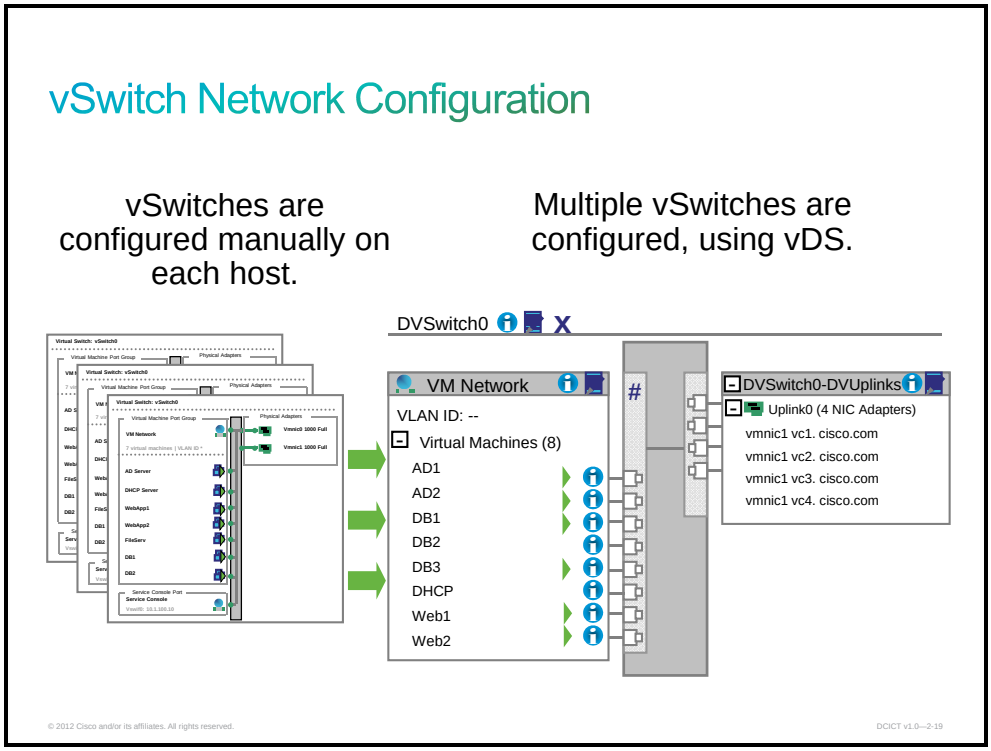
The requirement for a vDS is an installed vCenter Server.

https://t.me/learningnets



The vDS adds additional functionality and simplified management to the VMware network. The vDS adds the ability to use private VLANs (PVLANs), perform inbound rate limiting, and track VM port state with migrations. Additionally, the vDS is a single point of network management for VMware networks. The vDS is a requirement for the Cisco Nexus 1000V Series switch.

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



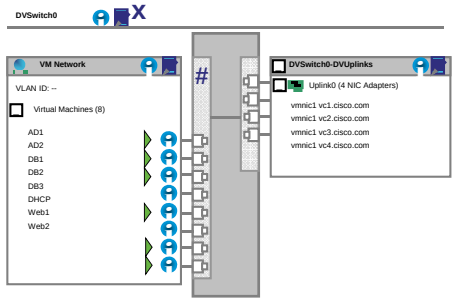
The figure shows the conceptual difference in management for a standard vSwitch environment versus a vDS environment. The standard vSwitch requires a separate configuration from a separate management panel. The vDS requires only one management panel for one switch that spans multiple hosts.

https://t.me/learningnets

vDS Enhancements

VMware vDS offers several enhancements to VMware switching:

- Port state migration (statistics and port state follow VM)
- Rx rate limiting
- PVLANs



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

PVLAN support enables broader compatibility with existing networking environments, using PVLAN technology. PVLANs enable users to restrict communication between VMs on the same VLAN or network segment. This feature significantly reduces the number of subnets that are needed for certain network configurations.

PVLANs are configured on a vDS with allocations made to the promiscuous PVLAN, the community PVLAN, and the isolated PVLAN. Within the subnet, VMs on the promiscuous PVLAN can communicate with all VMs. VMs on the community PVLAN can communicate among themselves and with VMs on the promiscuous PVLAN. VMs on the isolated PVLAN can communicate only with VMs on the promiscuous PVLAN.

Note Adjacent physical switches must support PVLANs and be configured to support the PVLANs that are allocated on the vDS.

Network vMotion is the tracking of VM networking state, such as counters or port statistics, as the VM moves from host to host on a vDS. This tracking provides a consistent view of a virtual network interface, regardless of the VM location or vMotion migration history. This feature greatly simplifies network monitoring and troubleshooting activities in which vMotion is used to migrate VMs between hosts.

The vDS expands upon the egress-only traffic-shaping feature of standard switches with bidirectional traffic-shaping capabilities. Egress (from the VM to the network) and ingress (from the network to the VM) traffic-shaping policies can now be applied on port group definitions.

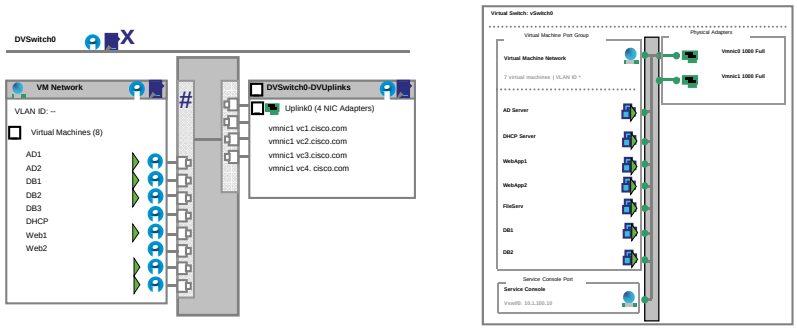
Traffic shaping is useful when you want to limit the traffic to or from a VM or group of VMs. This policy is usually implemented to protect a VM or other traffic in an oversubscribed network. Policies are defined by three characteristics: average bandwidth, peak bandwidth, and burst size.

https://t.me/learningnets

vSwitch and vDS

VMware vDS and vSwitch are not mutually exclusive:

- Physical NIC ports are assigned to either a vSwitch or vDS.
- Separate ports can be assigned to a vSwitch and vDS on the same ESXi host.



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

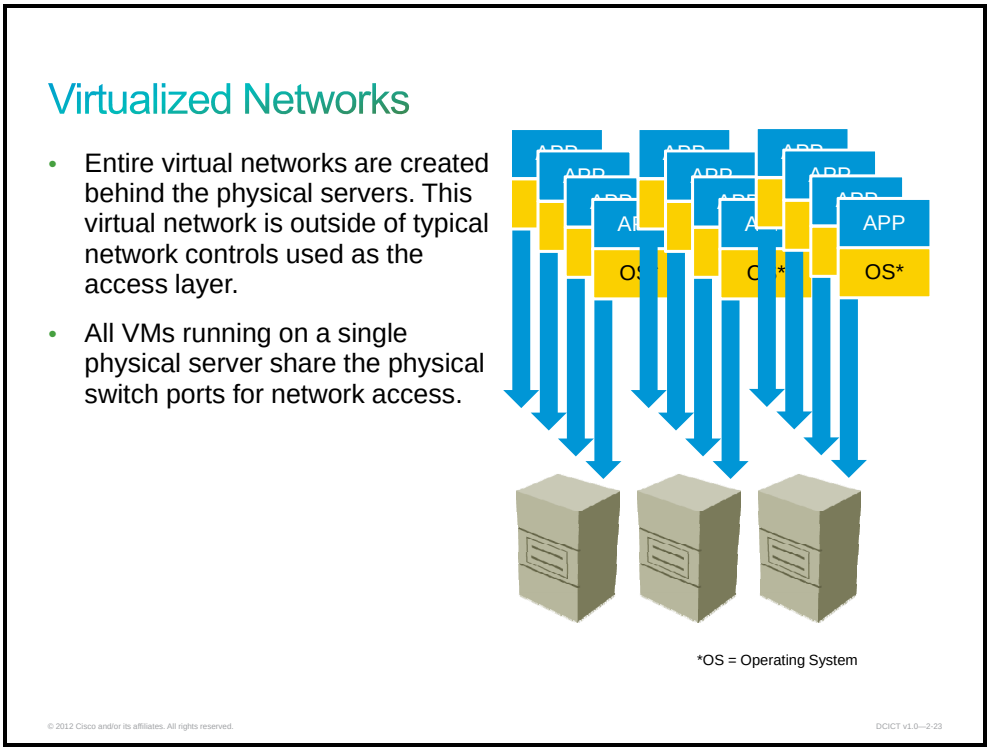
The VMware vSwitch and vDS are not mutually exclusive and can coexist within the same vCenter management environment. Physical NICs (VMNICs) may be assigned to either the vSwitch or the vDS on the same ESX or ESXi host.

You can also migrate the ESX service console and VMkernel ports from the vSwitch, where they are assigned by default during ESX installation, to the vDS. This arrangement facilitates the single point of management for all virtual networking within the vCenter data center object.

https://t.me/learningnets

How the Cisco Nexus 1000V Series Switch Brings Network Visibility to the VM Level

This topic describes the features that Cisco Nexus 1000V Series switches offer to VMs.



Networking is another resource that is being virtualized.

Having more than two physical network adapters in a host is seldom possible. VMware developed the vSwitch to help overcome this limitation. Providing each VM a virtual connection through a vSwitch achieves connectivity to the rest of the network.

VMs can have one or more vNICs. Those vNICs can be configured to attach to a vSwitch that uses one or more physical NICs in a host to act as uplinks to the external network.

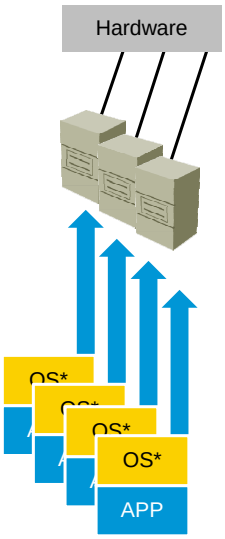
Essentially, VMs behave the same as if they were connected to a normal network switch.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Physical Switches in Virtualized Environments

- Multiple VMs running on a server connect through the same physical ports.
- Physical port configuration affects all attached VMs instead of a single server.
- Port configuration is static and not transportable with the VM.



*OS = Operating System

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

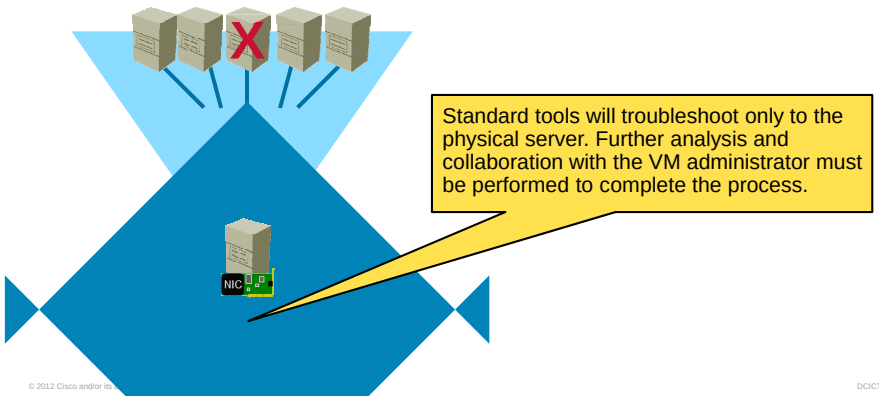
Some challenges are associated with virtual networking:

- Every VM requires some amount of network bandwidth.
- Most configurations on a physical network switch, which connects to a host, will affect all the VMs on that particular host.
- A VM that moves between hosts can be affected by a different configuration on a network switch port that is connected to the target host.

https://t.me/learningnets

Virtual Server Environment

- When troubleshooting network issues, the network administration team has no visibility into the virtual-server level. The internal vSwitch and underlying VMs are hidden from the physical network
- A compromised virtual server will be traced only as far as the physical port and further diagnosis will need to occur.



