Table of Contents

Volume 2

Maintaining a Cisco Unified Computing System 7-1

Overview 7-1
Module Objectives 7-1

Backup and Restore 7-3

Overview 7-3
Objectives 7-3
Backup Types 7-4
Full-State 7-4
Backup and Import Jobs 7-7
Administrative State 7-8
Backup Configuration 7-9
Import Configuration 7-11
Merge or Replace 7-11
Full-State or Configuration Only 7-11
Protocols 7-11
Summary 7-16

Managing High Availability 7-17

Overview 7-17
Objectives 7-17
High Availability Connectivity 7-18
High Availability Requirements 7-18
Cluster Interfaces 7-19
High Availability Terminology 7-20
Cisco UCS Controller 7-21
Local Storage 7-22
Chassis EEPROM 7-22
Agreement 7-23
Stability 7-23
Infrequent Elections 7-23
Stability under Quick Restart 7-23
Configuring High Availability 7-24
Configuring High Availability 7-24
Managing High Availability 7-29
Switchover 7-29
Failover 7-29
Partition in Space 7-30
Partition in Time 7-30
Summary 7-36

Monitoring System Events 7-37

Overview 7-37
Objectives 7-37
GUI Monitoring 7-38
Remote Monitoring 7-43
CLI Status Monitoring 7-46
Threshold Policies 7-51
Call Home Configuration 7-59
Exporting Data 7-66
Ethanalyzer 7-69
Summary 7-71
Module Summary 7-73
Module Self-Check 7-75

https://t.me/learningnets

Cisco Unified Computing System Deployments 8-1

Overview 8-1
Module Objectives 8-1

UCS Use Cases 8-3

Overview 8-3
Objectives 8-3
Use Case: Server Operations 8-4
Use Case: Virtualization 8-13
Expanded Memory Blade Servers (Cisco UCS B250-M1) 8-14
Virtualization Integration 8-14
Unified Fabric 8-14
Server Statelessness 8-15
Use Case: Databases 8-21
Summary 8-23

Operating System Provisioning 8-25

Overview 8-25
Objectives 8-25
SAN Based installation 8-26
Local Disk Installations 8-29
Network-Based Installation 8-30
Summary 8-31
Module Summary 8-33
Module Self-Check 8-35

Networks in a Virtual Server Environment 9-1

Overview 9-1
Module Objectives 9-1

Introduction to the Cisco Nexus 1000V 9-3

Overview 9-3
Objectives 9-3
Data Center Access Layer Networks 9-4
Virtual Environment Overview 9-9
Cisco Virtual Switching Overview 9-15
Feature Overview 9-19
Summary 9-25

Server Virtualization Primer 9-27

Overview 9-27
Objectives 9-27
Server Virtualization 9-28
VMware Virtualization 9-32
Virtual Networking with VMware 9-37
Virtual Storage Networking with VMware 9-40
VMFS 9-42
LUN Performance Considerations 9-43
Using Raw Device Mapping 9-44
VMware Environment Management 9-45
Summary 9-50

VMware Ethernet Networking 9-51

Overview 9-51
Objectives 9-51
Virtual Access Layer 9-52
The VMware vSwitch 9-61
VMware Network Configuration 9-69
Distributed Virtual Switching 9-75
Summary 9-80

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

The Cisco Nexus 1000V Architecture 9-81
Overview 9-81
Objectives 9-81
Cisco Nexus 1000V 9-82
Step 1 9-87
Step 2 9-87
Step 3 9-87
Step 4 9-87
Step 5 9-88
Step 1 9-89
Step 2 9-89
Step 3 9-89
Step 4 9-89
Step 5 9-89
Architectural Overview 9-91
Cisco Nexus 1000V Installation and Operation 9-99
Summary 9-109
Configure VMware Switching 9-111
Overview 9-111
Objectives 9-111
Configure vSwitches 9-112
Configuring VMware DVS 9-118
Summary 9-134
Installing and Configuring the Cisco Nexus 1000V 9-135
Overview 9-135
Objectives 9-135
Configure VSM vSwitch Networks 9-136
VSM Virtual Machine Creation 9-141
Initial VSM Configuration 9-159
Connect the VSM to vCenter 9-164
Cisco Nexus 1000V High Availability Configuration 9-171
Summary 9-175
Configure Basic Cisco Nexus 1000V Networking 9-177
Overview 9-177
Objectives 9-177
Port Profiles 9-178
VLAN Configuration 9-188
Creating Uplink Profiles 9-191
Creating VM Data Port Profiles 9-195
Configuring Cisco Nexus 1000V Port Channels 9-197
Symmetric 9-200
Asymmetric 9-200
Adding Hosts to the VSM 9-204
Adding VMs to a VSM 9-209
Summary 9-211
Module Summary 9-213
Module Self-Check 9-215

https://t.me/learningnets

Data Center Unified Computing Implementation Appendix 10-1

Overview 10-1
Module Objectives 10-1

Exploring the Unified Computing User Interfaces 10-3

Overview 10-3
Objectives 10-3
Command-Line Interface 10-4
Other Management Commands 10-13
The connect Commands 10-19
UCS Acronyms 10-25
Summary 10-29

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module 7

Maintaining a Cisco Unified Computing System

Overview

This module describes the maintenance of a Cisco Unified Computing System (Cisco UCS).

Module Objectives

Upon completing this module, you will be able to describe maintenance and high availability processes in a Cisco Unified Computing System. This includes being able to meet these objectives:

- Explain the backup and restore capabilities the Cisco UCS platform
- Describe how to maintain high-availability configurations
- Describe how to monitor system events

https://t.me/learningnets

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

Lesson 1

Backup and Restore

Overview

In this lesson, you will learn about the backup and restore capabilities of the Cisco Unified Computing System (Cisco UCS) platform.

Objectives

Upon completing this lesson, you will be able to explain the backup and restore capabilities the Cisco UCS platform. This includes being able to meet these objectives:

- Describe the supported backup types
- Describe backup and import jobs
- Describe how to configure a backup job
- Describe how to configure an import job

https://t.me/learningnets

Backup Types

This section describes the types of backups available in Cisco UCS.

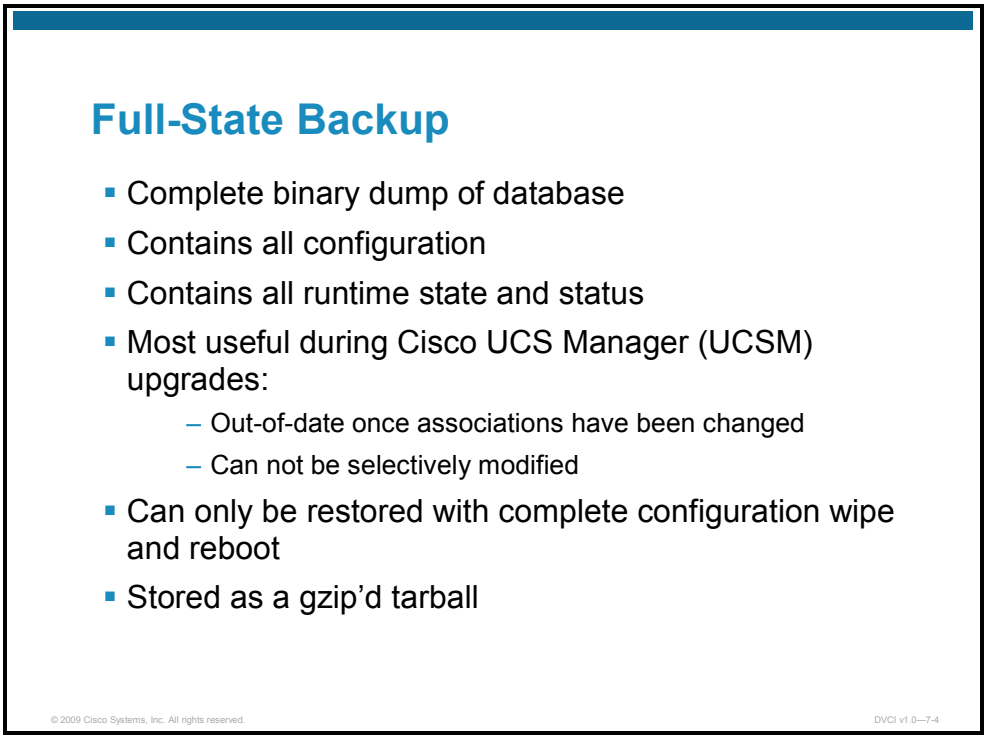

Full-State

Full-state is a disaster recovery-style backup that is intended for a catastrophic failure of both your UCSMs. This backup type includes all the runtime state information (for example, finite state machine [FSM] state of blades, associated state of Service Profiles, and so on), as well as configuration information (for example, users, policies, and so on).

Since the full-state backup includes runtime state, this type of backup can very quickly become out-of-date. Any changes to blade Service Profile associations will render this backup obsolete. If a full-state backup is applied that contains out-of-date association information, Cisco UCS does not make any effort to validate the associations. Only through a manual (and disruptive) rediscovery and re-association process can the runtime state be brought back into sync.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

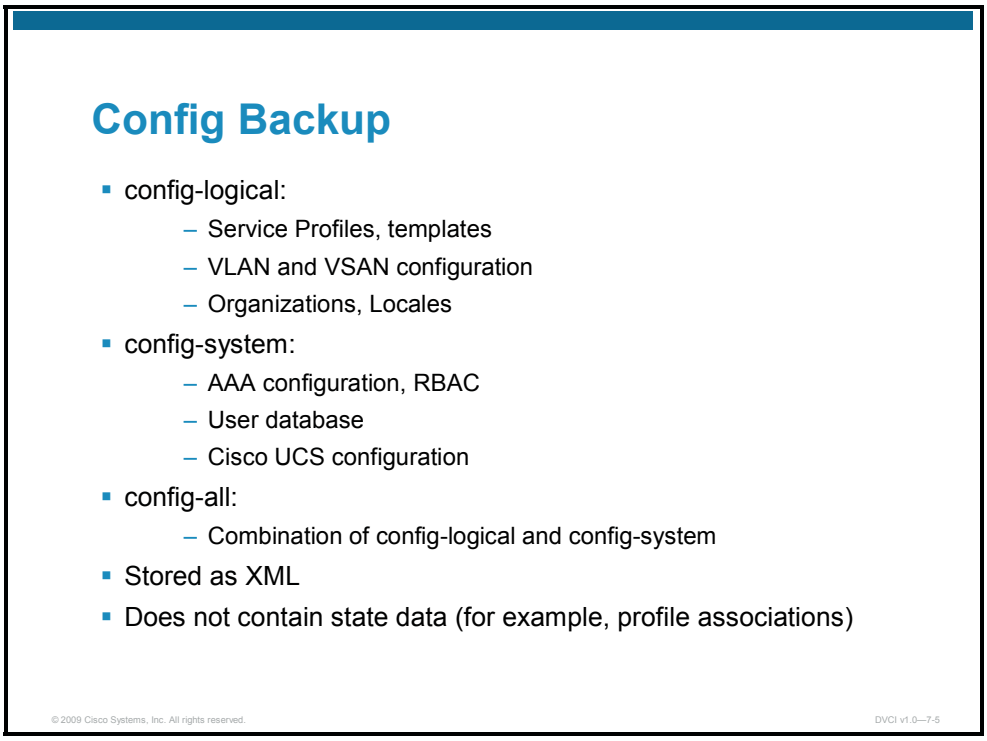

- **config-logical**: All configuration that is not associated with authentication, authorization, and accounting (AAA) (such as configured organizations, configured threshold policies, and configured VLANs and VSANs in your LAN and SAN clouds, respectively).
- **config-system:** All configuration that specifically pertains to the AAA role. Examples include RADIUS, Lightweight Directory Access Protocol (LDAP), TACACS, and users.
- **config-all:** A combination of config-logical and config-system.

As "config" style backups are simply Extensible Markup Language (XML) representations of the configuration of the Cisco UCS Manager, they can easily be edited using an XML editor or text editor. This makes them useful for creating templates for other Cisco UCS implementations, or for adjusting them if changes were made to the environment since the backup was last taken. No runtime state data (Service Profile associations, and so on) is stored within these backups.

https://t.me/learningnets

Use of Config Backups

XML configuration backups can be easily edited:

- Useful for duplicating configuration in other implementations
- Can be modified before import in DR scenarios
- Can be pre-populated in consulting engagements

© 2009 Cisco Systems, Inc. All rights reserved.

XML configuration-level backups can be very useful for more than simply recovery of a failed system. XML is easily edited by a variety of editors, and can be easily modified or customized before import.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Backup and Import Jobs

This section describes how Cisco UCS manages backup and import jobs.

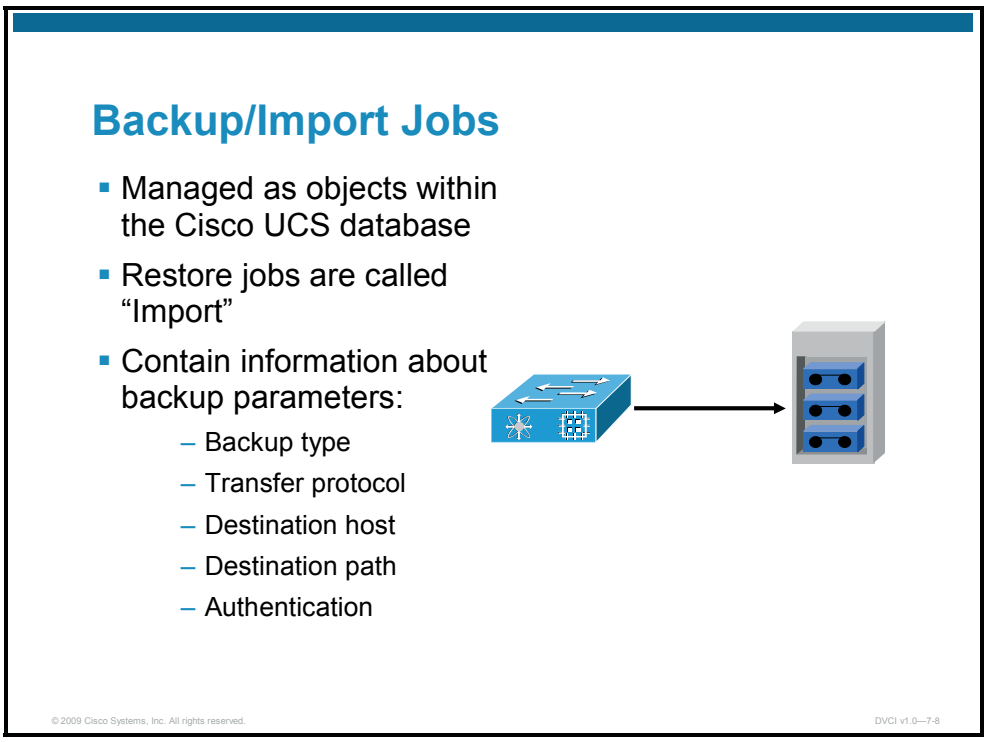

Cisco UCS handles backup and import operations as managed objects within the Cisco UCS Manager database. No backup data is stored within the database, only the information relating to the storage and transfer of the backup.

https://t.me/learningnets

Life-Cycle of Backup Jobs

- When created, object state becomes “enabled”
- Backup files are created, archived, and transferred to specified to destination
- Object state becomes “disabled”
- Backup can be re-run at any time by changing status to “enabled”
- A similar process is used for Import jobs

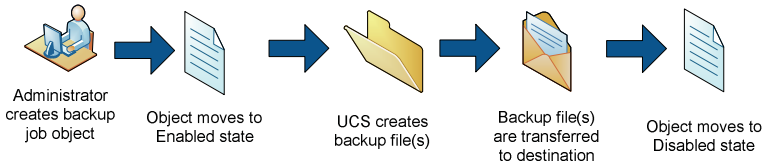

© 2009 Cisco Systems, Inc. All rights reserved.