© 2012 Cisco and/or its

Special challenges need to be considered in VM environments.

VM administrators do not like to give too much administrative access to too many people. However, VM administrators historically come from the data center server side and do not necessarily have much knowledge about network configuration.

As a result, network administrators are not always given sufficient access to the VM networking infrastructure. At the same time, they are expected to maintain different service level agreements (SLAs) on that infrastructure.

Moreover, it can be quite a challenge for a network administrator to contain any problematic servers in a VM environment. In a nonvirtualized network, the network administrator always has an action of last resort: shutting down the server port. In a VM environment, that action would mean shutting down network access for all the VMs, which might have a greater impact on production than doing nothing at all.

Also, troubleshooting might return indeterminate results that simply indicate a problem somewhere within the host but do not specify on which VM the problem originates.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Network Visibility and Configuration Comparison

Feature	Physical Network	Virtual Network
Network visibility	Individual server	Physical server
Port configuration	Individual server	Physical server
Network configuration	Network administrator	VM and network administrators
Security policies	Individual server	Physical server

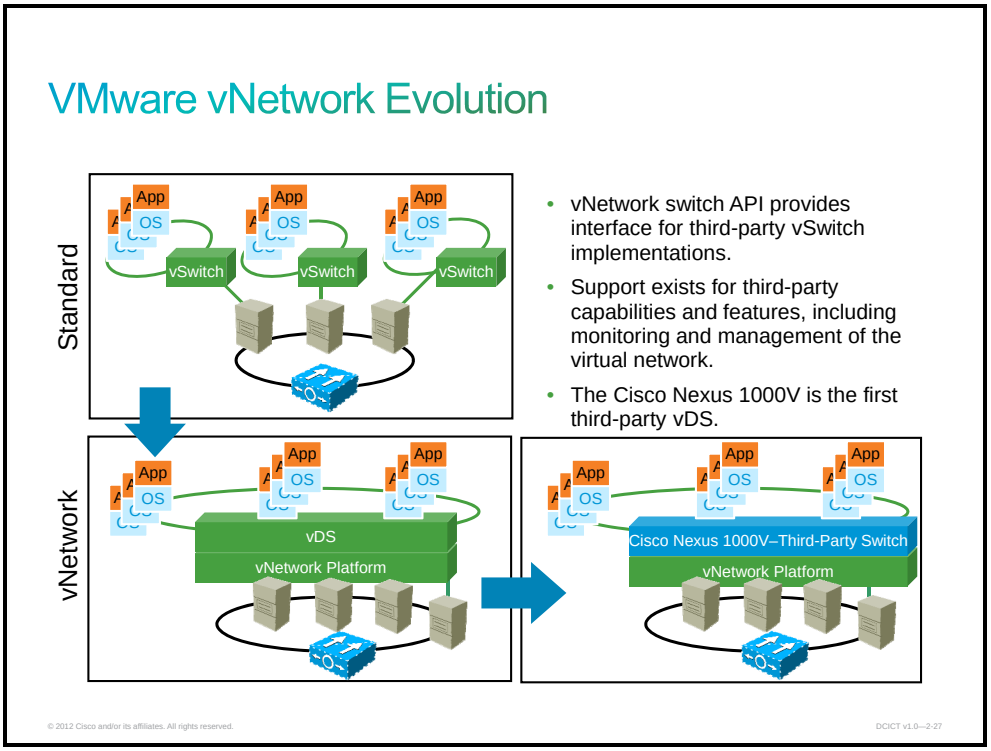
© 2012 Cisco and/or its affiliates. All rights reserved.

Because of differences in VM networking, a different mindset is required.

Certain tasks that were historically in the domain of network administrators can move to the administrators of the VM environment.

Doing so may affect internal SLAs within the company and the responsibilities of IT departments.

https://t.me/learningnets



The Cisco server virtualization solution uses a technology that Cisco and VMware developed jointly. The network access layer is moved into the virtual environment to provide enhanced network functionality at the VM level.

This feature can be deployed as a hardware- or software-based solution, depending on the data center design and demands. Both deployment scenarios offer VM visibility, policy-based VM connectivity, policy mobility, and a nondisruptive operational model.

Cisco Nexus 1000V Series Switches

The Cisco Nexus 1000V Series switch is a software-based solution that provides VM-level network configurability and management. The Cisco Nexus 1000V Series switch works with any upstream switching system to provide standard networking controls to the virtual environment.

VN-Link

Cisco and VMware jointly developed VN-Link technology, which has been proposed to the IEEE for standardization. The technology is designed to move the network access layer into the virtual environment, to provide enhanced network functionality at the VM level.

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

vSwitch Options with vSphere

vSwitch	Model	Details
vNetwork Standard Switch	Host based: 1 or more per ESX host	• Same as vSwitch in ESX 3.5
vDS	Distributed: 1 or more per data center	• Expanded feature set - PVLANs - Bi-directional traffic shaping - Network vMotion • Simplified management
Cisco Nexus 1000V Series switch	Distributed: 1 or more per data center	• Cisco IOS fashion CLI • Same remote management as Cisco or Nexus physical switches • Different feature set compared to Cisco or Nexus physical switches

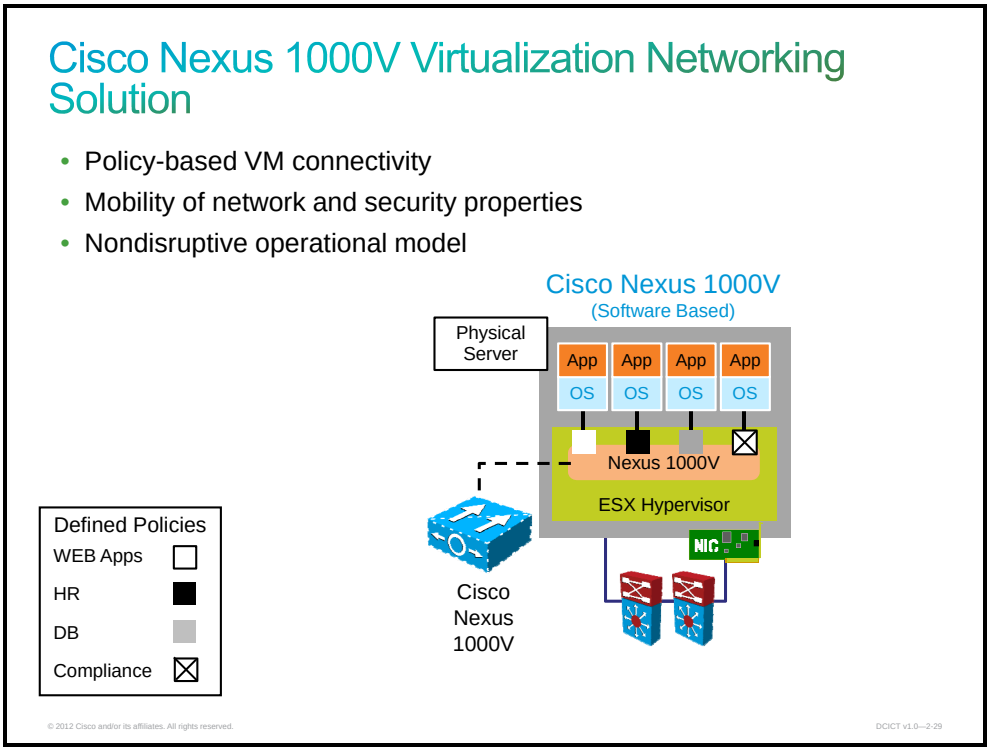
© 2012 Cisco and/or its affiliates. All rights reserved.

With vSphere 4, VMware customers can now enjoy the benefits of three virtual networking solutions: vSwitch, vDS, and the Cisco Nexus 1000V Series switch.

The Cisco Nexus 1000V Series switch bypasses the vSwitch by using a Cisco software switch. This model provides a single point of configuration for the networking environment of multiple ESX or ESXi hosts. Additional functionality includes policy-based connectivity for the VMs, network security mobility, and a nondisruptive software model.

VM connection policies are defined in the network and applied to individual VMs from within vCenter. These policies are linked to the universally unique identifier (UUID) of the VM and are not based on physical or virtual ports.

https://t.me/learningnets



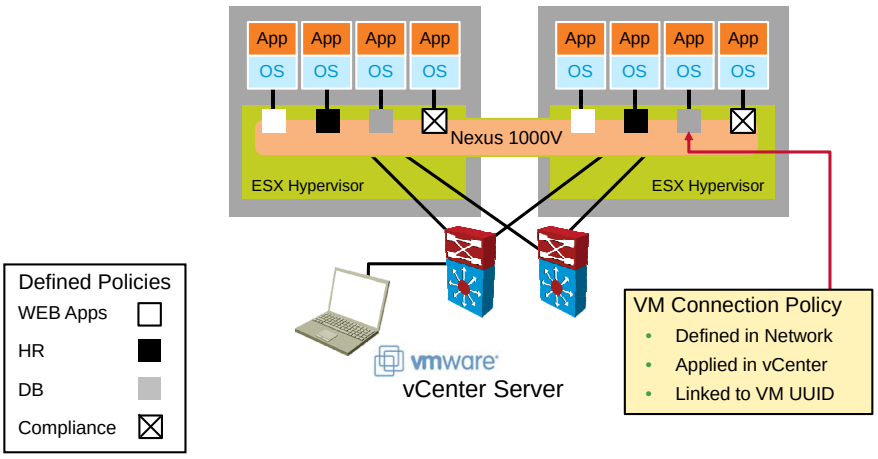
The Cisco Nexus 1000V Series switch provides the following important features:

- **Policy-based VM connectivity:** The network administrator can now apply a policy to the VM level.
- **Mobility:** When a VM is migrated from one host to another, the network and security policies can follow seamlessly, without network reconfiguration.
- **Nondisruptive:** The model is nondisruptive because the vDS is already in place. The VMware administrator continues normal tasks, but without needing to prepare the network-configuration side of the connection.

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

Cisco Nexus 1000V Distributed Virtual Switch

Policy-based VM connectivity



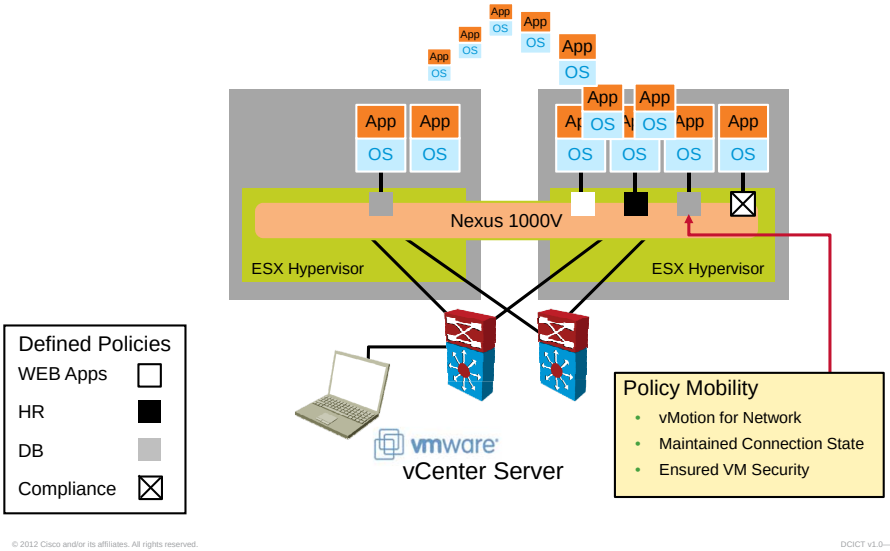
© 2012 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 1000V Series switch provides for policy-based connectivity of VMs. This policy is defined and applied by network administrators rather than by VMware administrators. This feature allows network administrators to regain control of their responsibilities and provide support at the level of the individual VM.

The figure illustrates how a policy is defined and pushed down to vCenter. The VMware administrator can then apply that policy to the VM on creation or modification.

Cisco Nexus 1000V Distributed Virtual Switch (Cont.)

Mobility of network and security properties



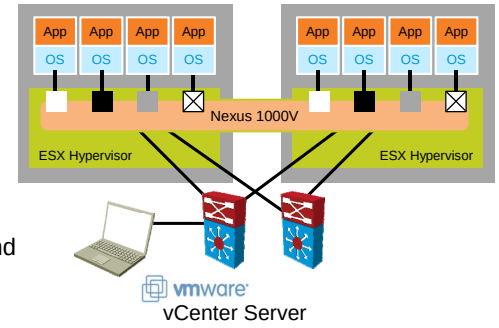
© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

All policies that are applied and defined fully support VMware mobility capabilities such as vMotion or high availability. Policies remain applied to the VM even as it moves from one host to another.

Cisco Nexus 1000V Distributed Virtual Switch (Cont.)

- Nondisruptive operational model
- Server benefits
  - Existing VM management preserved
  - Reduced deployment time and operational workload
  - Improved scalability
  - VM-level visibility
- Network benefits
  - Unified network management and operations
  - Improved operational security
  - Enhanced VM network features
  - Policy persistence
  - VM-level visibility



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

Cisco Nexus 1000V Series Switches can be introduced to an existing virtual environment. With proper planning and deployment, the migration from VMware native virtual networking can be nondisruptive for running services.

VMware administrators reduce their workload by relinquishing control of virtual networking to network administrators.

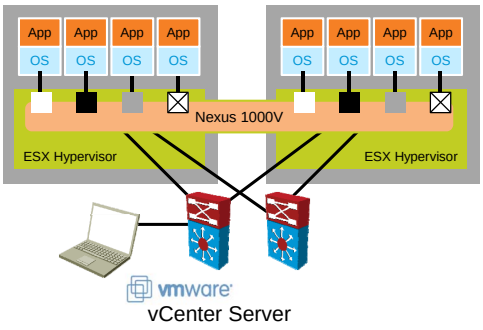
Cisco Nexus 1000V Series Switches introduce unified management with the rest of the IP network by using the same familiar techniques and commands that are available on other network platforms.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 1000V Distributed Virtual Switch Features

- Layer 2
  - VLAN, PVLAN, 802.1q
  - Link Aggregation Control Protocol (LACP)
  - vPC host mode
- QoS classification and marking
- Security
  - Layer 2, 3, 4 access lists
  - Port security
- SPAN and ERSPAN
- NetFlow
- Compatibility with VMware
  - vMotion, Storage vMotion
  - Distributed Resource Scheduler (DRS) , high availability (HA), FT



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

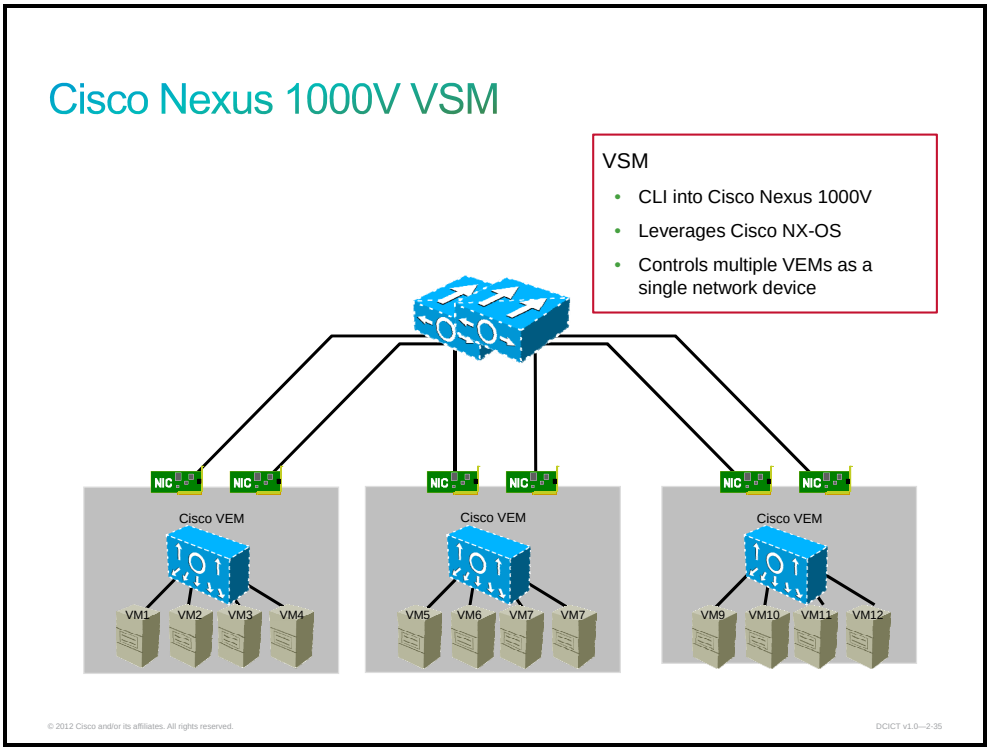
Cisco Nexus 1000V Series Switches also introduce additional functionality to virtual network connectivity that standard vSwitches or vDSs do not have or have to a lesser extent:

- More detailed VLAN and PVLAN control is available.
- QoS capabilities are greatly enhanced.
- Port-security capabilities are introduced.
- Access control list (ACL) functionalities are available.
- Monitoring tools such as Switched Port Analyzer (SPAN), Encapsulated Remote SPAN (ERSPAN), and Wireshark are a part of the Cisco Nexus 1000V Series switch.
- NetFlow provides visibility into how, when, and where network traffic is flowing.

https://t.me/learningnets

How the VSM and VEM Integrate with VMware ESX or ESXi and vCenter

This topic describes how the VSM and VEM integrate with ESX (or ESXi) and vCenter.



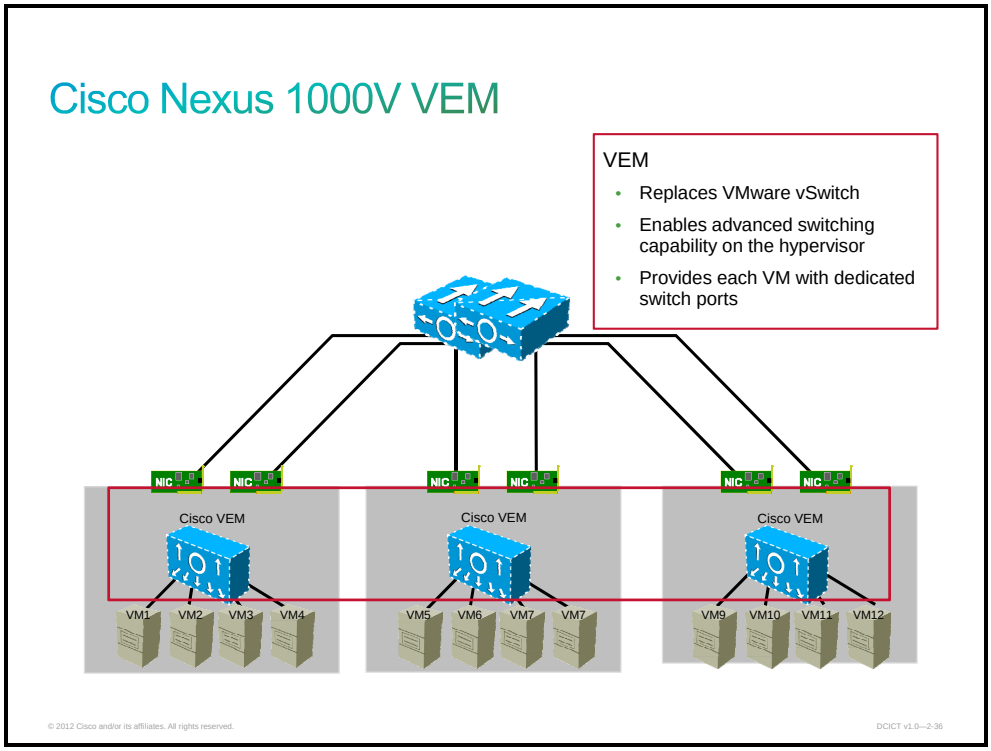
The VSM is a virtual equivalent of the supervisor modules that can be found in other Cisco Nexus Operating System (NX-OS) devices. The VSM provides the platform on which Cisco NX-OS runs and that interacts with other components that are part of the Cisco Nexus 1000V Series switch.

Management access to the VSM is provided through a Cisco NX-OS CLI. This CLI has the same syntax and behavior as the CLI on other Cisco Nexus devices.

All the line cards, such as VEMs, that connect to the VSM behave as a single network device.

The VSM can reside on a VM or a Cisco Nexus 1010 appliance.

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



The VEM is the virtual equivalent of a line card of a standard switch. The VEM resides on every VMware host on the hypervisor layer. A VEM provides connectivity among the VMs and between the VMs and the outside network, through the physical NIC ports in the host. Multiple VEMs that communicate with one VSM or a set of VSMs correspond to one logical switch.

VEMs on different hosts do not have a direct line of communication to one another. Rather, they require an outside switch to link them. The VEM-to-VSM communication path carries only control traffic.

https://t.me/learningnets

Cisco Nexus 1000V Virtual Chassis

- Virtual chassis is a logical representation of Cisco Nexus 1000V
- Contains up to:
  - 2 VSMs
  - 64 VEMs
  - Virtual service blades

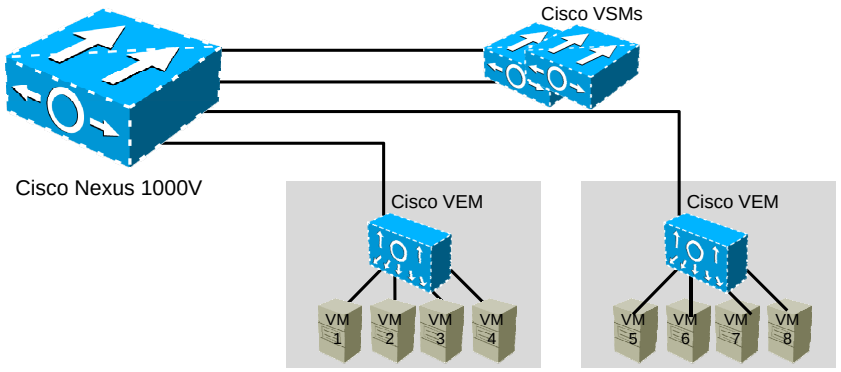


© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco Nexus 1000V Series virtual chassis is an expression that encompasses the Cisco Nexus 1000V Series components such as VSMs and VEMs.