Administrative State

After you create the managed object that represents the backup operation, the administrative state of the backup becomes “enabled”. In the background, the state and/or configuration is collected and transferred to the backup server using the protocol that you choose. You can follow the workflow of the backup by examining the FSM that is associated with the backup. Once the backup is complete, its administrative state changes to “disabled”.

https://t.me/learningnets

Backup Configuration

This section describes how to perform backup operations in a Cisco UCS environment.

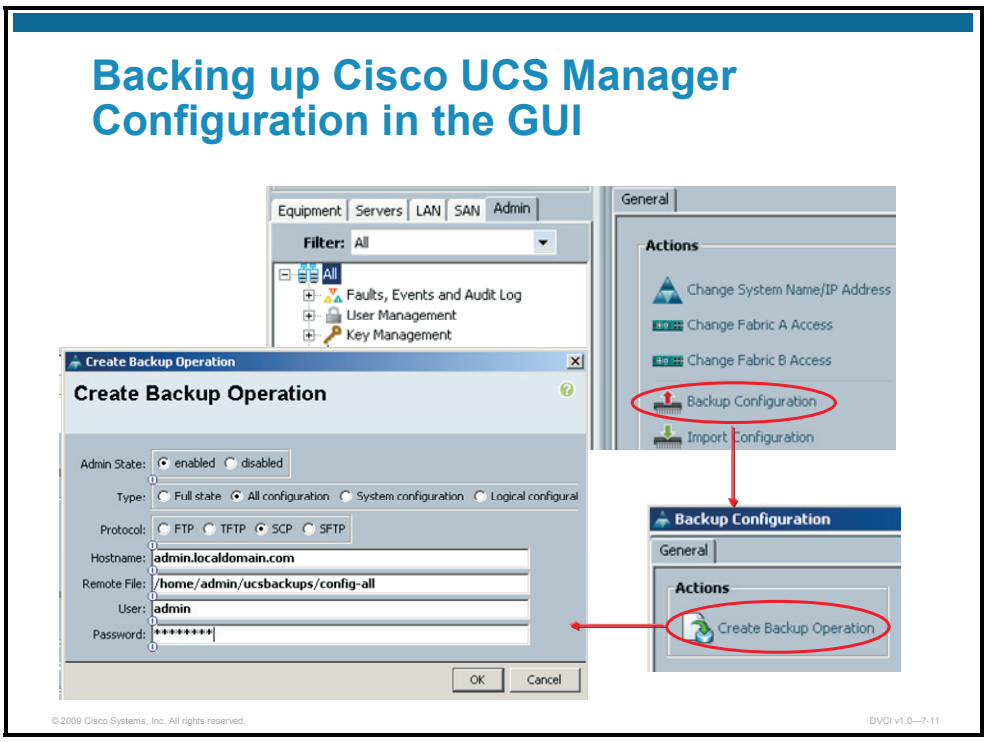

Currently, only one backup configuration can be created per backup server. Backups are identified in the Cisco UCS Manager using the hostname or IP address of the backup server in the configuration.

Click the **Admin** tab in the navigation pane, and select the **All** object. Click the **Backup Configuration** link in the actions pane, and then click **Create Backup Operation** in the following window. Finally, provide the proper configuration and connectivity values for the backup job. A backup job created in the Admin State of "enabled" will run immediately. After the backup job is completed, it changes to a "disabled" state. If created as "disabled", the job does not run until enabled.

https://t.me/learningnets

Backing up the Cisco UCS Manager Configuration in the CLI

```
s6100-A# scope system
s6100-A /system # create backup
  scp://admin@admin.localdomain.com/backups full-state enabled
Password:******
s6100-A /system/backup* # commit-buffer
s6100-A /system/backup # up
s6100-A /system # show backup detail
Backup:
Hostname: admin.localdomain.com
    Remote File: /backups
    Type: Full State
    User: admin
    Protocol: Scp
    Administrative State: Enabled
    Description:
```

© 2009 Cisco Systems, Inc. All rights reserved.

Log into the Cisco UCS Manager command-line interface (CLI) using Secure Shell (SSH) or Telnet using the cluster IP address. If in a single-Fabric Interconnect environment, use the management IP address of the single node.

Change the CLI scope to **system**.

The command executed in the figure sends the backup file (via SCP) to the backup server in the /backups directory. The backup file is a gzip-compressed tar archive. Repeat the **show backup detail** command until the administrative state changes to disabled, at which point the backup will have successfully completed.

https://t.me/learningnets

Import Configuration

This section describes how to perform import operations in a Cisco UCS environment.

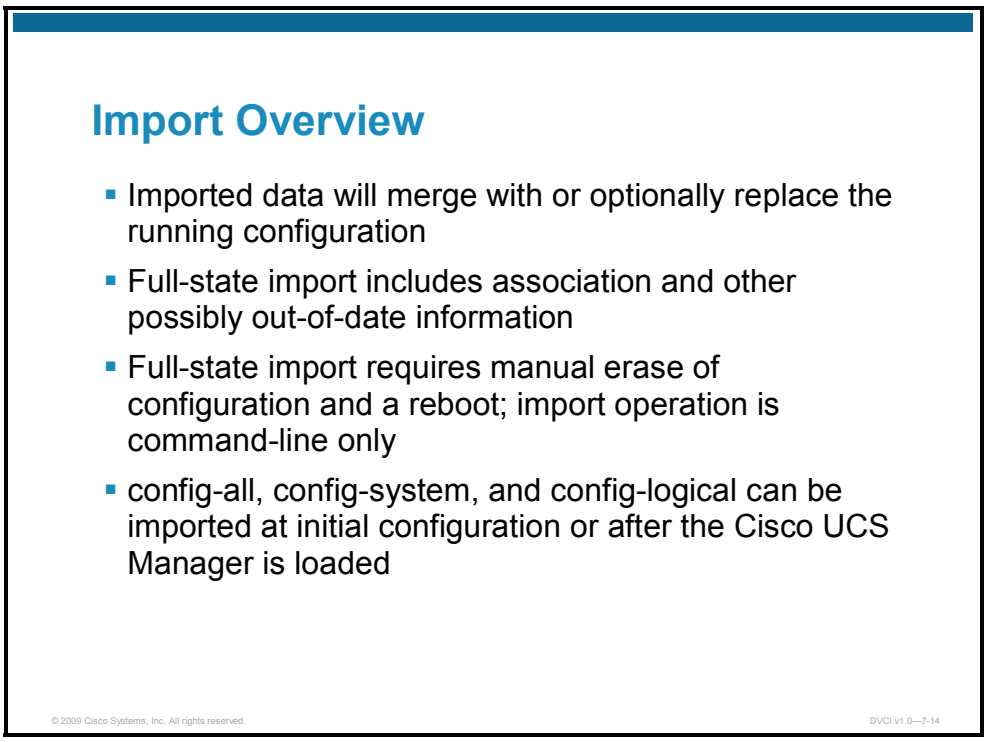

Merge or Replace

You can merge or replace a backup into your configuration at runtime without service interruption. The "replace" option will completely remove the current configuration and replace it with the backup configuration. The "merge" option will attempt to create managed objects that exist in the backup but that do not exist in the current configuration.

Full-State or Configuration Only

If you restore a full-state backup, then you must erase the configuration, reboot the Fabric Interconnect, and specify the following information at startup:

- IP address/netmask of Fabric Interconnect
- IP address of gateway
- IP address of backup server
- Protocol/user credentials/path/filename

If you choose to restore a configuration-style backup, you can either perform the restore at Fabric Interconnect boot or at runtime.

Protocols

The protocol that you use to restore a backup need not be the same as the protocol that was used to transfer the backup. However, the same caveats affecting the backup will apply, with regard to the running protocols, user credentials, and paths.

https://t.me/learningnets

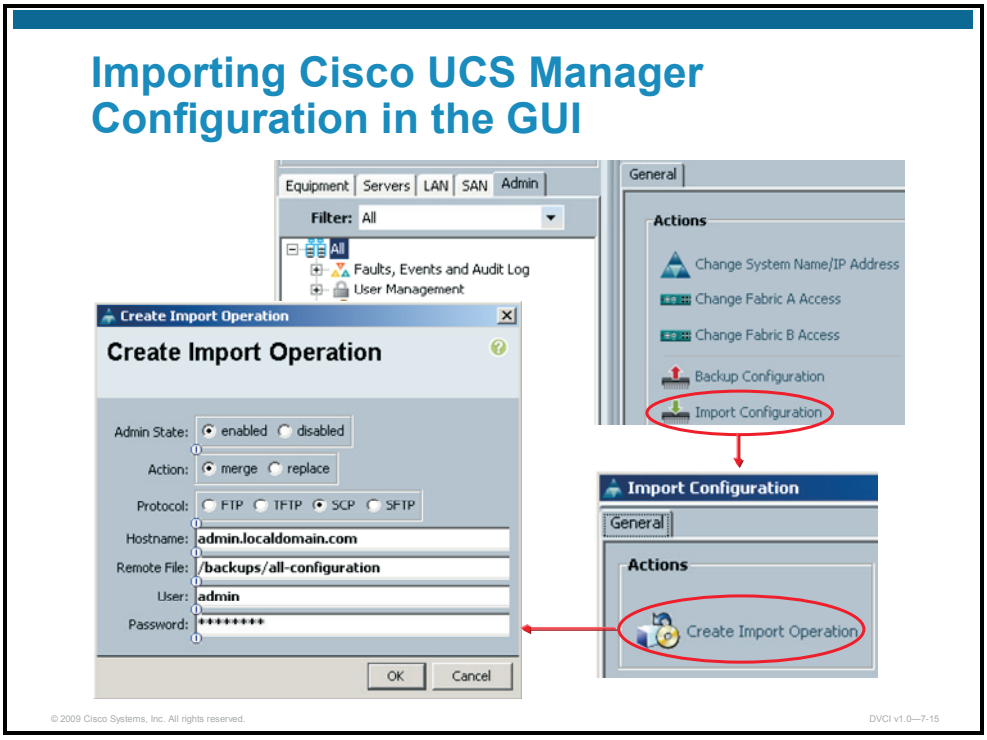

The Cisco UCS Manager configuration can be restored through an import operation Configuration backups can be restored at runtime. Full-state backups can not be restored at runtime. Full-state backups can only be restored at startup.

To create the Import job, click the **Admin** tab in the navigation pane. Select the **All** object and click the **Import Configuration** link. In the data pane, click the **Create Import Operation** link and define the parameters for the import operation.

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

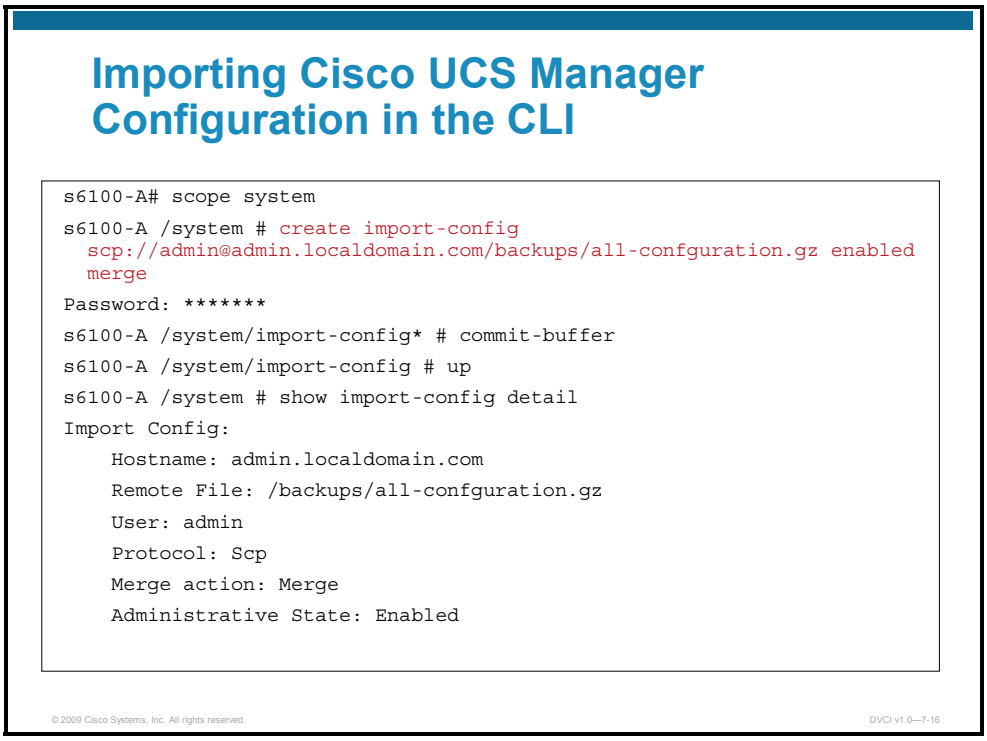

You can also restore backup configurations at runtime, using the CLI. The default import operation is a merge (rather than a replace), unless otherwise specified in the command syntax. The **create import-config** command creates an import job object, in the same way as the GUI.

https://t.me/learningnets

Importing Full-State Backup

```
s6100-A# erase configuration
Configuration will be erased. Are you sure? (yes/no): yes
Removing all the configuration. Please wait....
Configurations are cleaned up.
Rebooting....
…
---- Basic System Configuration Dialog ----
…
 Enter the setup mode; setup newly or restore from backup. (setup/restore) ?
  restore
…
  Continue to restore this Fabric interconnect from a backup file (yes/no) ?
  yes

  Physical Switch Mgmt0 IPv4 address : 192.168.10.101
  Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
  IPv4 address of the default gateway : 192.168.10.254
```

© 2009 Cisco Systems, Inc. All rights reserved.