Cisco Nexus 1000V Virtual Chassis (Cont.)

```
vsm# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    0      Virtual Supervisor Module        Cisco Nexus1000V   active *
2    0      Virtual Supervisor Module        Cisco Nexus1000V   ha-standby
3    248    Virtual Ethernet Module          NA                 ok
```

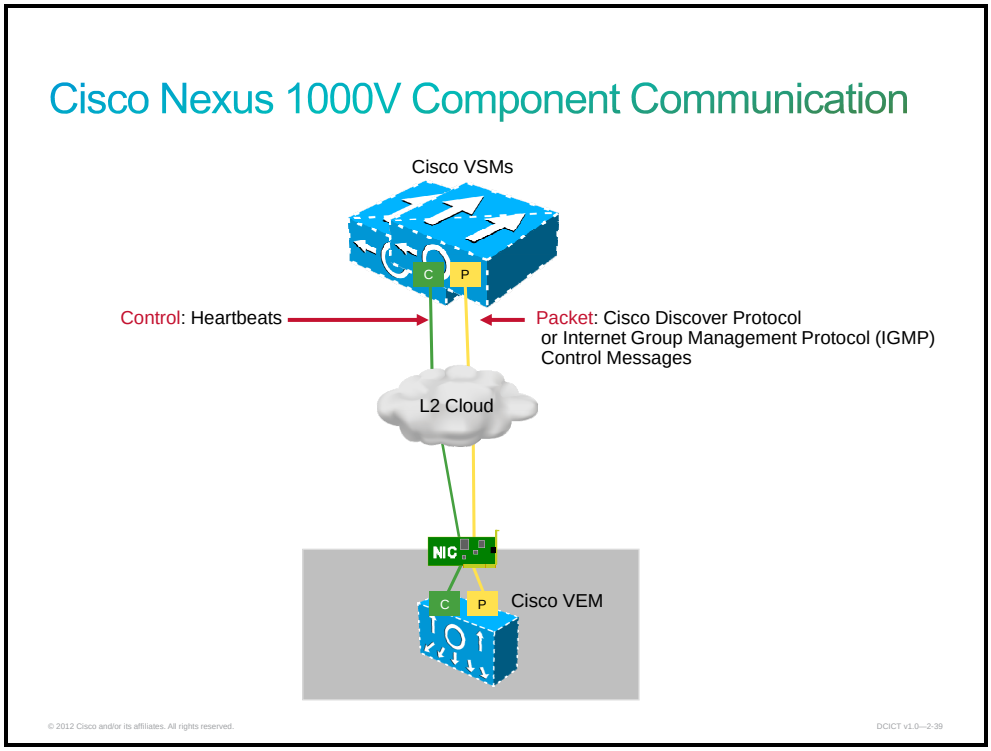


© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco Nexus 1000V Series virtual chassis behaves as if it were a physical device with multiple line cards. For example, the **show module** command in the Cisco Nexus 1000V CLI displays the VSMs and VEMs in the same way that it would display supervisors and line cards on a Cisco Nexus 7000 Series switch.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets



Communication between the VSM and VEM is provided through two distinct virtual interfaces: the control and packet interfaces.

The control interface carries low-level messages to each VEM, to ensure proper configuration of the VEM. A 2-second heartbeat is sent between the VSM and the VEM, with a 6-second timeout. The control interface maintains synchronization between primary and secondary VSMs. The control interface is like the Ethernet out-of-bound channel in switches such as the Cisco Nexus 7000 Series switch.

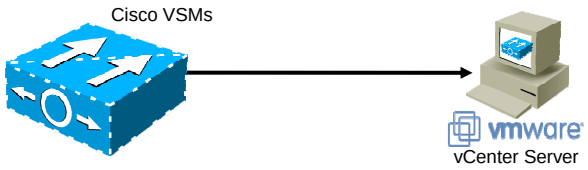
The packet interface carries network packets, such as Cisco Discovery Protocol or Internet Group Management Protocol (IGMP) control messages, from the VEM to the VSM.

You should use one or two separate VLANs for the control interface and for the packet interface.

Being VLAN interfaces, the control and packet interfaces require Layer 2 connectivity.

https://t.me/learningnets

Cisco Nexus 1000V Component: vCenter Communication



- Communication using the VMware VIM API over SSL
- Connection configured manually on the VSM or through the installer application
- Requires installation of vCenter plug-in (downloaded from VSM)
- After establishment, the Cisco Nexus 1000V is created in vCenter

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

Communication between the VSM and vCenter is provided through the VMware Virtual Infrastructure Methodology (VIM) application programming interface (API) over Secure Sockets Layer (SSL). The connection is set up on the VSM and requires installation of a vCenter plug-in, which is downloaded from the VSM.

After communication between the two devices is established, the Cisco Nexus 1000V vDS is created in vCenter.

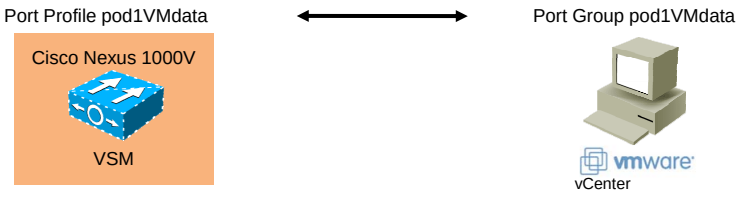
This interface is known as the out-of-band (OOB) management interface. Although not required, best practice is to have this interface, vCenter, and host management in the same VLAN.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Port Profiles

- Port profiles are the Cisco Nexus 1000V variant of port groups from VMware.
- Although you can configure individual ports on the Cisco Nexus 1000V, doing so through port profiles is preferred.
- The next figure shows an example of a port profile configuration.



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

Port profiles are used to configure interfaces in the Cisco Nexus 1000V Series switch with a common set of configuration commands. A port profile can be assigned to multiple interfaces. Any changes to a port profile are automatically propagated across all interfaces that are associated with that port profile.

In vCenter Server, the port profile is represented as a port group. Both virtual and physical interfaces are assigned in vCenter Server to a port profile and perform these functions:

- Define port configuration by policy
- Apply a single policy across a large number of ports
- Support both virtual Ethernet and Ethernet ports

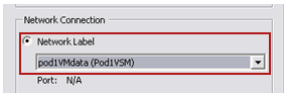
Note	Any manual configuration of an interface overrides the port profile configuration. Manual configuration is not recommended for general use but rather for tasks such as quick testing of a change.

https://t.me/learningnets

Port Profiles and Port Groups

Port profiles correspond to port groups within VMware. By default, the port group created within VMware for each port profile has the same name. VMware administrators use the port group to assign network settings to VMs and uplink ports.

```
N1000v-VSM(config)# port-profile pod1VMdata
N1000v-VSM(config-port-prof)# switchport mode access
N1000v-VSM(config-port-prof)# switchport access vlan 102
N1000v-VSM(config-port-prof)# vmware port-group pod1VMdata
N1000v-VSM(config-port-prof)# no shut
N1000v-VSM(config-port-prof)# state enabled
N1000v-VSM(config-port-prof)# vmware max-ports 12
```





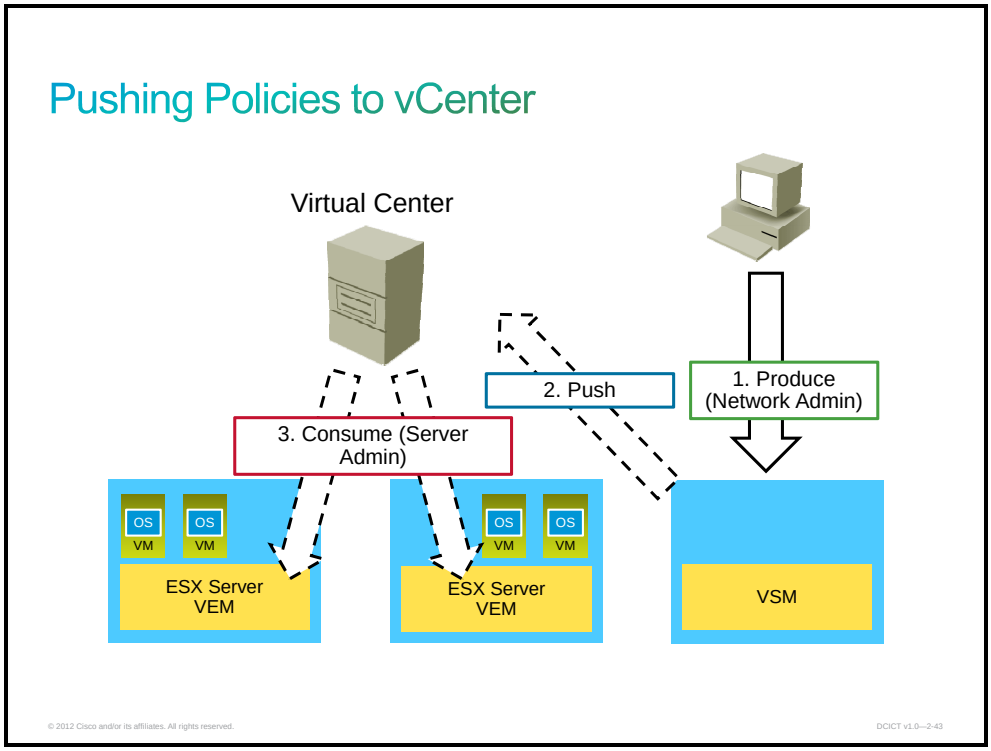
© 2012 Cisco and/or its affiliates. All rights reserved.

When a port profile is created and enabled, a corresponding port group is created in vCenter. By default, this port group has the same name as the profile, but this name is configurable. VMware administrators use the port profile to assign network settings to VMs and uplink ports. When a VMware ESX or ESXi host port (VMNIC) is added to a vDS that the Cisco Nexus 1000V Series switch controls, an available uplink port group is assigned and those settings are applied. When a NIC is added to a VM, an available VM port group is assigned. The network settings that are associated with that profile are inherited.

A NIC in VMware is represented by an interface that is called a VMNIC. The VMNIC number is allocated during VMware installation.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



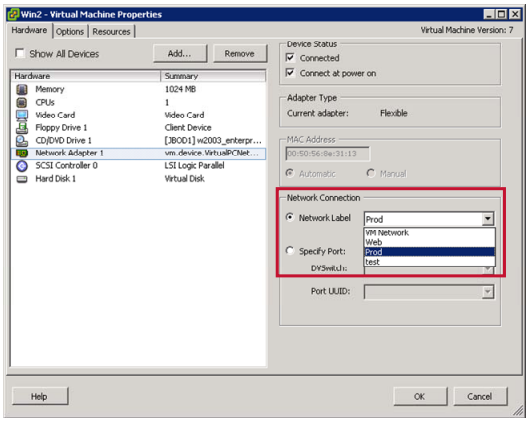
Cisco Nexus 1000V Series Switches provide an ideal model in which network administrators define network policies that virtualization or server administrators can use as new VMs are created. Policies that are defined on the Cisco Nexus 1000V Series switch are exported to vCenter and assigned by the server administrator as new VMs require access to a specific network policy. This concept is implemented on the Cisco Nexus 1000V Series switch by using a feature called port profiles. The Cisco Nexus 1000V Series switch with the port profile feature eliminates the requirement for the server administrator to create or maintain a vSwitch and port group configurations on any of their ESX or ESXi hosts.

Port profiles separate network and server administration. For network administrators, the Cisco Nexus 1000V feature set and the ability to define a port profile by using the same syntax as for existing physical Cisco switches help to ensure consistent policy enforcement without the burden of managing individual switch ports. The Cisco Nexus 1000V solution also provides a consistent network management, diagnostic, and troubleshooting interface to the network operations team, allowing the virtual network infrastructure to be managed like the physical infrastructure.

https://t.me/learningnets

VMware Administrator View

Consistent workflow: Continue to select port groups when configuring a VM in VMware virtual infrastructure client.