You must erase the current configuration (which invokes a reboot) in order to perform a restoration of a full-state backup.

The network configuration for the Fabric Interconnect must be provided in order for the request to go over the network to the backup server to retrieve the backup file. Note that any network configuration applied here will be overwritten by the information stored in the backup.

https://t.me/learningnets

Importing Full-State Backup (Cont.)

```
Enter the protocol to get backup file (scp/ftp/tftp/sftp) ?  ftp
 Enter the IP address of backup server: 10.100.100.30
 Enter fully qualified backup file name: /backups/ucs-fullbackup.tgz
 Enter user ID: adminuser
 Enter password: *******
   Retrieved backup configuration file.
 Configuration file - Ok


Cisco UCS 6100 Series Fabric Interconnect
s6100-A login:
```

- After retrieving and applying the backup, the system is ready for use

© 2009 Cisco Systems, Inc. All rights reserved.

You must erase the current configuration (which invokes a reboot) in order to perform the restoration of a full-state backup.

https://t.me/learningnets

Summary

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

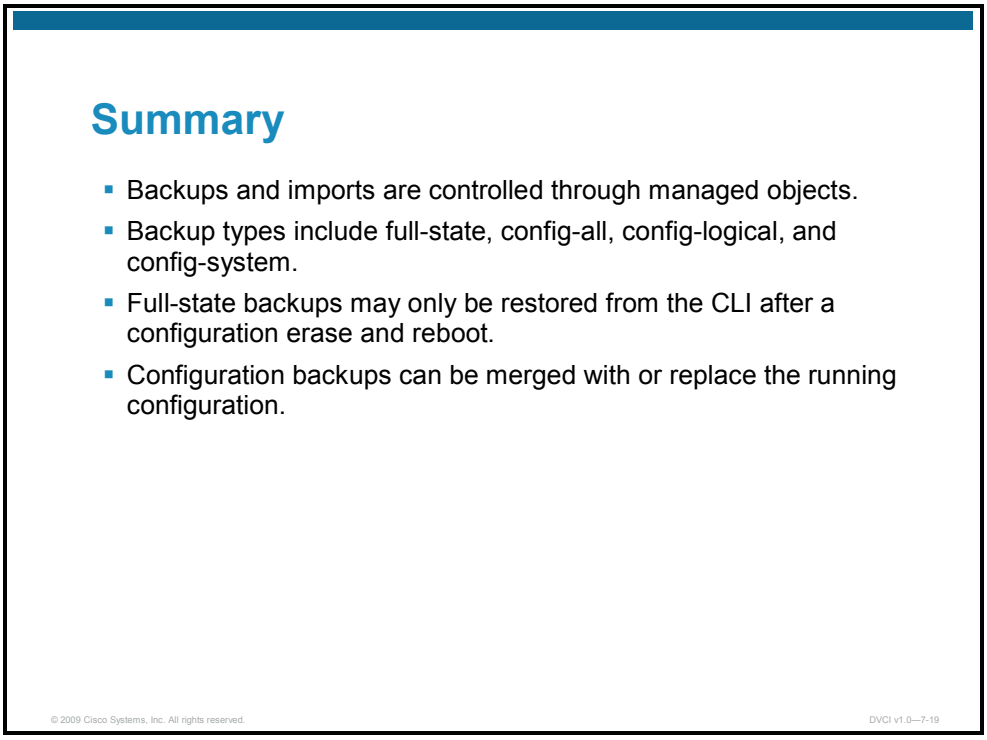

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Managing High Availability

Overview

In this lesson, you will learn how to configure and maintain high availability options within the Cisco Unified Computing System (Cisco UCS).

Objectives

Upon completing this lesson, you will be able to describe how to maintain high-availability configurations. This includes being able to meet these objectives:

- Describe the high availability connectivity of Fabric Interconnects
- Explain the terms used by Cisco UCS regarding high availability
- Describe how to configure high availability
- Describe how to manage high availability

https://t.me/learningnets

High Availability Connectivity

This topic describes how Fabric Interconnects connect in a highly available configuration.

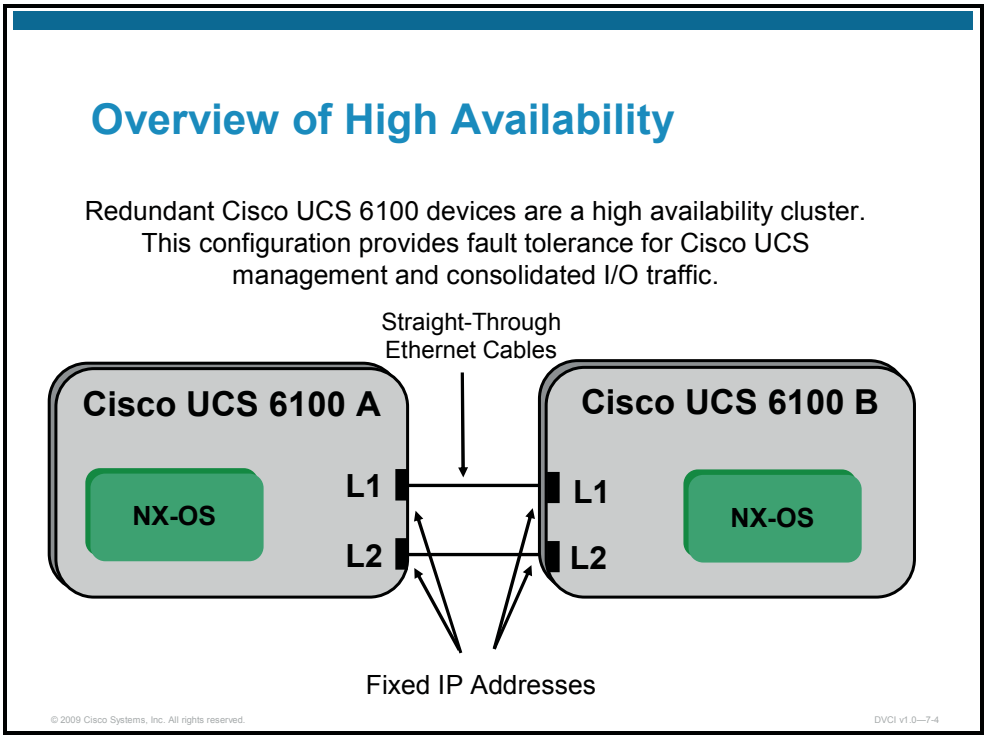

Configuring two Fabric Interconnects in your Cisco UCS system can provide redundancy for both the Cisco UCS Manager functionality and the switching functionality that the Fabric Interconnect provides.

High Availability Requirements

- Connect L1 of Fabric Interconnect A to L1 of Fabric Interconnect B
- Connect L2 of Fabric Interconnect A to L2 of Fabric Interconnect B.
- Connect Fabric Interconnect A to IOM A of each chassis using 1,2, or 4 uplinks
- Connect Fabric Interconnect B to IOM B of each chassis using 1,2, or 4 uplinks

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

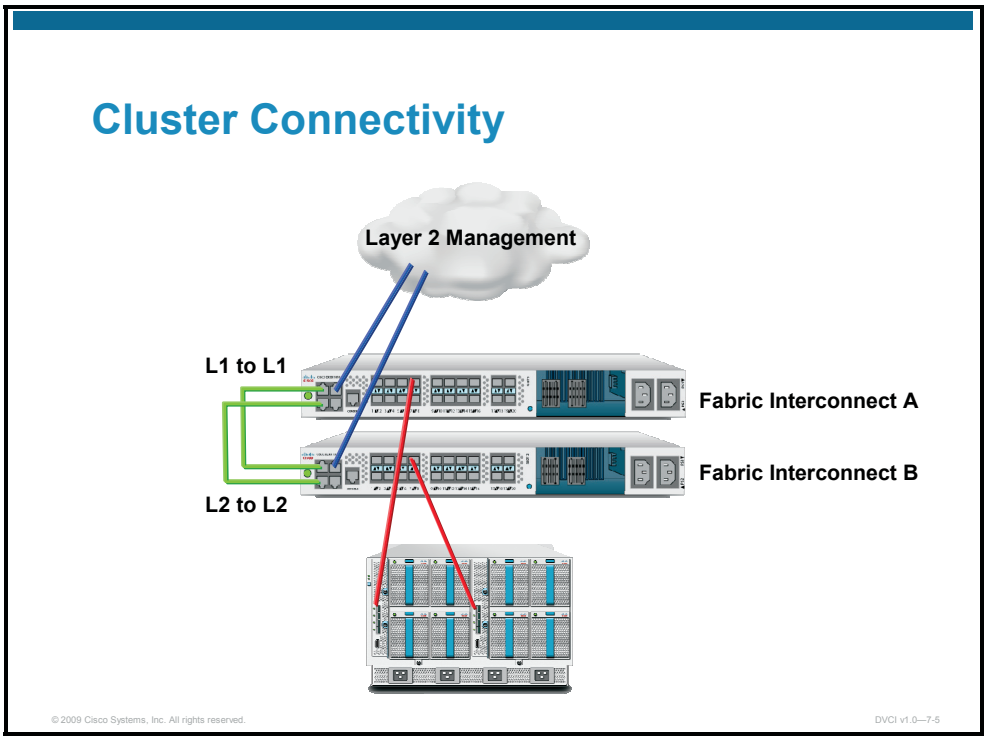

Cluster Interfaces

These interfaces provide a cluster link between two Cisco UCS 6100 Fabric Interconnects. They carry the cluster heartbeat messages between the two Fabric Interconnects as well as high-level messages between Cisco UCS Manager elements. These links are part of an 802.3ad bond managed by the underlying operating system. The bond is configured to run Link Aggregation Control Protocol (LACP). The IP addresses on these links are fixed. The ports are labeled “L1” and “L2” on the Fabric Interconnect and operate at 10/100/1000 speeds.

The management port (mgmt0) of each Fabric Interconnect should be connected to the same Layer 2 network to facilitate failover and failback of the management IP address.

Each Fabric Interconnect should connect to only one “side” of each chassis.

https://t.me/learningnets

High Availability Terminology

This topic explains the terminology used in Cisco UCS high-availability configurations.

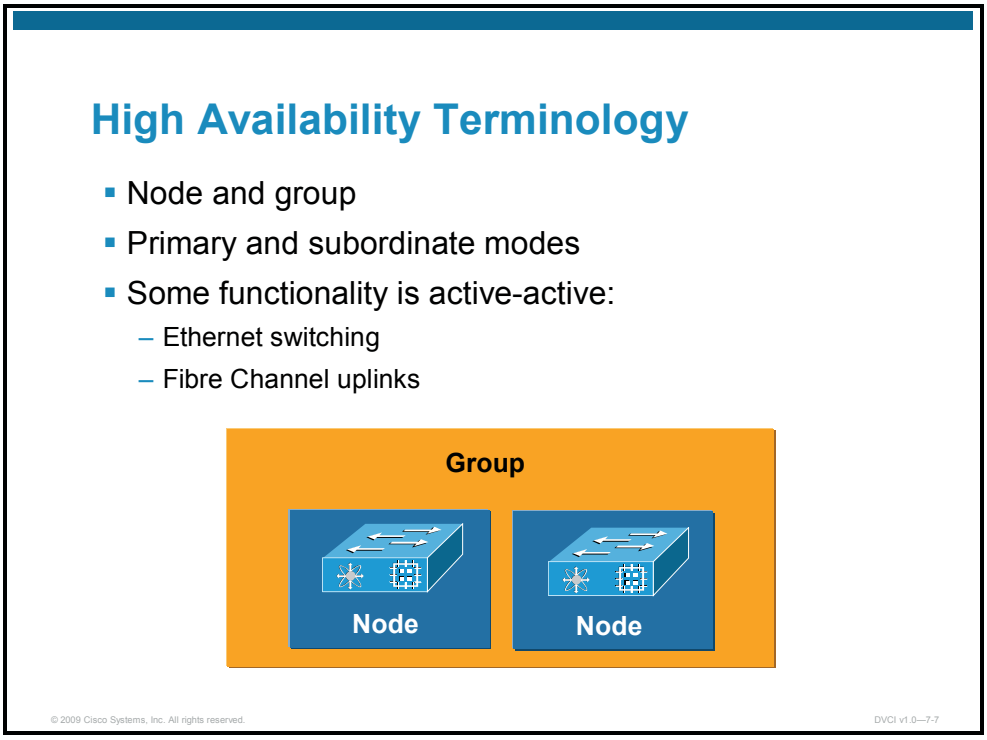

A node is one of the two Fabric Interconnects. The set of two nodes is called a group. Each node is identified by a unique node ID, generated by the Cisco UCS Manager, and stored in the Cisco UCS Manager database.

All interactions with the external world, such as the user interface (UI), command-line interface (CLI), Keyboard/Video/Mouse (KVM) access to nodes, and so on, go through the primary the Cisco UCS Manager. The standby node runs an instance of the Cisco UCS Manager with reduced functionality; this node is designated as “subordinate”.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

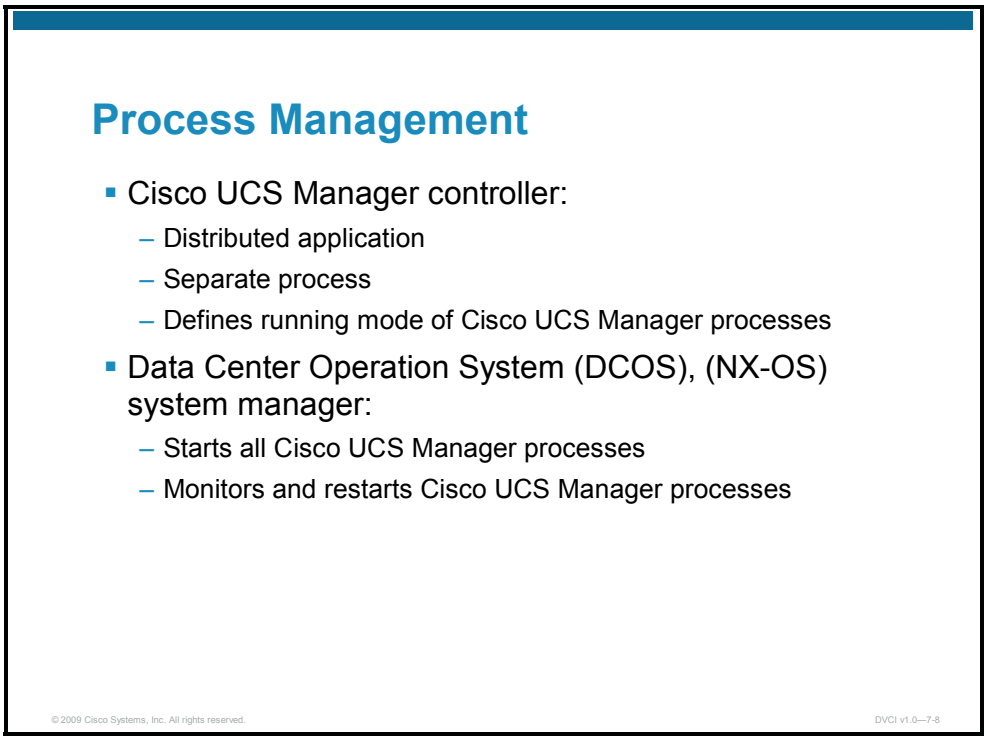

Cisco UCS Controller

The Cisco UCS Manager controller is a distributed application running on both primary and subordinate Cisco UCS Manager. Each instance is represented by a unique ID (the same as the node ID). The Cisco UCS Manager controller is implemented as a separate process; the address space separation guarantees a higher degree of fault isolation. It also allows the controller to distinguish between a crash of other system processes and a crash of the controller itself. The Cisco UCS Manager controller decides which Cisco UCS Manager components should run in primary or subordinate mode.

All the Cisco UCS Manager processes are always started on both nodes.

https://t.me/learningnets

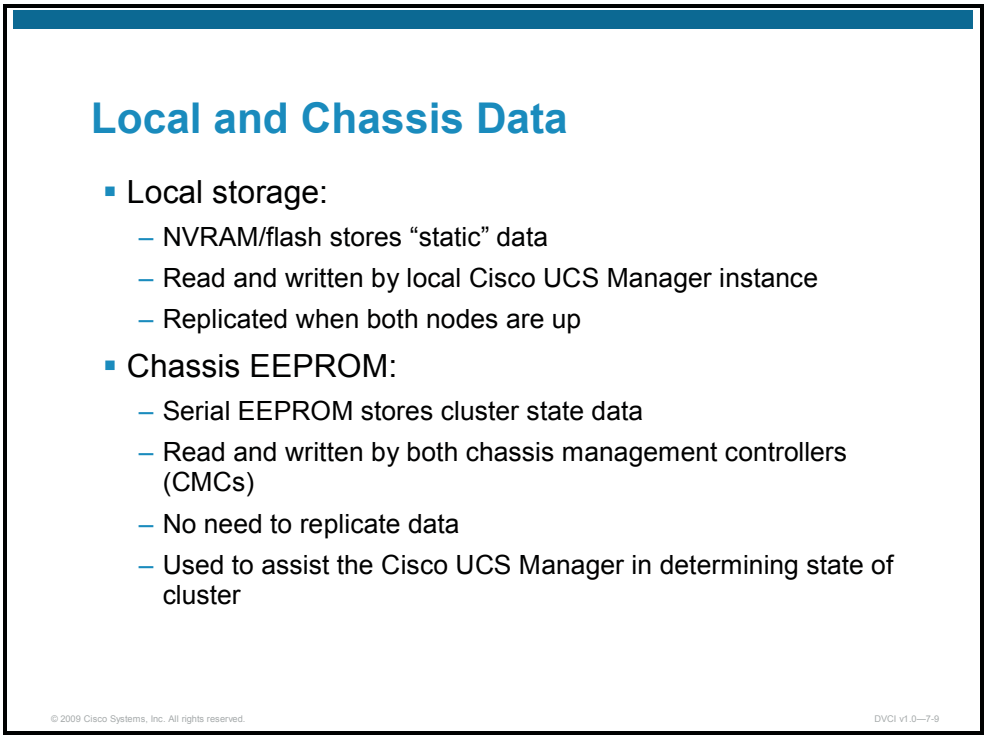

Local Storage

Each Cisco UCS Fabric Interconnect maintains its own local storage in NVRAM and flash memory. Local storage contains static data—that is, storage that does not change with cluster membership changes. For example, installable images are stored in the /bootflash portion of the flash memory. Data such as installable images are replicated at runtime, while both cluster members are present in the cluster. You do not need to (nor can you) download images via the Cisco UCS Manager interface to individual nodes. The download is replicated to both nodes. If a node is not present during an image download, then that image is replicated to that node when it rejoins the cluster.

Chassis EEPROM

Each CMC maintains its own part of the shared chassis storage in the serial EEPROM (SEEPROM). Chassis storage contains a combination of static and dynamic information. For example, the static portion contains the node ID for each node configured in the cluster. The dynamic portion contains the version of the configuration as seen by that node. There is no need to replicate the contents of the SEEPROM. Each node maintains its own portion while both nodes may read from both topics.

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

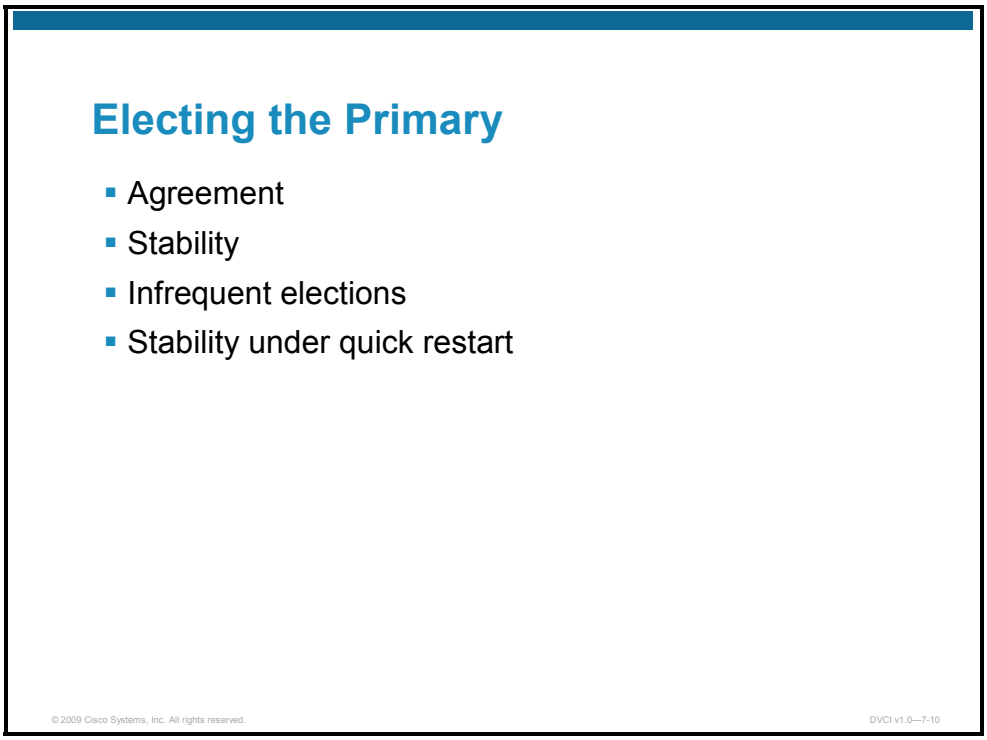

Agreement

A process declares a new leader when the following conditions apply:

- It has received acknowledges that his election request has been processed
- By checking the election counter in the incoming messages
- All processes propose the same new leader

Stability

The leadership should only change in one of the following two cases:

- An administrative change in the configuration requires the leader to be moved
- The leader process has crashed

Infrequent Elections

Elections are only caused by the following events:

- Administrative configuration change
- New process joining the group
- Process exiting the group
- Process crash

Stability under Quick Restart

We allow a leader process to crash, restart, and still join the group as the leader. This prevents a change of leadership when a process (or a node) is the subject of a quick restart. The rationale here is that a change of leadership and therefore a switchover can be more expensive than waiting for the leader process (or node) to re-initialize.

https://t.me/learningnets

Configuring High Availability

This topic describes how to configure high availability.

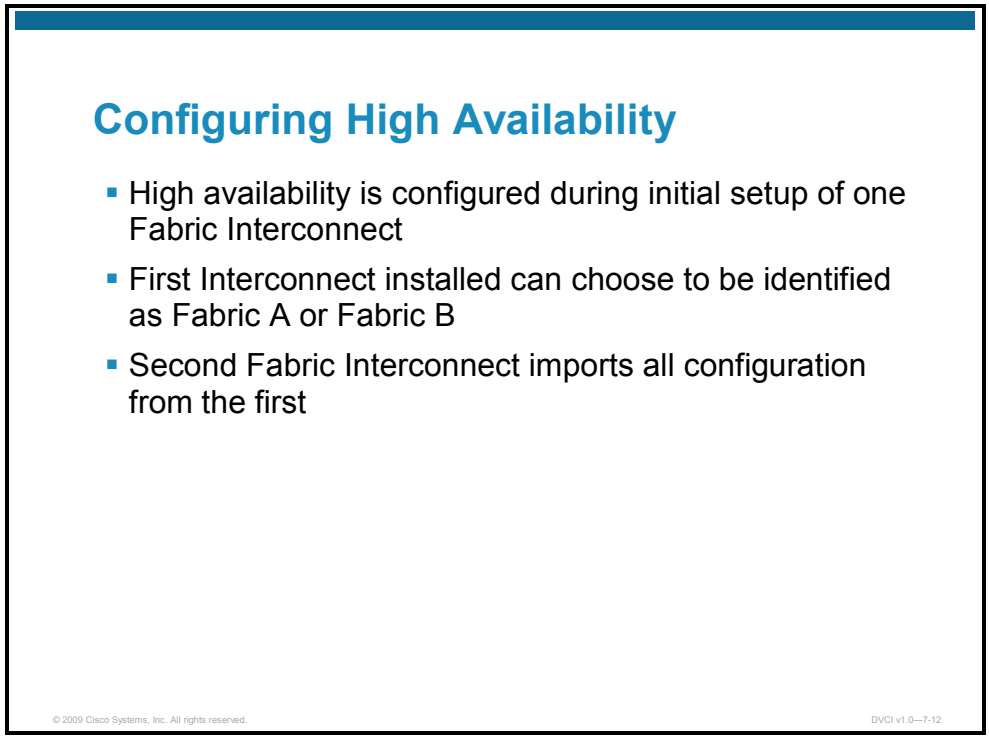

Configuring High Availability

High availability is configured during the initial configuration of the Fabric Interconnects. As shown in the following figures, only the first configured node requires any cluster-specific configuration. The first Fabric Interconnect configured will initially be the primary node, only because it is the first node in the cluster.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

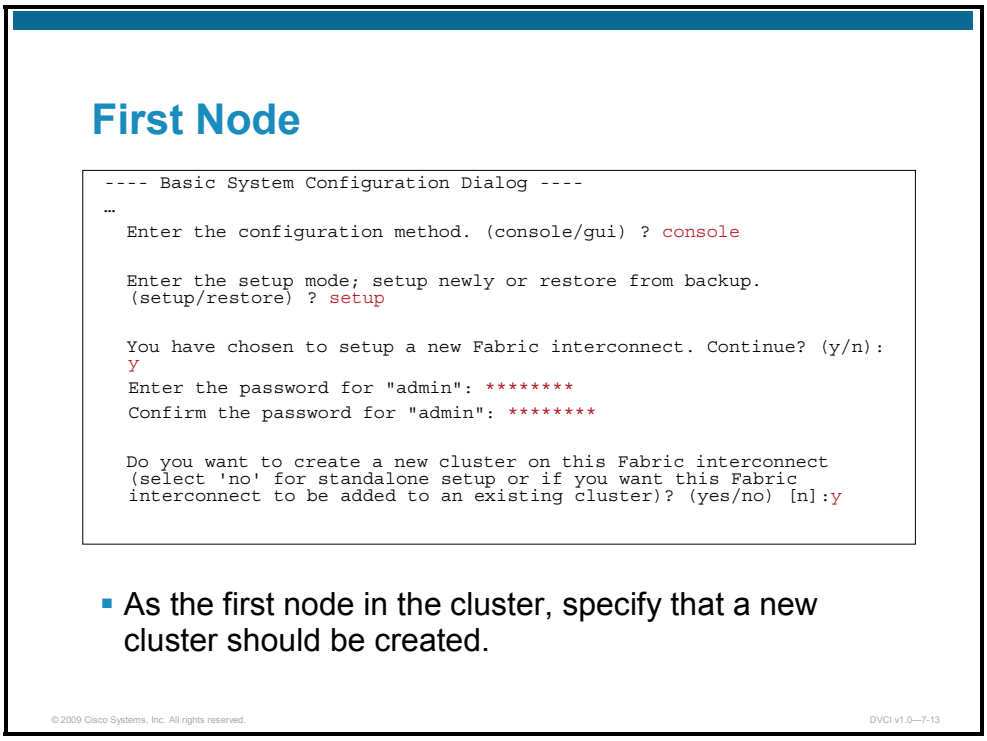

During initial configuration of the first node, specify that a new cluster is being created.

https://t.me/learningnets

First Node (Cont.)

```
Enter the switch fabric (A/B) []: A

 Enter the system name:  s6100

 Physical Switch Mgmt0 IPv4 address : 192.168.10.101

 Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

 IPv4 address of the default gateway : 192.168.10.254

 Cluster IPv4 address : 192.168.10.200

 Configure the DNS Server IPv4 address? (yes/no) [n]:
```

- Only the first node installed receives cluster-specific configuration.

© 2009 Cisco Systems, Inc. All rights reserved.