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

For consistent workflow, continue to choose Port Groups when configuring a VM in vCenter.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Standard VMware Admin Roles

Task	VMware Admin	Network Admin
vSwitch config	Per ESX Host	N/A
Port Group Config	Per ESX Host	N/A
Add ESX Host	vCenter-Based	N/A
NIC Teaming Config	Per vSwitch	N/A
Virtual Machine Creation	vCenter-Based	N/A
Security	N/A	N/A
Visibility	vCenter	N/A
Management	vCenter	N/A

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

The figure shows typical administrative tasks before the introduction of vDS or Cisco Nexus 1000V Series switch. There is no participation by the network administrator.

Cisco Nexus 1000V Admin Roles

Task	VMware Admin	Network Admin
vSwitch Config	Automated	Same as Physical Network
Port Group Config	Automated	Policy-Based
Add ESX Host	Unchanged–vCenter-Based	N/A
NIC Teaming Config	Automated	Port Channel Optimized
Virtual Machine Creation	Unchanged–vCenter-Based	N/A
Security	Policy-Based	ACL, PVLAN, Port Security, TrustSec
Visibility	VM-Specific	VM-Specific
Management	Unchanged–vCenter-Based	Cisco CLI, XML API, SNMP, DCNM

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

The figure redefines the administrative tasks with the introduction of the vDS and Cisco Nexus 1000V. Some tasks are unchanged, whereas others are now the responsibility of the network team.

https://t.me/learningnets

Summary

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

Summary

- To provide connectivity for VMs in the virtual server environment, vendors such as VMware have developed a software-based switch called a vSwitch, which resides in the physical server host. VMware administrators create, modify, and manage this vSwitch, which is outside of the control of the network team in the data center.
- To enhance the capabilities of the vSwitch, VMware developed the vDS. This provided more functionality, but was still managed by the VMware administrators, through the vCenter Server GUI.
- To provide better visibility to the network team for the growing virtual server access layer, Cisco worked with VMware to develop Cisco's first software-based switch, the Cisco Nexus 1000V switch. This switch is managed by the network team, providing the capability of pushing network-based policies down to the virtual server layer.
- The Cisco Nexus 1000V switch comprises a VSM and a VEM. These modules create the vSwitch. This vSwitch integrates into the VMware environment and has connectivity to the vCenter Server. Policies that the network team creates can be pushed down to vCenter for applying to the virtual servers that the VMware administrators manage.

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

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 5

Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch

Overview

The purpose of this lesson is to identify how to use the VMware ESX and Cisco Nexus 1000V CLIs to validate connectivity of the Cisco Nexus 1000V Virtual Supervisor Module (VSM) to Virtual Ethernet Modules (VEMs) and VMware vCenter.

Objectives

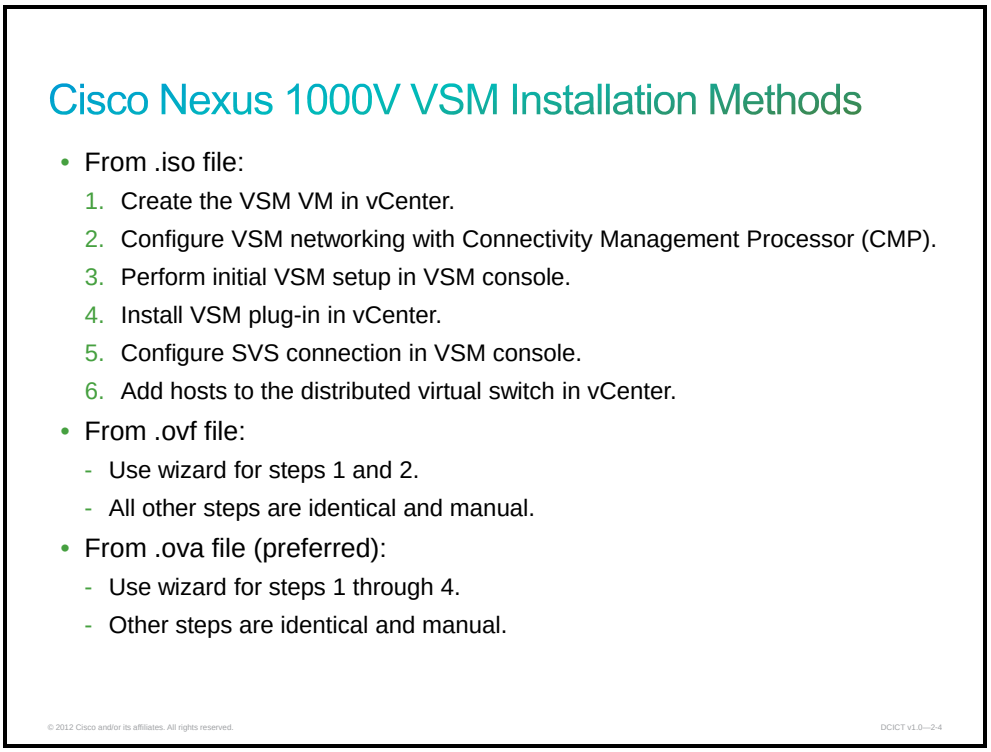
Upon completing this lesson, you will be able to verify the initial setup and operation for Cisco Nexus 1000V Series Switches. You will be able to meet these objectives:

- Identify the commands that are used to verify the initial configuration and module status on the Cisco Nexus 1000V Series switch
- Identify how to verify VEM status on the VMware ESX or ESXi host
- Identify how to validate VM port groups

https://t.me/learningnets

Verifying the Initial Configuration and Module Status on the Cisco Nexus 1000V Series Switch

This topic describes how to verify the initial configuration of the Cisco Nexus 1000V Series switch and confirm that the VSM and VEMs are available.



There are several ways to install and deploy the VSM. The preferred method is to use an Open Virtualization Appliance (OVA) file. This method provides the highest degree of guidance and error-checking for the user.

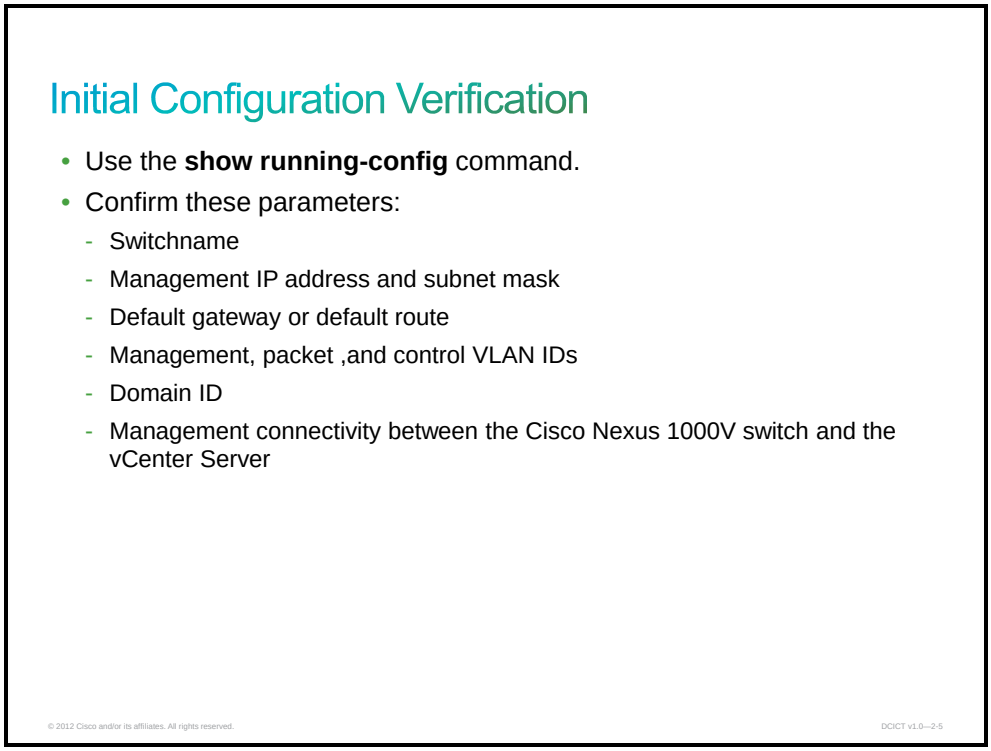
All other methods are less streamlined and require the administrator to be knowledgeable. However, these other methods work well in certain situations.

Open Virtualization Format (OVF) files are standardized file structures that are used to deploy virtual machines (VMs). You can create and manage OVF files by using the VMware OVF Tool.

OVA files are similar to OVF files.

© 2012 Cisco Systems, Inc.

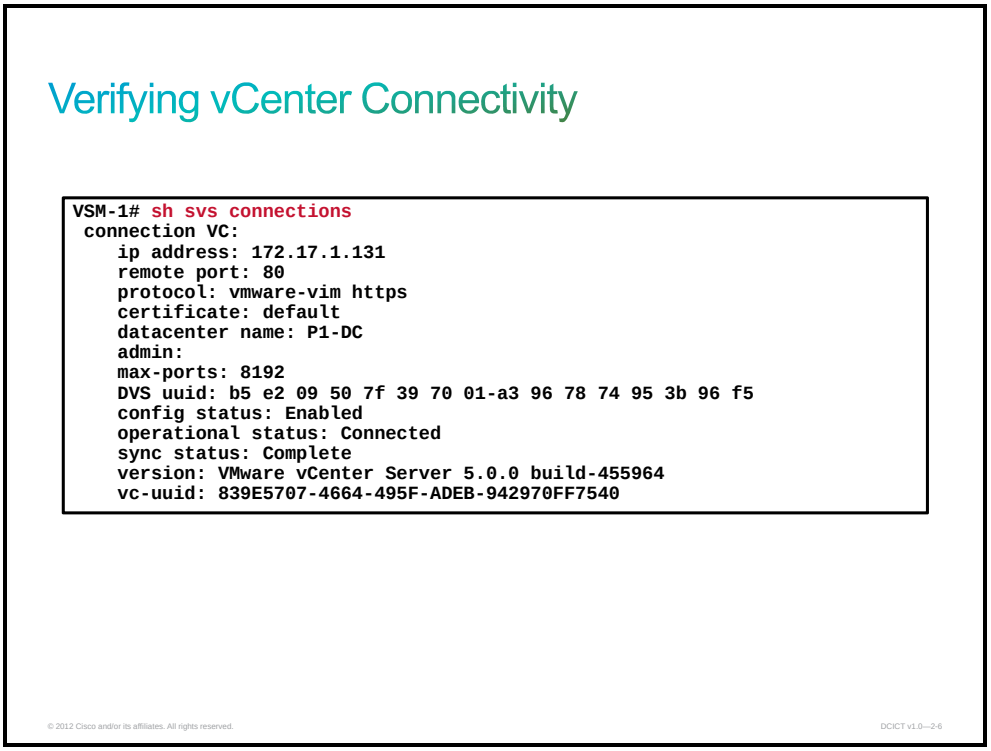
https://t.me/learningnets



After the Cisco Nexus 1000V Series switch has been installed on a VM or Cisco Nexus 1010 appliance, the initial configuration dialog box is displayed. The network administrator performs the initial configuration to provide a basic configuration for the Cisco Nexus 1000V Series Switch. All further configurations, such as port profiles, are configured at the CLI in configuration mode.

To verify the initial configuration and subsequent modifications to the configuration, use the **show running-config** command at the Cisco Nexus 1000V CLI.

https://t.me/learningnets



To establish a connection between the Cisco Nexus 1000V Series switch and vCenter, the network administrator configures a software virtual switch (SVS) connection. This connection must be in place for the Cisco Nexus 1000V Series switch to push configuration parameters such as port profiles to vCenter.

To verify that the SVS connection is in place, use the **show svs connections** command.

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

Verifying vCenter Connectivity (Cont.)