During the creation of a new cluster, the setup wizard requires a cluster IP address. This will act as the floating IP address that is used for accessing the UI, CLI, and API interfaces.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Second Node

```
 ---- Basic System Configuration Dialog ----
…
Enter the configuration method. (console/gui) [console] ? console
Installer has detected the presence of a peer Fabric interconnect. This
  Fabric interconnect will be added to the cluster. Continue (y/n) ? y

  Enter the admin password of the peer Fabric interconnect:******
Connecting to peer Fabric interconnect... done
    Retrieving config from peer Fabric interconnect... done
    Peer Fabric interconnect Mgmt0 IP Address: 192.168.10.101
    Peer Fabric interconnect Mgmt0 IP Netmask: 255.255.255.0
    Cluster IP address          : 192.168.10.200
```

- The second node will retrieve the Cisco UCS database and other configuration information, such as the cluster IP address, from the peer node.

© 2009 Cisco Systems, Inc. All rights reserved.

When the second node is brought online and properly cabled, the system configuration wizard will detect the presence of the peer node. Providing the password to the peer node allows the newly added node to log in and retrieve the configuration information it needs.

https://t.me/learningnets

Second Node (Cont.)

```
  Physical Switch Mgmt0 IPv4 address : 192.168.10.102
  Apply and save the configuration (select 'no' if you want to re-enter)?
  (yes/no): yes
  Applying configuration. Please wait.

  Configuration file - Ok


Cisco UCS 6100 Series Fabric Interconnect
s6100-B login:
```

- Only this node's specific IP address must be configured.
- All other database and configuration information have been retrieved from the peer.

© 2009 Cisco Systems, Inc. All rights reserved.

The only information required other than what is retrieved from the peer node will be this node's management IP address. If this node were being added to a previously configured cluster (such as in the case of replacing a failed node) this information would have been retrieved from the peer as well.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Managing High Availability

This topic describes how to manage high availability.

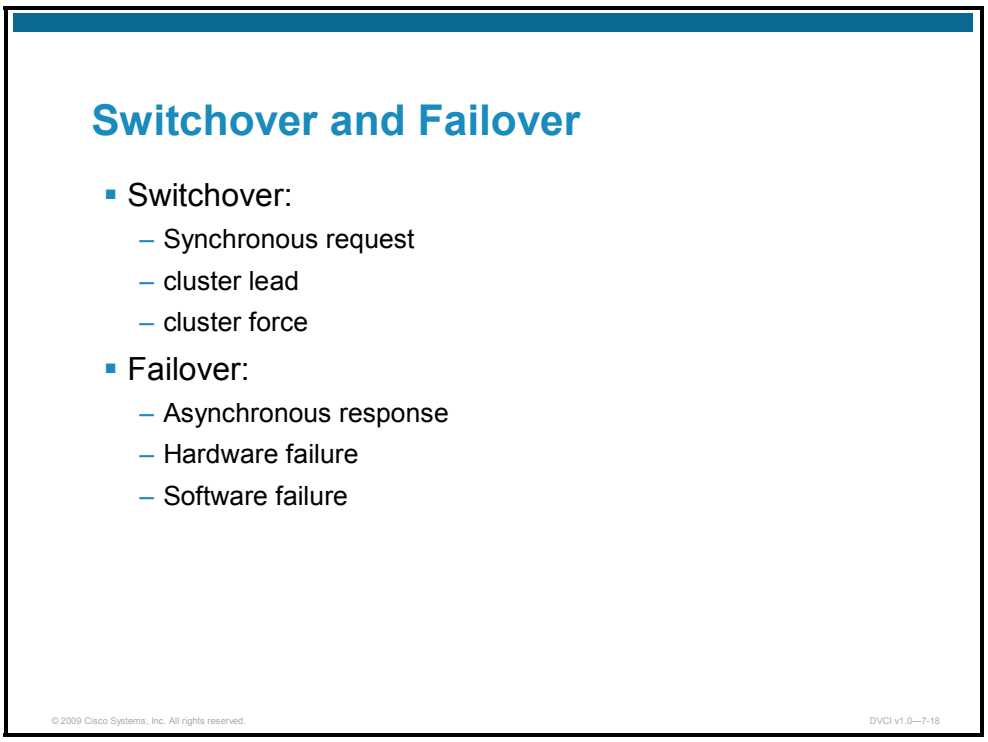

Switchover

A switchover is a synchronous request issued by the administrator to change the primary to subordinate, and the subordinate to primary. Use the **cluster lead** command (or cluster force, if necessary) to invoke this change within the local-mgmt interface.

Failover

A failover is an asynchronous response by the Cisco UCS Manager system (specifically, the system controller) to promote the subordinate node to the primary role. This can be a response to a software failure (in other words, individual Cisco UCS Manager process(es) failing on the primary node) or gross hardware failure, such as a catastrophic switch failure or split brain.

https://t.me/learningnets

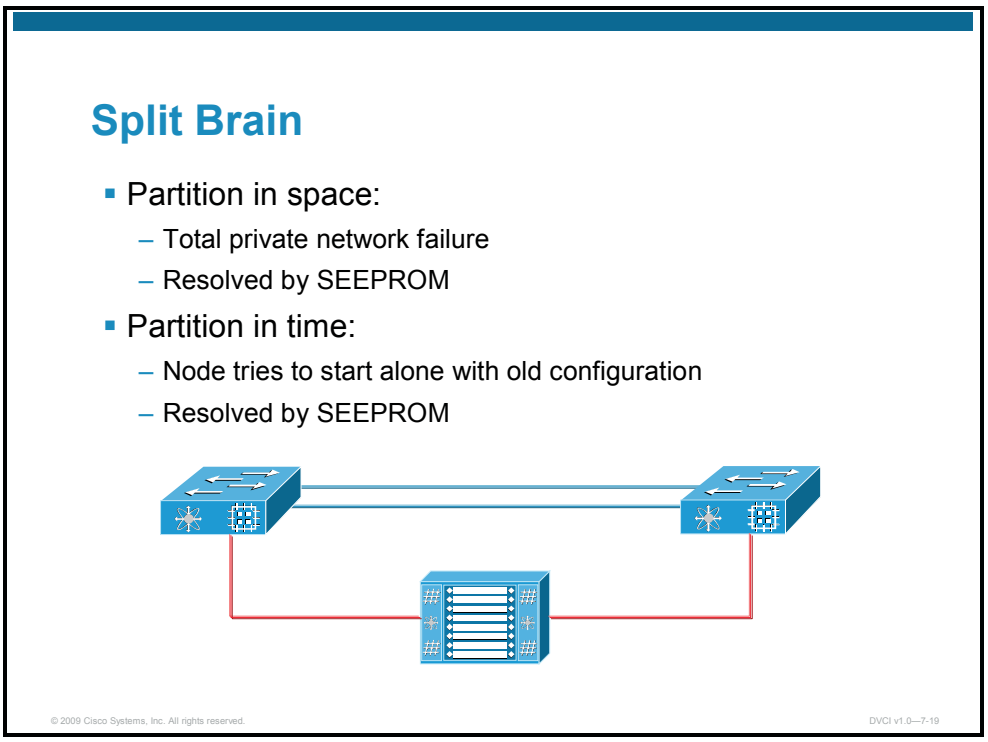

Partition in Space

A “partition in space” occurs when nodes fail to communicate with each other over the private network (L1 and L2 links both failed). To resolve this split brain (and assuming both switches are active at the time of the private network failure), each CMC will act on the behalf of its fabric’s Cisco UCS Manager instance to reach the SEEPROM first and write its node ID in the primary field. This is known as a “quorum race”. The winner will remain in the cluster and the loser will abort. When the links are restored, the losing node can rejoin the cluster and act as subordinate.

Partition in Time

A “partition in time” occurs when one of the nodes is down for a period of time, during which changes to the configuration are made on the active primary node. These changes would not have been replicated to the down node.

If the primary node were to shut down after having made configuration changes to the database, but prior to being able to replicate them to the other (downed) node, and that downed node were to try to join the cluster alone, then that condition is referred to as a “partition in time”.

To resolve this split brain, a version number that represents the configuration is written to the EEPROM. On solo startup, a node compares its version to that of the other node (recall that both nodes can read both parts of the EEPROM). If the version is the same or higher, it can start the cluster.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Checking Cluster Status

```
…
s6100-A# show cluster state
Cluster Id: 0xf867f6ec8f5b11de-0x9c99000decb21744

A: UP, PRIMARY
B: UP, SUBORDINATE

HA READY
S6100-A#
```

- The **show cluster state** command can quickly provide simple status information about the cluster.

© 2009 Cisco Systems, Inc. All rights reserved.

The **show cluster state** command can quickly provide information about the health of the cluster, including which node (A or B) is currently acting as the primary node.

https://t.me/learningnets

Checking Extended Cluster Status

```
s6100-A# show cluster extended-state
Cluster Id: 0xf867f6ec8f5b11de-0x9c99000decb21744
Start time: Sat Aug 22 20:42:38 2009
Last election time: Sat Aug 22 20:49:47 2009
A: UP, PRIMARY
B: UP, SUBORDINATE
```

- The **show cluster extended-state** command provides the same information as **show cluster state**, plus additional information (shown on following slide).

© 2009 Cisco Systems, Inc. All rights reserved.

The first portion of the **show cluster extended-state** command is the same as the previously shown **show cluster state** command.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Checking Extended Cluster Status (Cont.)

```
…
A: memb state UP, lead state PRIMARY, mgmt services state: UP
B: memb state UP, lead state SUBORDINATE, mgmt services state: UP
   heartbeat state PRIMARY_OK

INTERNAL NETWORK INTERFACES:
eth1, UP
eth2, UP

HA READY
Chassis detailed state:
Chassis, serial: FOX1307H0M8, state: active
s6100-A#
```

- The **show cluster extended-state** provides additional information that may be useful in troubleshooting cluster errors.

© 2009 Cisco Systems, Inc. All rights reserved.

The additional information provided by **show cluster extended-state** includes the status of the individual cluster links (shown here as eth1 and eth2, corresponding to L1 and L2 respectively), as well as information about the chassis used to complete the high-availability configuration.

https://t.me/learningnets

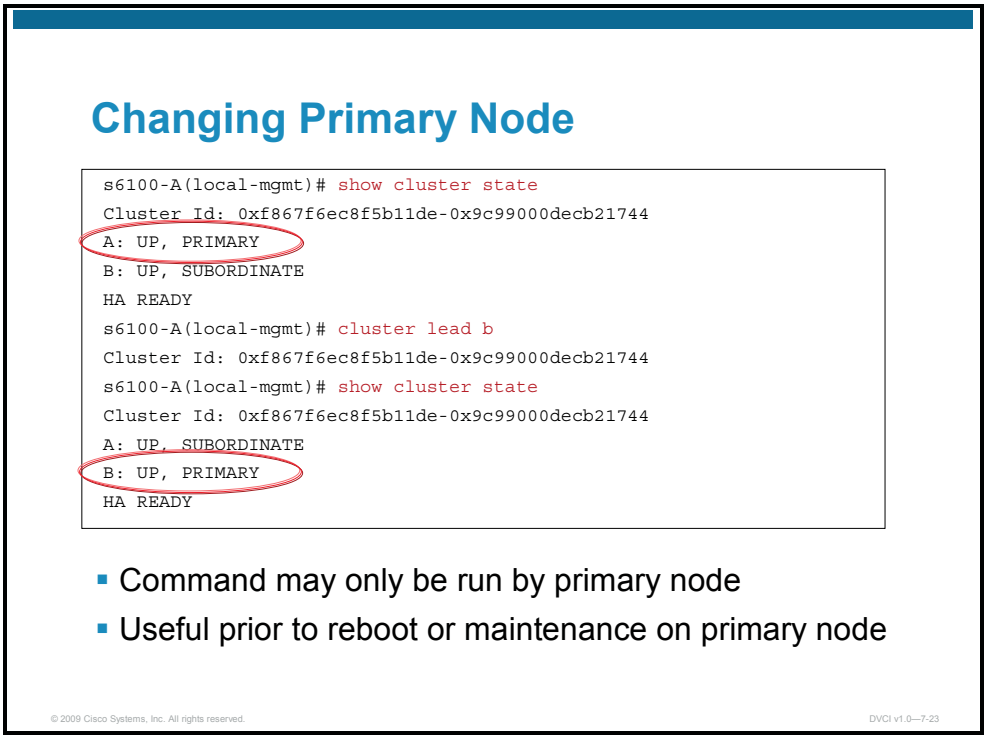

To manually failover the primary node, run the **cluster lead** command from the local-mgmt CLI on the primary node. Only the primary node may run this command.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Forcing Primary Node

```
s6100-A(local-mgmt)# show cluster state
Cluster Id: 0xf867f6ec8f5b11de-0x9c99000decb21744

A: UP, SUBORDINATE
B: UP, PRIMARY

HA READY
s6100-A(local-mgmt)# cluster force primary
Cluster Id: 0xf867f6ec8f5b11de-0x9c99000decb21744
request failed: cannot accept  force command when election has successfully
  completed
s6100-A(local-mgmt)#
```

- Command will only function when election process has failed, for example due to a partition in time issue.

© 2009 Cisco Systems, Inc. All rights reserved.

To force the local node to become primary in the case of a partitioned cluster (in which case neither node is primary), use the **cluster force primary** command on the node to become primary. This example shows the failure that occurs if the election process was already successful.

https://t.me/learningnets

Summary

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

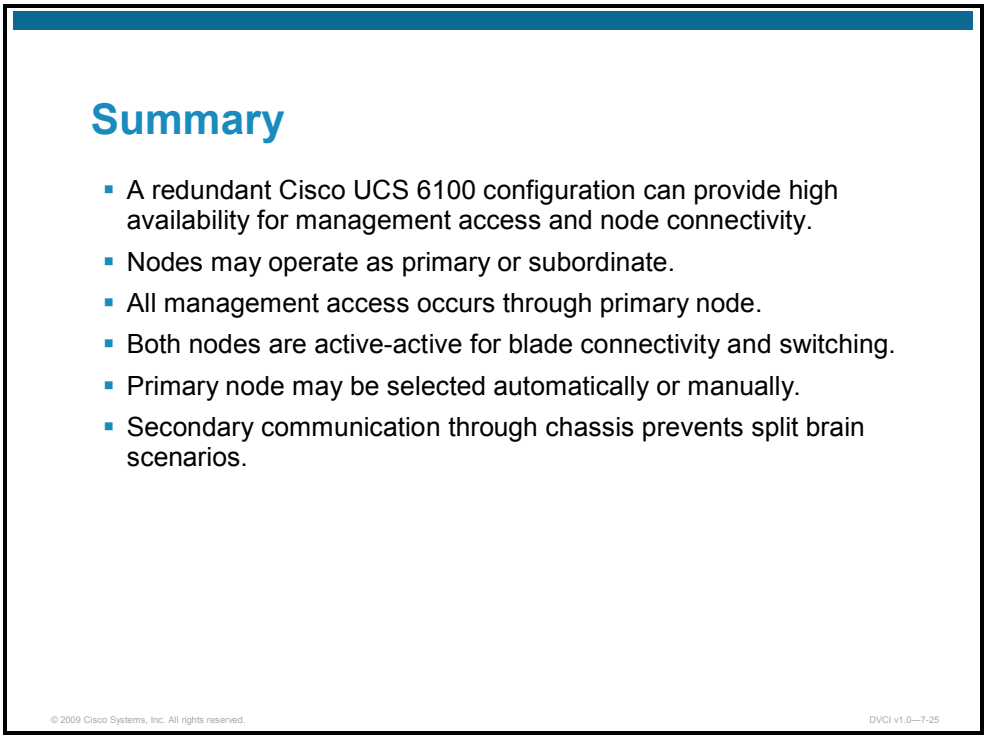

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Monitoring System Events

Overview

This topic describes how to monitor system events.

Objectives

Upon completing this lesson, you will be able to describe how to monitor system events. This includes being able to meet these objectives:

- Describe event and failure monitoring using the GUI
- Describe options for monitoring Cisco UCS remotely
- Describe the monitoring of Cisco UCS from the CLI
- Describe Cisco UCS Threshold Policies
- Describe configuration of Call Home Policies
- Describe data export capabilities of the Cisco UCS Manager
- Describe the Ethanalyzer feature

https://t.me/learningnets

GUI Monitoring

This topic describes event and failure monitoring using the Cisco UCS Manager GUI.

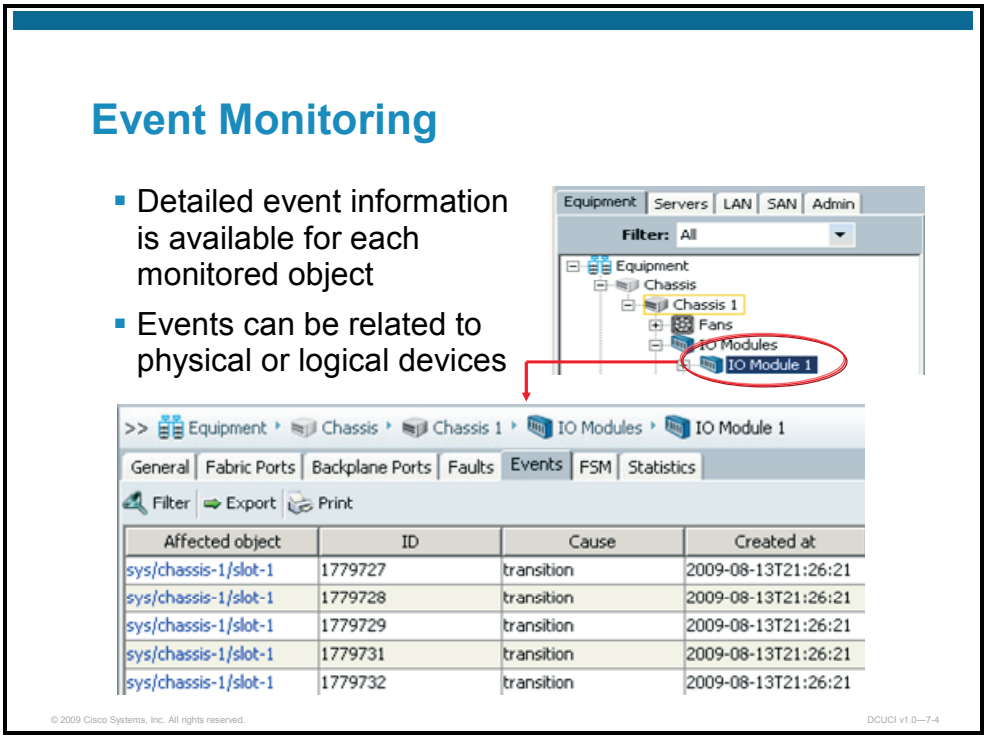

Cisco UCS Manager tracks many events and conditions related to every monitored object within the Cisco Unified Computing System (Cisco UCS) environment. This figure shows an example of events for an I/O module (IOM).

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Event Monitoring (Cont.)

- Events can also be viewed on a system-wide basis from the Admin tab.
- Double-clicking any event will open a detailed description of the event.

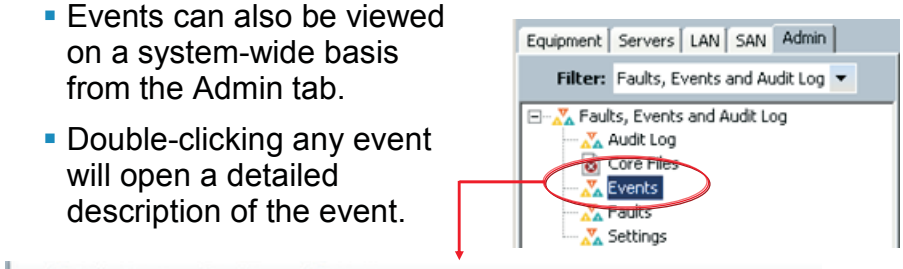

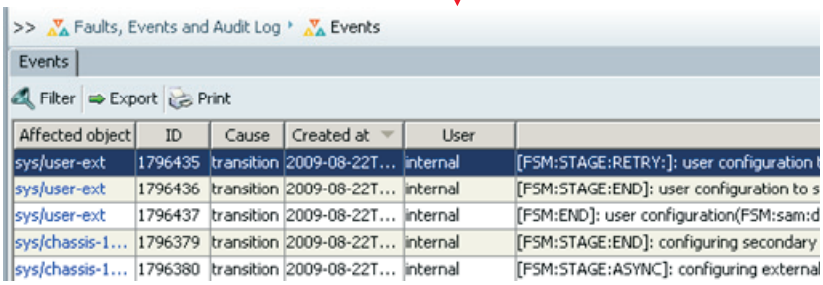

© 2009 Cisco Systems, Inc. All rights reserved.

Events can be viewed on a system wide basis from the Admin tab.

https://t.me/learningnets

Event Export

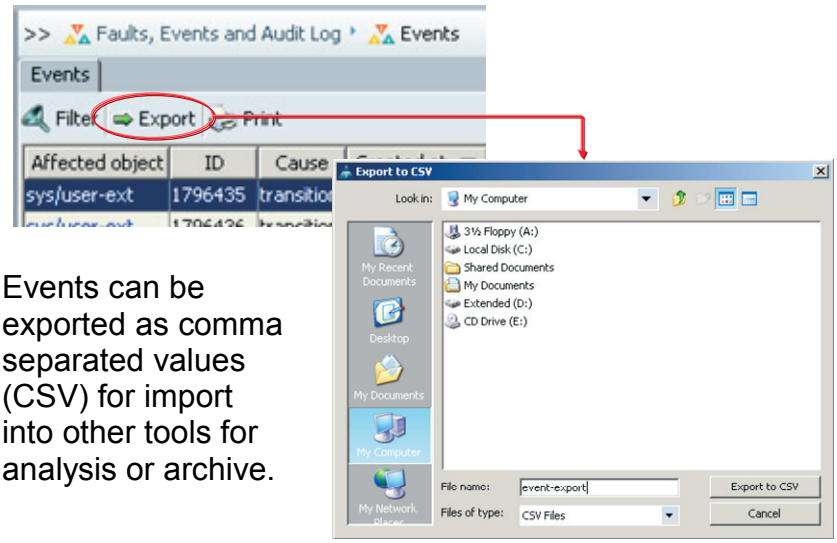

- Events can be exported as comma separated values (CSV) for import into other tools for analysis or archive.

© 2009 Cisco Systems, Inc. All rights reserved.

Events can be exported for use by other tools or for archiving.

https://t.me/learningnets

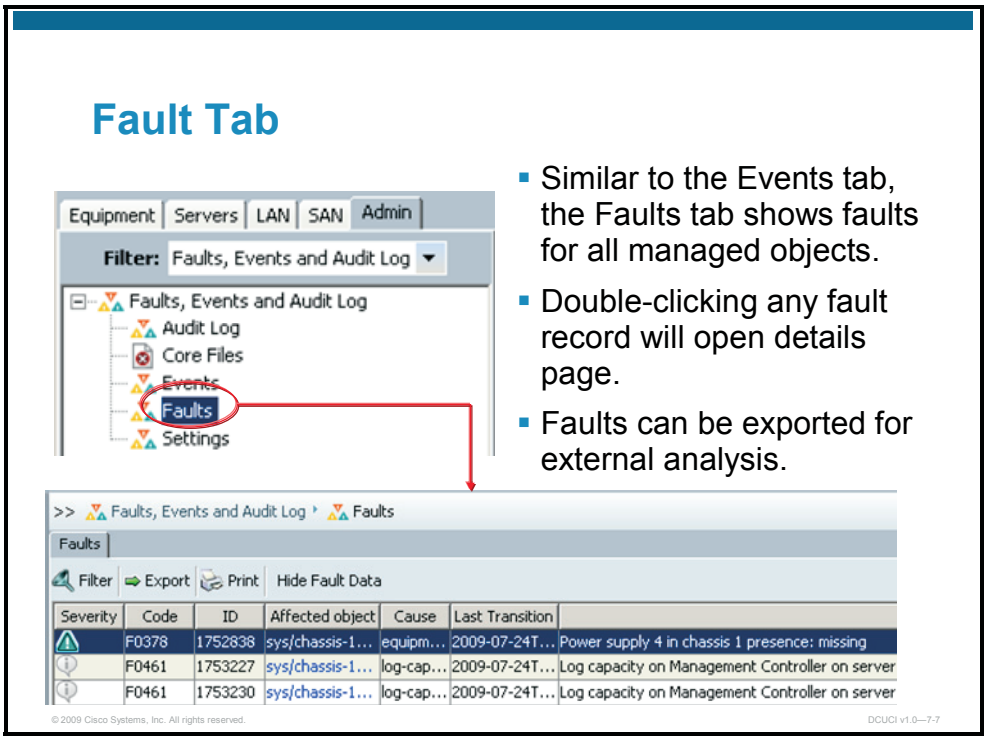

The GUI Faults tab provides information on system faults. These faults can either be defaults or user-defined.

https://t.me/learningnets

Fault Details

- In some cases, faults may be acknowledged to prevent repeated notifications.
- The detail page shows the specific reason for failure, timestamps, and other detailed information about the fault.

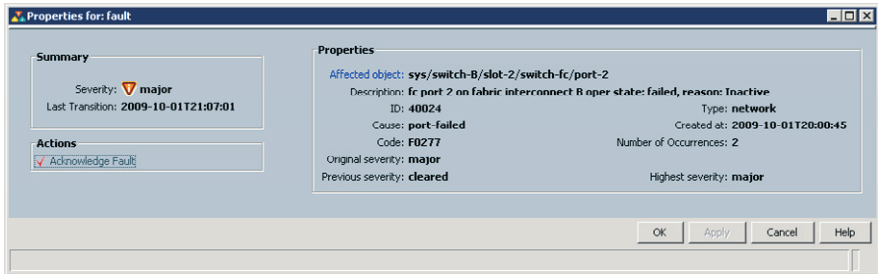

© 2009 Cisco Systems, Inc. All rights reserved.

Fault reporting occurs for on-board diagnostics, Cisco UCS Manager failures, chassis management controller (CMC) logging, baseboard management controller (BMC) logging, and blade Power On Self Test (POST) and operating system (OS) events. Faults are managed objects within the Cisco UCS system. Various fault types exist within Cisco UCS:

- Hardware:
  - GPIO pins
  - I2C requests
- Network:
  - Link failure
  - Ping test
- Software:
  - System manager
  - Core dump

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Remote Monitoring

This topic describes configuration of various protocols for remote monitoring of Cisco UCS.

External Status Monitoring

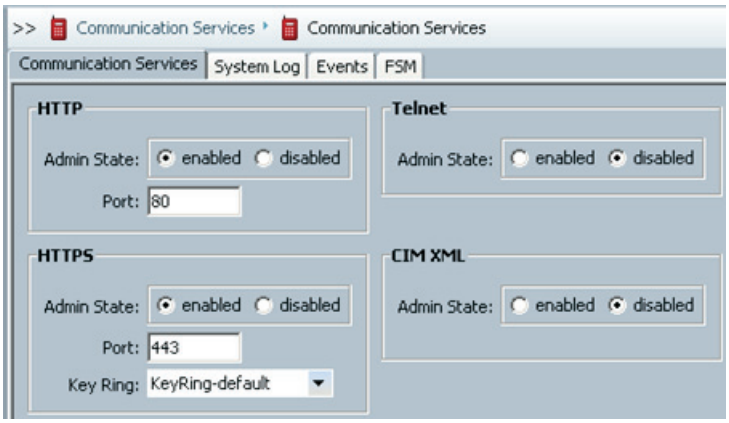

- Multiple protocols are supported for monitoring the Cisco UCS Manager environment from external entities.

© 2009 Cisco Systems, Inc. All rights reserved.