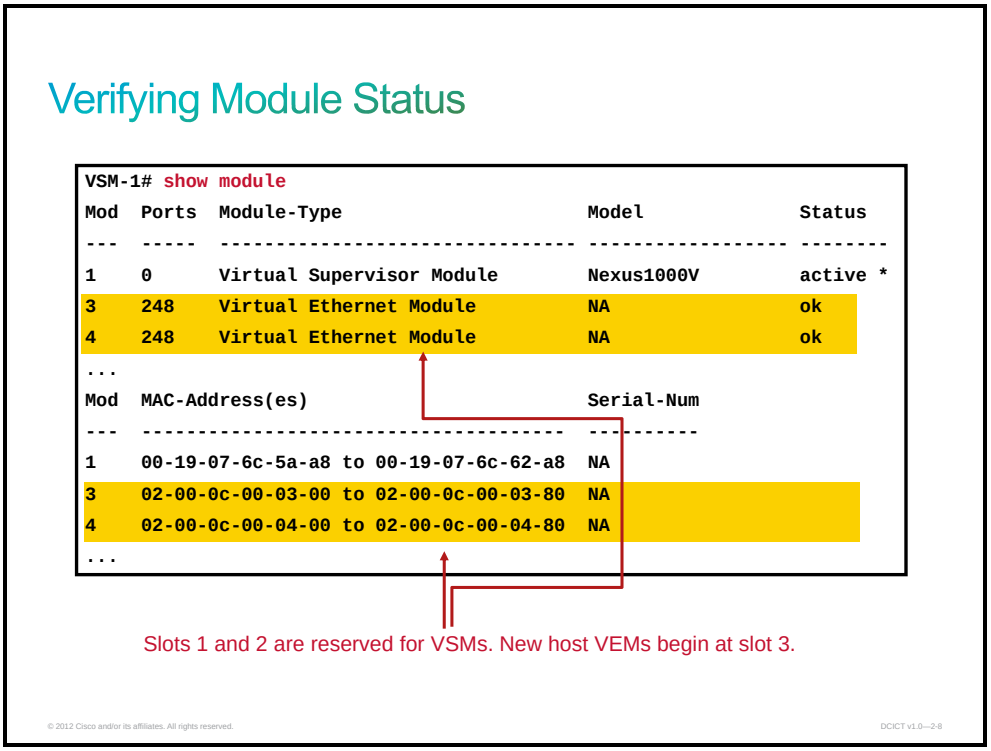
```
VSM-1# sho svs domain
SVS domain config:
  Domain id:    1
  Control vlan: 110
  Packet vlan:  110
  L2/L3 Control mode: L2
  L3 control interface:  NA
  Status: Config push to VC successful
```

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

You can perform further verification of the connection between the Cisco Nexus 1000V Series switch and vCenter, use the **show svs domain** command. Each Cisco Nexus 1000V Series switch uses one domain ID. All ESX or ESXi hosts that have a VEM installed listen to updates from one domain ID: the virtual chassis in which they reside. This domain ID is in all updates from vCenter to the ESX or ESXi hosts on which the VEMs are installed.

To verify the domain ID parameters, use the **show svs domain** command on the Cisco Nexus 1000V Series switch.

https://t.me/learningnets



The output of the **show module** command shows the primary supervisor in slot 1, the high-availability standby supervisor in slot 2, and the very first ESX or ESXi host that has been added to the Cisco Nexus 1000V Series switch instance in slot 3.

After a host has been added and the VEM has been successfully installed, the VEM appears as a module on the VSM CLI. This appearance is similar to modules that are added to a physical chassis.

Note	Slots 1 and 2 are reserved for VSMs. New host VEMs start from slot 3.

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

Verifying Module Status (Cont.)

The **show module vem map** command shows the status of all VEMs as well as the universally unique identifier (UUID) of the host on which the VEM runs.

```
VSM-1# show module vem map
Mod     Status          UUID
---     -----------     ----
3       powered-up      34343937-3638-3355-5630-393037415833
4       powered-up      34343937-3638-3355-5630-393037415834
```


This is the UUID of the host on which the VEM resides.

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

The **show module vem map** command shows the status of all VEMs, as well as the universally unique identifier (UUID) of the host on which the VEM runs. This command can be used to verify that the VEM is installed and tied to the UUID of the host.

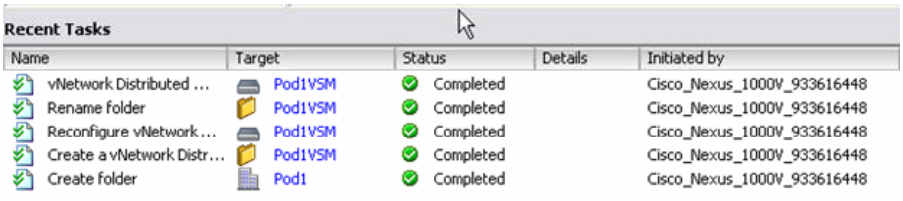
https://t.me/learningnets

Verifying the VEM Status on the ESX or ESXi Host

This topic discusses how to verify the status of the VEM on the ESX or ESXi host by using vCenter or the host CLI.

Verifying VSM Connectivity

After the VSM connects properly, you should see output that shows the creation of a vDS. The vDS also appears in the vCenter Inventory networking pane.

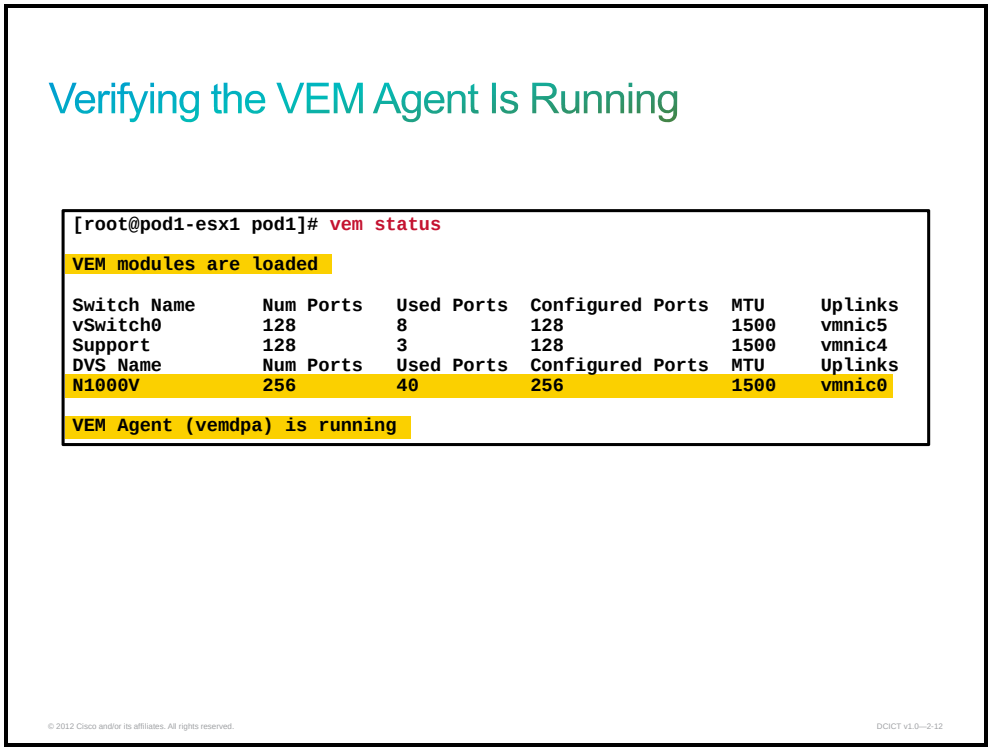


© 2012 Cisco and/or its affiliates. All rights reserved.

After the VSM connects to vCenter, you see the vNetwork Distributed Switch (vDS) appear in the vCenter Networking inventory panel. You should see the port groups that you configured for control, management, and packet traffic. (These port groups are required to provide connectivity between vCenter, the VSM, and the VEMs.) Some other port groups are created by default. One is the Unused_Or_Quarantined DVUplinks port group, which connects to physical NICs. Another is the Unused_Or_Quarantined VMData port group, which faces the VM.

 © 2012 Cisco Systems, Inc.

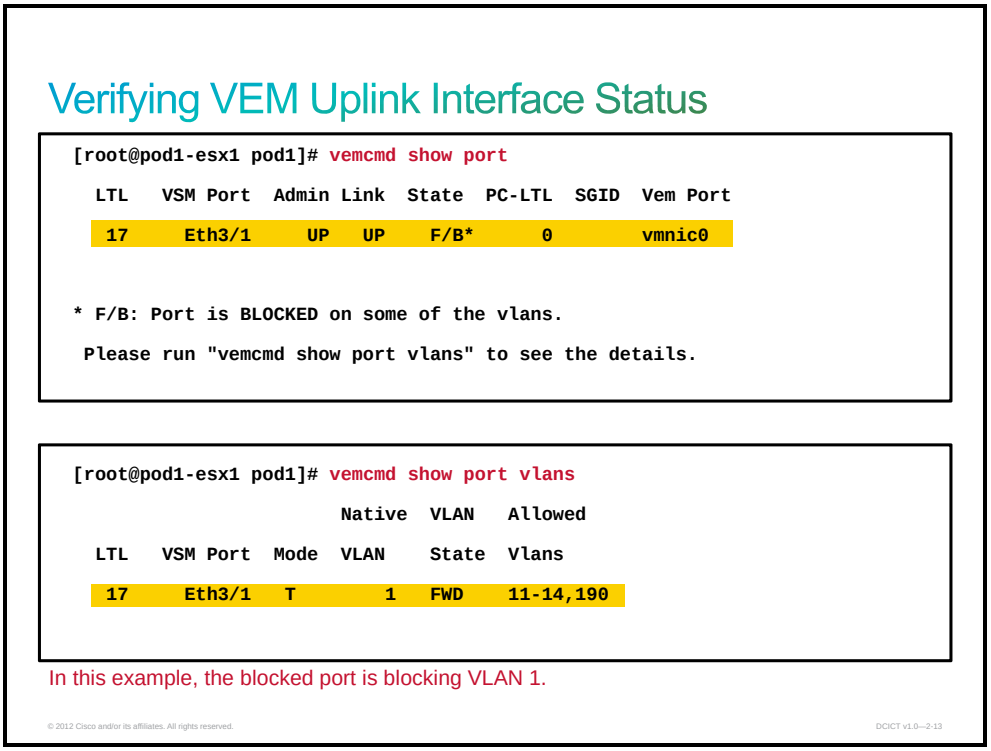
https://t.me/learningnets



You can verify the VEM status on the CLI of the ESX or ESXi host. To perform the verification, open a connection to the host and log in, using the correct credentials.

At the CLI of the host, run the **vem status** command. This command verifies that the VEM module is loaded and that the VEM Agent is running on this host. The command also confirms which interface is being used as the uplink.