Various export methods exist within Cisco UCS to tie internal events to external systems. HTTP, HTTPS, Extensible Markup Language (XML) and Simple Network Management Protocol (SNMP) are among the most common monitoring protocols used for system management.

https://t.me/learningnets

SNMP Monitoring

- Cisco UCS Manager can send SNMP traps to external monitoring tools
- SNMP user accounts can be created for monitoring tools that will pool Cisco UCS for events
- SNMP v3 is supported, includes same MIBs as Cisco Nexus 5000
- Does not provide MIBs downstream of Fabric Interconnect

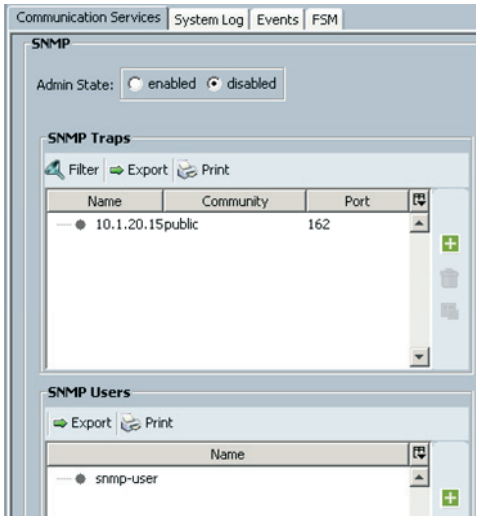

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco UCS Manager GUI provides robust monitoring and reporting features which can also be used to export data to external status monitoring tools such as SNMP-based software or syslog servers.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

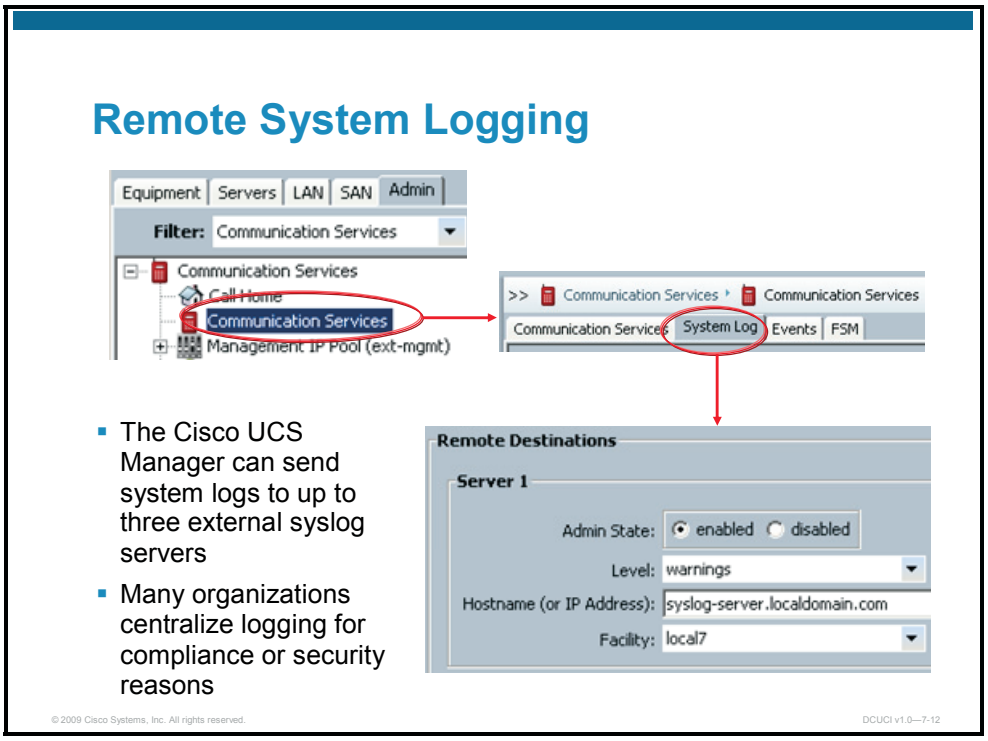

The Cisco UCS Manager can send system logs to up to three remote syslog servers. Each can be configured with different alert levels and remote facility names.

https://t.me/learningnets

CLI Status Monitoring

This topic describes Cisco UCS status monitoring from the CLI.

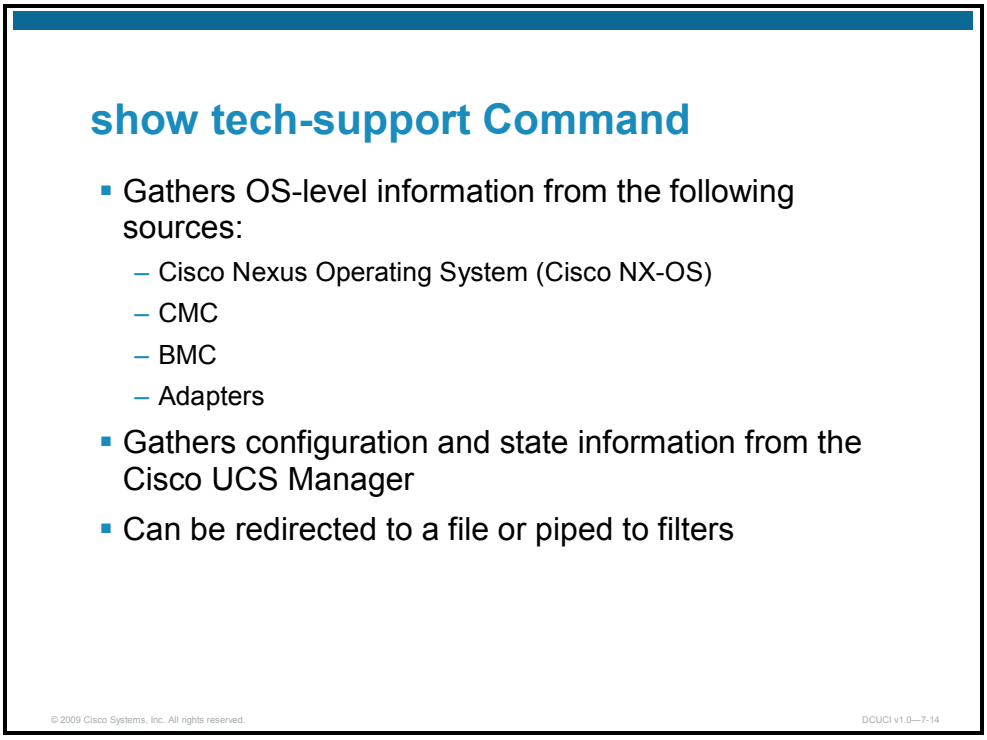

The **show tech support** commands provide detailed information for use in troubleshooting system issues as well as gathering information for the Cisco Technical Assistance Center (Cisco TAC).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

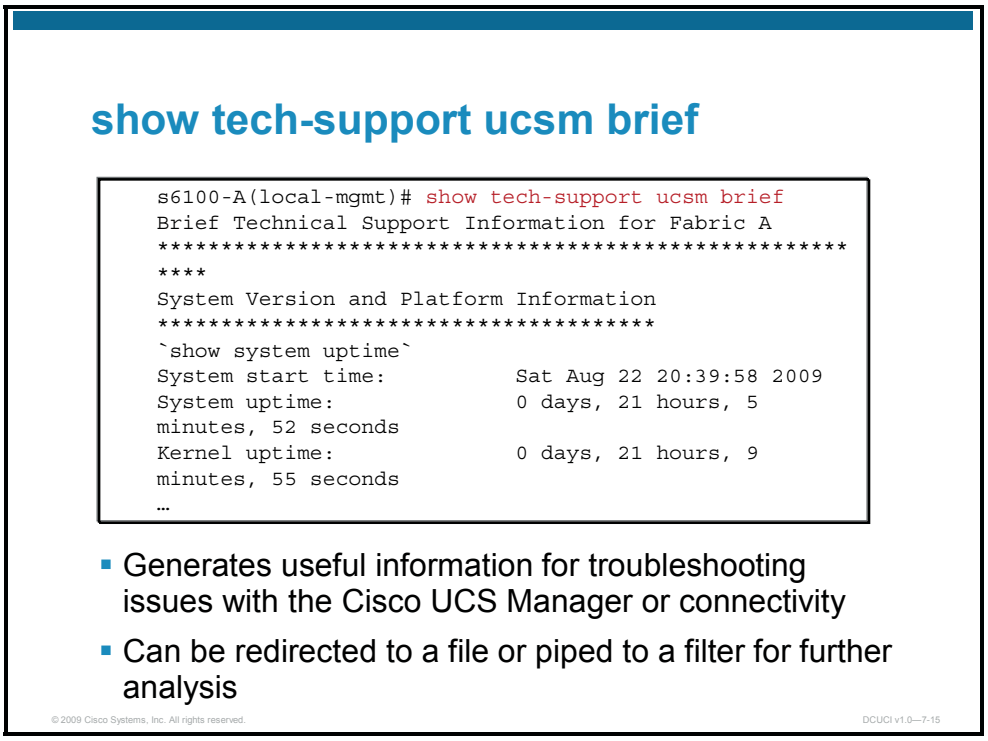

The figure shows an example of the output form the **show tech-support ucsm brief** command.

https://t.me/learningnets

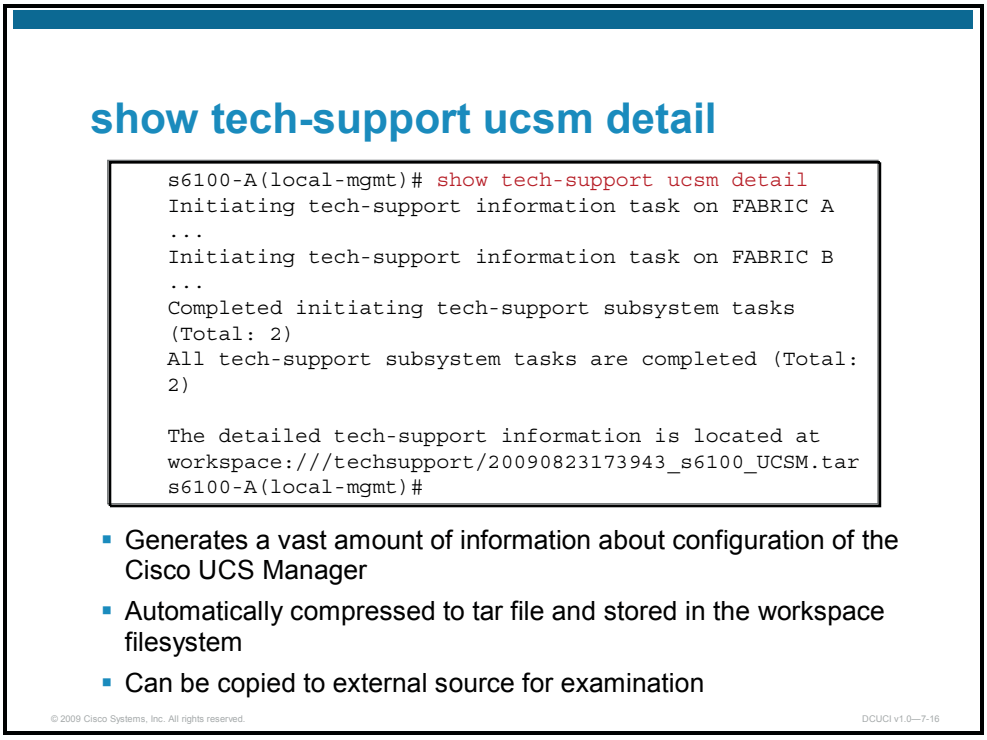

The figure shows an example of the output from the **show tech-support ucsm detail** command. The created file can be copied with any supported protocol (FTP, Secure Copy Protocol [SCP], Secure FTP [SFTP], and so on) to an external host for analysis or transfer to TAC.

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

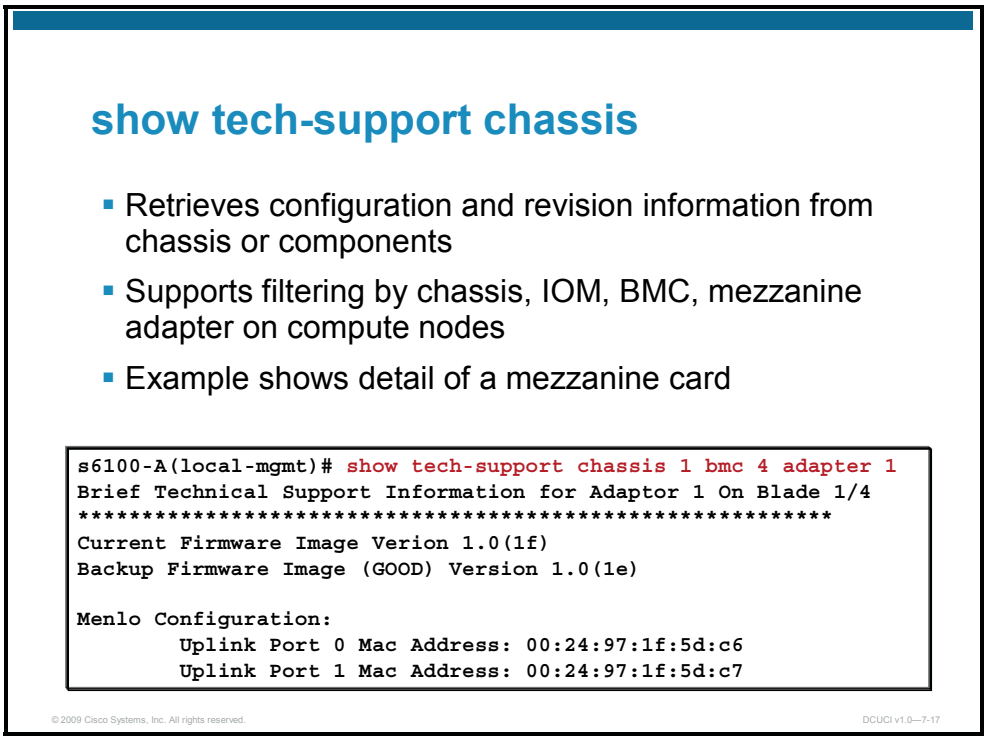

The figure shows an example of the output from the **show tech-support chassis** command. In this example, the output is filtered to include only the first mezzanine adapter on the fourth compute node of the first chassis. Other options include:

`show tech-support chasiss x`

- Shows all information about chassis x including all subcomponents

`show tech-support chassis x iom y`

- Shows all information about IOM y in chassis x

`show tech-support chassis x bmc y`

- Shows all information about compute node y in chassis x

`show tech-support chassis x bmc y adapter z`

- Shows all information about mezzanine adapter z in compute node y in chassis x

https://t.me/learningnets

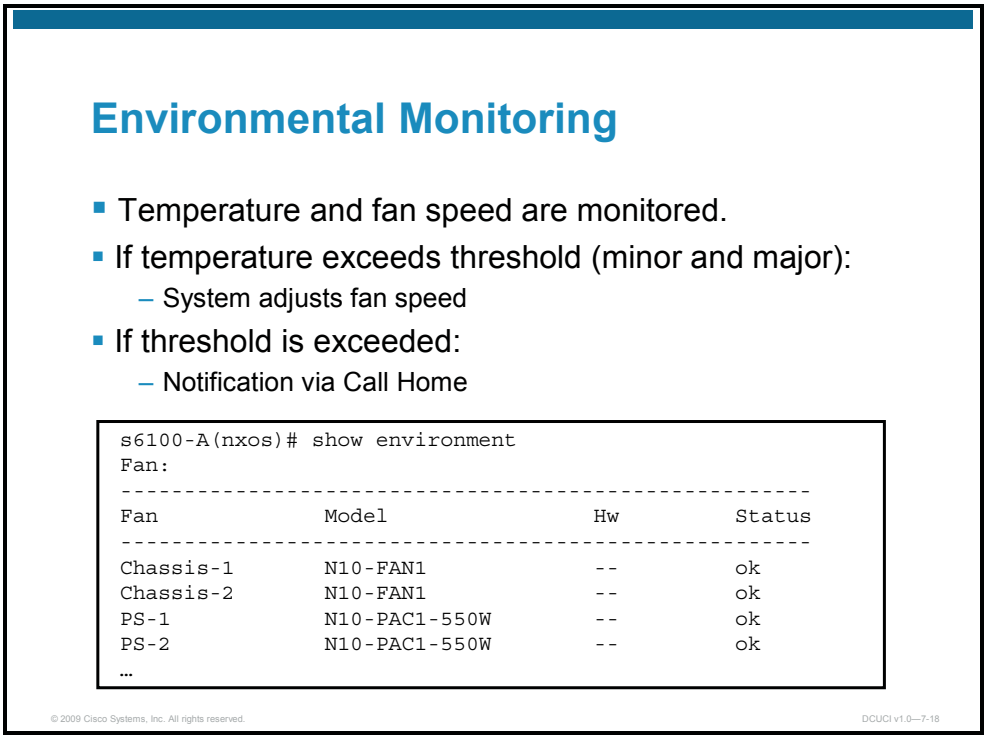

The figure summarizes the environmental monitoring provided by the Fabric Interconnect. The **show environment** command from within the Cisco NX-OS CLI displays the following environmental conditions:

- Fan
- Temperature
- Power Supply
- Power Usage Summary

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

Threshold Policies

This topic describes Cisco UCS Threshold Policies.