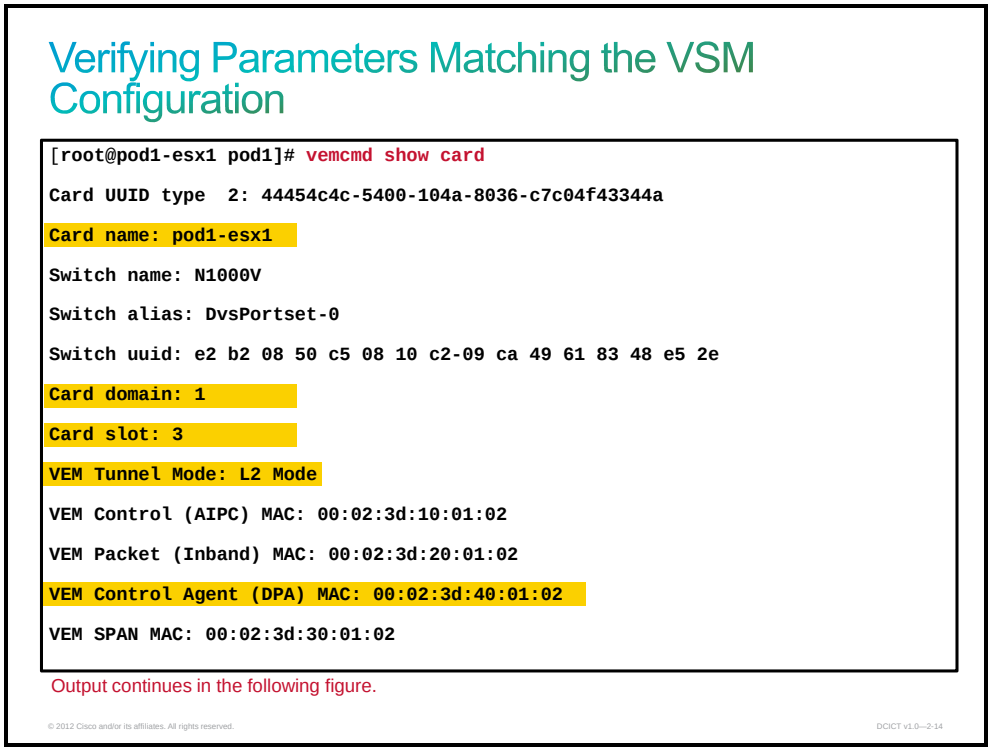
https://t.me/learningnets



The following commands can be used on the CLI of the ESX or ESXi host to provide further verification of the uplink interfaces on that host:

- **vemcmd show port:** This command verifies the VEM port that is used on the host and the Cisco Nexus 1000V Series switch. The command provides details of the port state and whether any issues need to be highlighted.
- **vemcmd show port vlans:** In the figure, the **vemcmd show port** command identified that the uplink port was blocked. You can use this command to verify which VLANs are carried across the uplink and whether any VLANs are missing. In the figure, the uplink is blocking for VLAN 1 because it is being used only at one end of the connection.

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



The following command is used to verify that the parameters on the ESX or ESXi host for the VEM match the configuration on the Cisco Nexus 1000V Series switch:

- **vemcmd show card:** Use this command to verify these components:
  - Card name
  - Card domain ID
  - Card slot
  - Control traffic connectivity mode between the VSM and VEM
  - VEM control agent MAC ID
  - Used VSM MAC addresses
  - Used control and packet VLANs

Note	The next two figures show some of the output of this command.

https://t.me/learningnets

Verifying Parameters Matching the VSM Configuration (Cont.)

```
Primary VSM MAC : 00:50:56:88:00:01
Primary VSM PKT MAC : 00:50:56:88:00:03
Primary VSM MGMT MAC : 00:50:56:88:00:02
Standby VSM CTRL MAC : 00:50:56:88:00:04
Management IPv4 address: 172.16.10.11
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
Secondary VSM MAC : 00:00:00:00:00:00
Secondary L3 Control IPv4 address: 0.0.0.0
Upgrade : Default
Max physical ports: 32
Max virtual ports: 216
```

Output continues in the following figure.

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

The output in the figure is a continuation of the output from the previous figure.

Verifying Parameters Matching the VSM Configuration (Cont.)

```
Card control VLAN: 12
Card packet VLAN: 13
Card Headless Mode : No
       Processors: 4
  Processor Cores: 4
Processor Sockets: 1
  Kernel Memory:   3885056
Port link-up delay: 5s
Heartbeat Set: True
PC LB Algo: source-mac
```

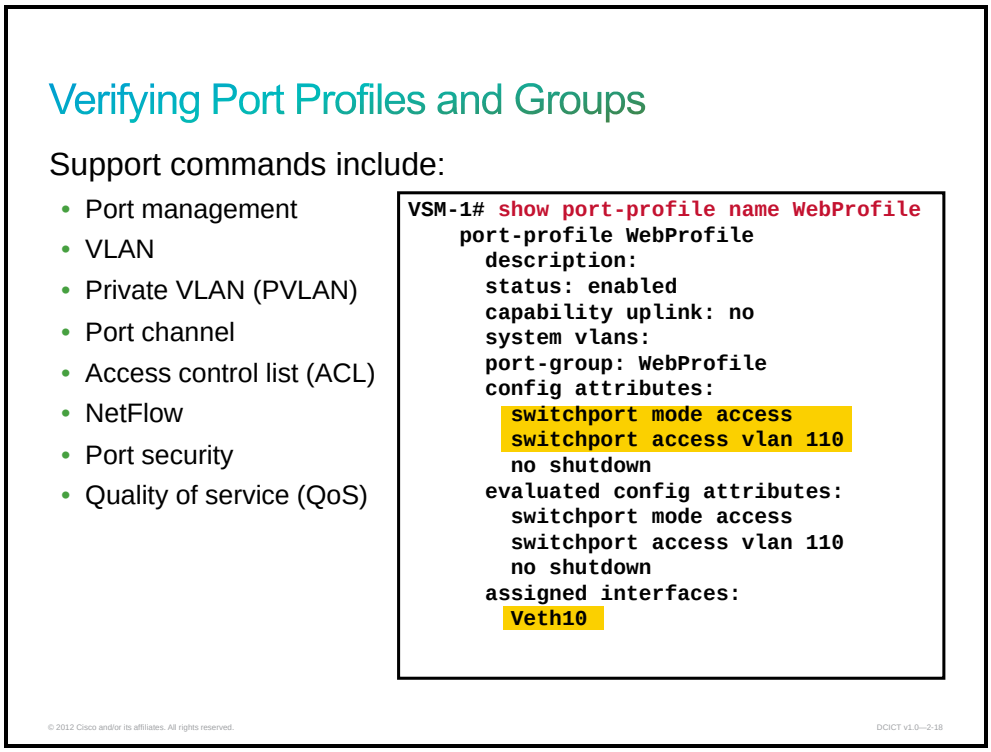
© 2012 Cisco and/or its affiliates. All rights reserved.

The output in the figure is a continuation of the output from the previous two figures.

https://t.me/learningnets

Validating VM Port Groups

This topic discusses how to validate that VMs are using the correct vDS port groups, by using vCenter and the Cisco Nexus 1000V CLI.

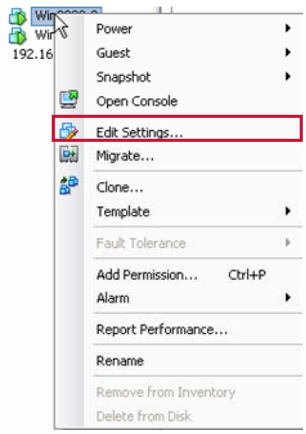


Use the **show port-profile name** *name* command to verify the profile configuration and parameters. From this command, you can check which switchport mode this port profile is using. You can also verify which VLANs are being used and which interfaces are assigned to this port profile.

https://t.me/learningnets

Verifying Port Profiles and Groups (Cont.)

- Right-click the desired VM.
- Choose **Edit Settings.**



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

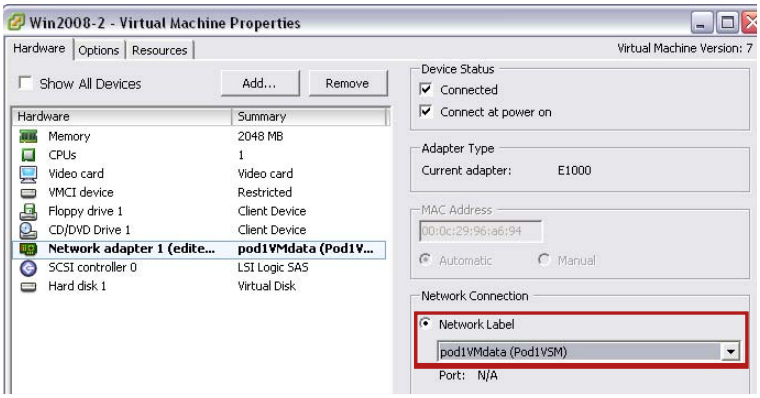
To add a VM to a VSM port group, right-click the VM and choose **Edit Settings**.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying Port Profiles and Groups (Cont.)

- The pod1VMdata port profile has been created and enabled, and now becomes available in the port group configuration of VMware vCenter.
- The VMware administrator needs to assign proper port profiles to VMs to achieve desired connectivity.



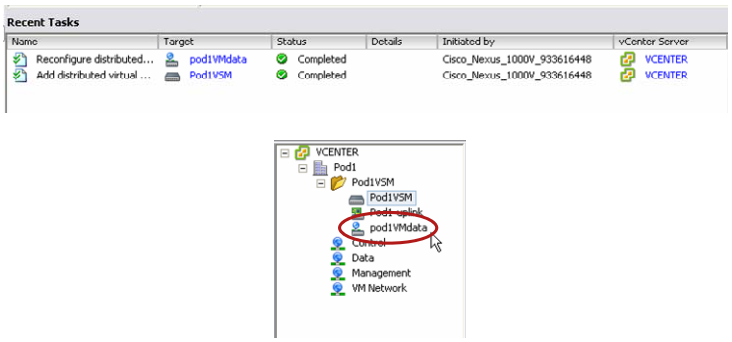
© 2012 Cisco and/or its affiliates. All rights reserved.

After the port profile has been created on the Cisco Nexus 1000V Series switch and pushed down to vCenter, it is available for the VMware administrator to use when creating or modifying VMs.

The figure shows that a port profile pod1VMdata has been created on the Cisco Nexus 1000V Series switch. As the figure shows, the VMware administrator is modifying the VM properties. In the Network Connection section, the administrator has chosen the pod1VMdata port profile as a port group. This port group can be applied to the VM.

https://t.me/learningnets

Verifying Port Profiles and Groups (Cont.)



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

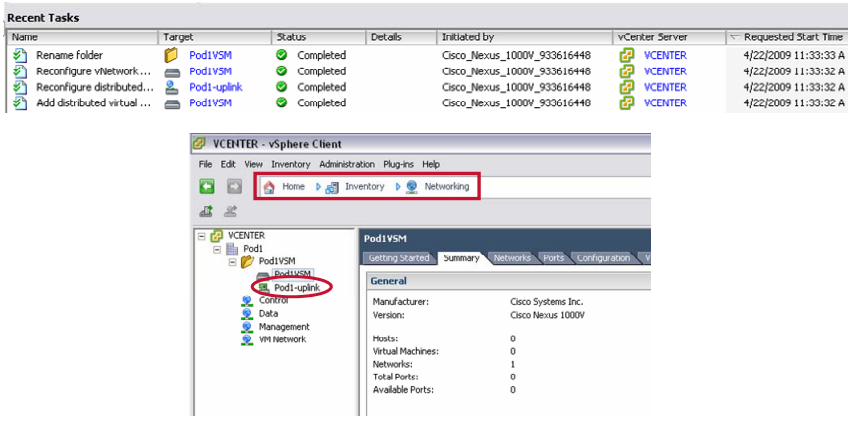
The figure describes the process of verifying the VMware data port profile configuration. From within VMware vSphere, choose **Inventory > Networking** within the navigation pane. The network inventory objects, including the newly created port profile pod1VMdata, appear. The vSphere Recent Tasks window shows that the creation of the new port profile has been completed.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Verifying Port Profiles and Groups (Cont.)

For external connectivity, an uplink port profile is created.



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

Uplink port profiles are created on the Cisco Nexus 1000V Series switch and are pushed to vCenter so that the virtual switch can provide external connectivity for the VMs that reside on that host.

Use the same verification method that you used for VM port profiles to verify that the port profile is available for the VMware administrator to use.

https://t.me/learningnets

Summary

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

Summary

- The status of the VSM and VEMs in the Cisco Nexus 1000V switch virtual chassis can be verified at the CLI on the switch.
- To verify the VSM creation and the VEM installation on the ESX or ESXi host, check the vCenter GUI and the CLI of the host.
- When creating or modify VMs on the vCenter Server, verify that the vDS port groups created on the Cisco Nexus 1000V switch are available for use. On the Cisco Nexus 1000V switch CLI, use the **show port-profile** command to verify which VMs are using which port profiles.

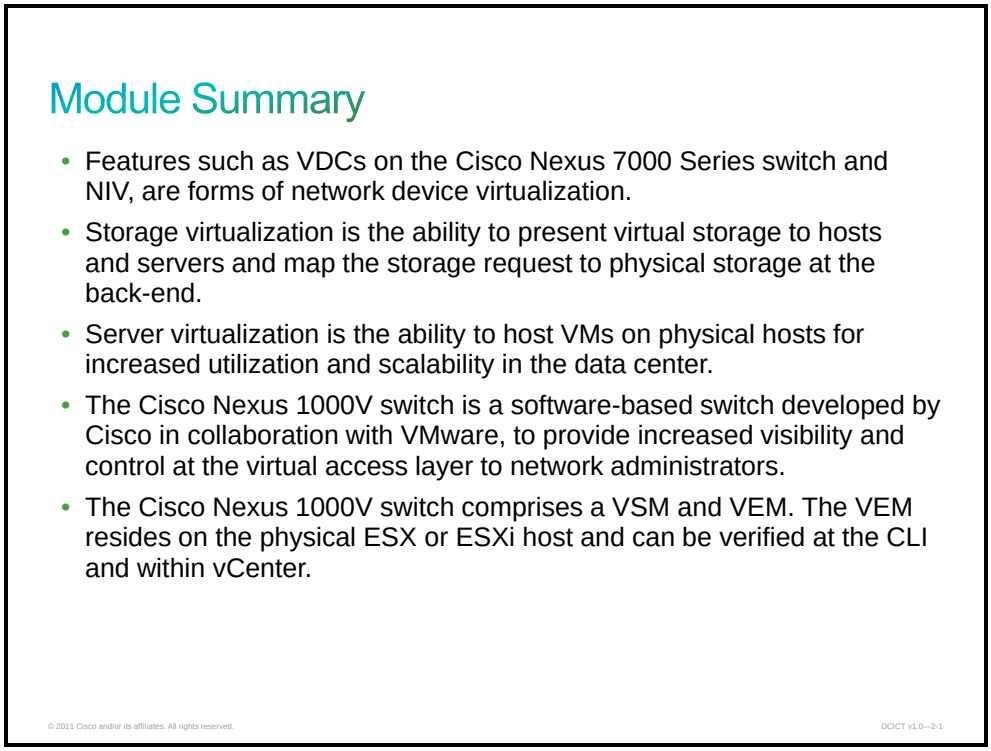
© 2012 Cisco and/or its affiliates. All rights reserved.

© 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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



Network device virtualization includes features such as virtual device contexts (VDCs) and Network Interface Virtualization (NIV). VDCs are used on the Cisco Nexus 7000 Series Switch, to provide the ability to consolidate onto few physical switches yet still retain the separation of domains. The Cisco Nexus 7000 Series Switch can be partitioned into a maximum of four VDCs. Each VDC is a logical switch that resides on one physical switch. This feature allows customers to consolidate various physical switches onto one physical infrastructure for greater flexibility and a reduced footprint in the data center. Each VDC runs its own processes and has its own configuration. The controlling VDC is VDC1, which is the default VDC.

The Cisco Nexus 2000 Fabric Extender is an external module that can be connected to a Cisco Nexus 7000 or 5000 Series Switch for greater port density without increased management overhead. The parent switch makes all the switching decisions, and the Cisco Nexus 2000 Fabric Extender provides the port count. Through the use of NIV and the VN-Link technology, the parent switch can recognize the remote ports on which traffic is sourced. From this determination, the switch can make the correct policy and switching decision.

Virtualization provides a process for presenting a logical grouping or subset of computing resources. Storage virtualization is used to create a common pool of all storage resources and to present a subset of those resources to the host or server. Storage virtualization is basically a logical grouping of logical unit numbers (LUNs).

Three main storage-system virtualization options are available: host-, array-, and network-based virtualization.

Several challenges exist inside the data center: management of physical resources, space constraints, cabling, power, and cooling, to name a few. To help reduce some of those challenges, servers can be moved from a physical environment to a virtual environment. Available technologies include VMware ESX or ESXi servers, Microsoft Hyper-V, and Linux enterprise virtualization. By virtualizing servers, companies can get better utilization of

https://t.me/learningnets

physical equipment and reduce their physical footprint, cabling, and power and cooling requirements. Virtual servers have the same capabilities and requirements as physical servers but can be managed through a center management infrastructure such as VMware vSphere vCenter Server.

To provide connectivity for the virtual server infrastructure, virtual switches can be used inside the physical host on which the virtual servers reside. VMware provides a software-based virtual Ethernet switch (vSwitch) or a vNetwork Distributed Switch (vDS). The vSwitch is a standalone switch that is managed individually. The vDS is a distributed switch that spans multiple hosts. This switch is managed through vCenter Server.

One disadvantage of the VMware implementation is that VMware administrators manage the virtual switch layer. These technicians do not always have a full understanding of network requirements and policies. In addition, the VMware-based switch does not necessarily have all the same features and policies as a normal network-based switch.

Cisco has worked with VMware to develop a software-based switch that can take advantage of the vDS architecture. The Cisco Nexus 1000V Series Switch provides a full Cisco Nexus Operating System (NX-OS) CLI, with all the features of a regular network-based switch, and is managed by the network team. Therefore, the network team has full visibility to the virtual access layer. The team can impose policies to the virtual server port as well as to the physical port of the host on which the virtual server resides.

The Cisco Nexus 1000V Series Switch comprises one or two Virtual Supervisor Modules (VSMs) and one or more Virtual Ethernet Modules (VEMs). The VSM is installed on either a virtual machine (VM) or a Cisco Nexus 1010 appliance. The VEM is installed on the ESX or ESXi host. Port profiles are created on the Cisco Nexus 1000V Series Switch, to provide network policies that the virtual servers can use. These port profiles are known as port groups on the vCenter Server. The VMware administrator applies these port groups to VMs that are created or modified. The port profiles are pushed from the Cisco Nexus 1000V Series Switch to vCenter, by using the secure connection that is established between the devices.

Verification of VSM and VEM installation can be performed on the Cisco Nexus 1000V Series Switch, vCenter, and the CLI of the ESX or ESXi host. Port profile configuration can be verified on the Cisco Nexus 1000V Series Switch, and port group creation can be verified on the vCenter Server.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Storage Networking. http://www.cisco.com/en/US/products/hw/ps4159/index.html.
- Cisco Systems, Inc. *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide.* San Jose, California, October 2011. http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/virtual_device_context/configuration/guide/vdc_nx-os_cfg.html.
- Cisco Systems, Inc. *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide.* San Jose, California, April 2012. http://www.cisco.com/en/US/docs/switches/datacenter/nexus2000/sw/configuration/guide/rel_6_0/b_Configuring_the_Cisco_Nexus_2000_Series_Fabric_Extender_rel_6_0.html.

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

- Cisco Systems, Inc. Cisco Nexus 5000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.1(3)N1(1), Configuring the Fabric Extender. http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/layer2/513_n1_1/b_Cisco_n5k_layer2_config_gd_rel_513_N1_1_chapter_010100.html.
- Cisco Systems, Inc. Cisco Nexus 1000V Series Switches Configuration Guides. http://www.cisco.com/en/US/products/ps9902/products_installation_and_configuration_guides_list.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) How many VDCs are supported on a Cisco Nexus 7000 Series switch? (Source: Virtualizing Network Devices)

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

Q2) Which two features can be configured only in the default VDC? (Choose two.) (Source: Virtualizing Network Devices)

A) VLANs
B) CoPP
C) VDC resource allocation
D) VRFs
E) management IP address

Q3) Which ports comprise a port group on the N7K-M132XP-12 module? (Source: Virtualizing Network Devices)

A) 1–4
B) 1–8
C) 1, 3, 5, and 7
D) 1 and 2

Q4) Which command do you use to move from the default VDC to a nondefault VDC on the Cisco Nexus 7000 Series switch? (Source: Virtualizing Network Devices)

A) **changeto vdc**
B) **moveto vdc**
C) **skipto vdc**
D) **switchto vdc**

Q5) How many days are available in the grace license period on the Cisco Nexus 7000 Series switch? (Source: Virtualizing Network Devices)

A) 180
B) 120
C) 90
D) 60

Q6) Which command must you enable in the default VDC on a Cisco Nexus 7000 Series switch to enable a Cisco Nexus 2000 Fabric Extender to be attached and configured in a nondefault VDC? (Source: Virtualizing Network Devices)

A) **feature-set fex**
B) **feature fex**
C) **install feature fex**
D) **install feature-set fex**

Q7) In which mode are interfaces that connect to the Cisco Nexus 2000 Fabric Extender listed in the **show interface brief** command on a Cisco Nexus 5000 or 7000 Series switch? (Source: Virtualizing Network Devices)

https://t.me/learningnets

A) fabric
B) access
C) trunking
D) host

Q8) Which feature do you configure on storage arrays to provide basic LUN-level security? (Source: Virtualizing Storage)

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

Q9) Which form of server virtualization simulates some, but not all, of the hardware environment? (Source: Virtualizing Server Solutions)

A) partial virtualization
B) paravirtualization
C) full virtualization
D) host virtualization

Q10) Which feature requires VMware vCenter Server? (Source: Virtualizing Server Solutions)

A) NIC teaming
B) vSwitch
C) port groups
D) vDS

Q11) How many VSMs can be installed in the Cisco Nexus 1000V virtual chassis? (Source: Using the Cisco Nexus 1000V Series Switch)

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

Q12) How many VEMs does the Cisco Nexus 1000V virtual chassis support? (Source: Using the Cisco Nexus 1000V Series Switch)

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

Q13) What defines a logical group of ports with the same configuration as on the VMware vCenter Server? (Source: Using the Cisco Nexus 1000V Series Switch)

A) port profiles
B) network connection policy
C) port groups
D) group policies

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

Q14) What is the maximum number of network ports that can be assigned to a VMware vSwitch? (Source: Using the Cisco Nexus 1000V Series Switch)

A) 64
B) 32
C) 16
D) 8

Q15) Which method is preferred for installing the Cisco Nexus 1000V Series switch? (Source: Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch)

A) OVA file
B) OVF file
C) manual installation
D) ISO file

Q16) Which is the first slot in which a VEM is installed on the Cisco Nexus 1000V virtual chassis? (Source: Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch)

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

Q17) Which three port groups need to be configured on vCenter to support the Cisco Nexus 1000V Series switch? (Source: Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch)

A) management, control, packet
B) management, control, data
C) control, packet, data
D) control, packet, default

Q18) On which two devices can you install a Cisco Nexus 1000V Series switch? (Choose two.) (Source: Verifying Setup and Operation of the Cisco Nexus 1000V Series Switch)

A) desktop computer
B) standalone server
C) VM
D) Cisco Nexus 1010 appliance
E) Windows 2008 Server

https://t.me/learningnets

Module Self-Check Answer Key

Q1) A

Q2) B, C

Q3) C

Q4) D

Q5) B

Q6) D

Q7) A

Q8) A

Q9) A

Q10) D

Q11) B

Q12) B

Q13) C

Q14) B

Q15) A

Q16) A

Q17) A

Q18) C, D

© 2012 Cisco Systems, Inc.

https://t.me/learningnets