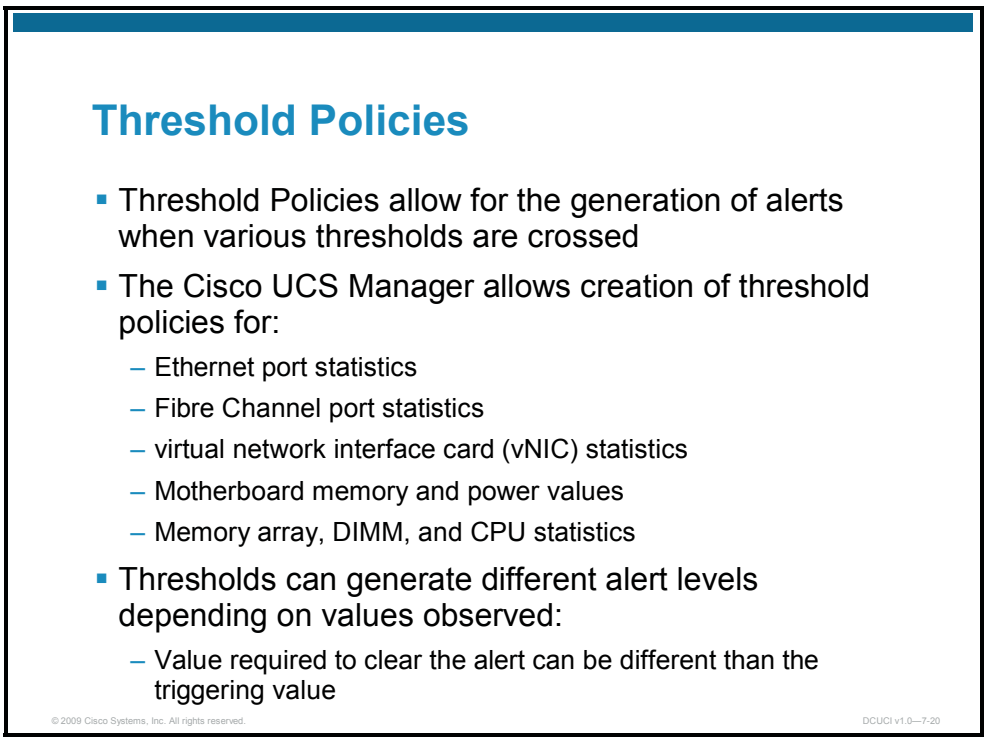

Threshold Policies allow administrators to configure alerts based on the observed statistics of various system components.

https://t.me/learningnets

Threshold Policy Wizard

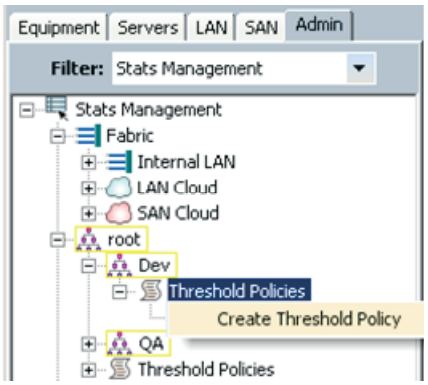

- To create a threshold policy, select the appropriate organization object, and right click **Threshold Policies**
- Click **Create Threshold Policy**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the correct location for creating a Threshold Policy within an organization.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Threshold Policy Wizard (Cont.)

- Provide a name for the Threshold Policy and an optional description

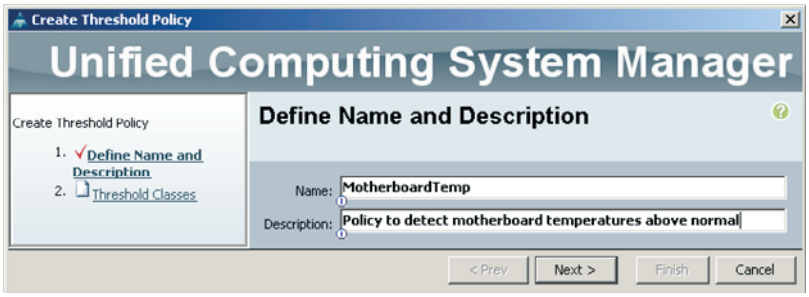

© 2009 Cisco Systems, Inc. All rights reserved.

Threshold Policies allow administrators to create notification events for various system properties.

https://t.me/learningnets

Threshold Policy Wizard (Cont.)

- Click **Add** to add a threshold class to the new policy.
- Note that multiple threshold classes can be added to the same policy.

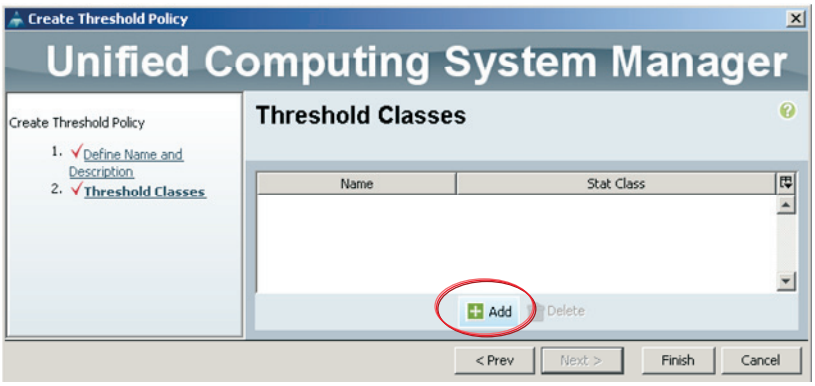

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the location of the Add button for adding a Threshold Class to the new policy.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Threshold Policy Wizard (Cont.)

- Choose the appropriate Statistics Class.
- Each Statistics Class will have various threshold definitions beneath it.

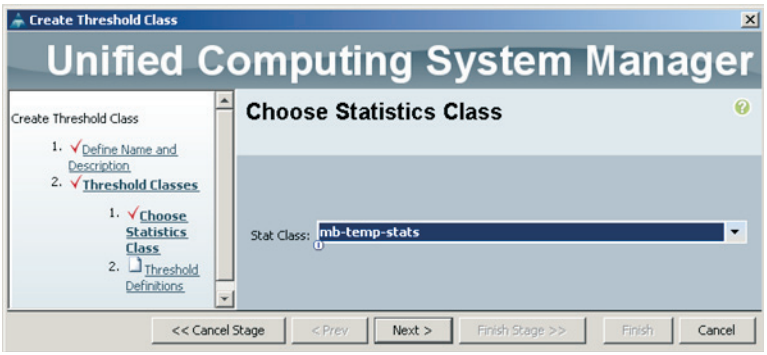

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the selection of a Statistics Class for this policy definition.

https://t.me/learningnets

Threshold Policy Wizard (Cont.)

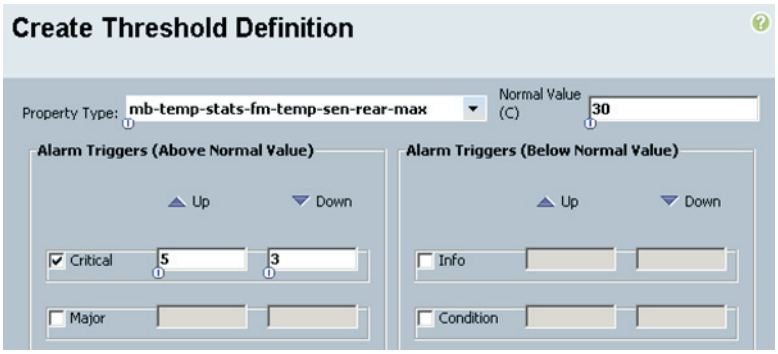

- This example shows the creation of a Threshold Definition for the rear temperature sensor on the motherboard
- A normal value of 30 degrees has been specified, with the Threshold Policy triggering at a critical alert level if the temperature exceeds 5 degrees over normal

© 2009 Cisco Systems, Inc. All rights reserved.

Note that for clarity, not all options are shown in the figure. Each of the alert levels (Critical, Major, Minor, Warning, Condition, and Info) can be configured independently in both the "Above Normal Value" and "Below Normal Value" columns.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Applying Threshold Policy

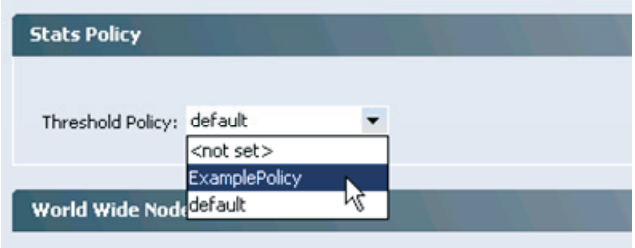

- Applying a threshold policy can be done during creation of a service profile, or after the profile is created.
- This figure shows selection of a policy from a previously created service profile
- To reach this setting, select the service profile in the Servers tab of UCS Manager, and select the Policies tab in the content pane

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the process for selecting a threshold policy for a previously defined service profile.

https://t.me/learningnets

Applying Threshold Policy (Cont.)

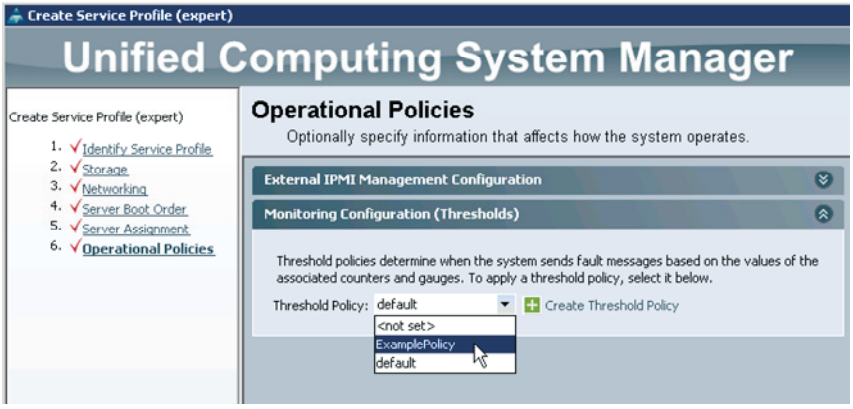

- This figure shows selection of a threshold policy during the creation of a service profile using the expert service profile wizard

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the process for selecting a threshold policy during the creation of a service profile using the expert service profile wizard.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Call Home Configuration

This topic describes the configuration of the Call Home feature within Cisco UCS.

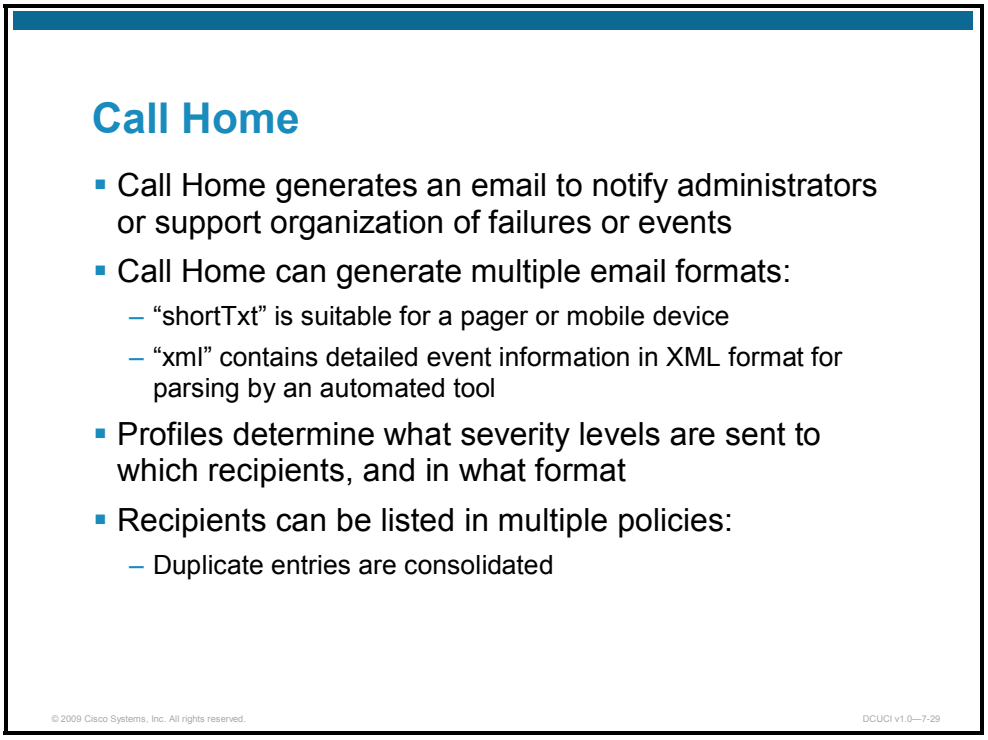

The Cisco UCS Manager GUI provides robust monitoring and reporting features which can also be used to export data to external status monitoring tools such as SNMP-based software or syslog servers.

https://t.me/learningnets

Contact Configuration

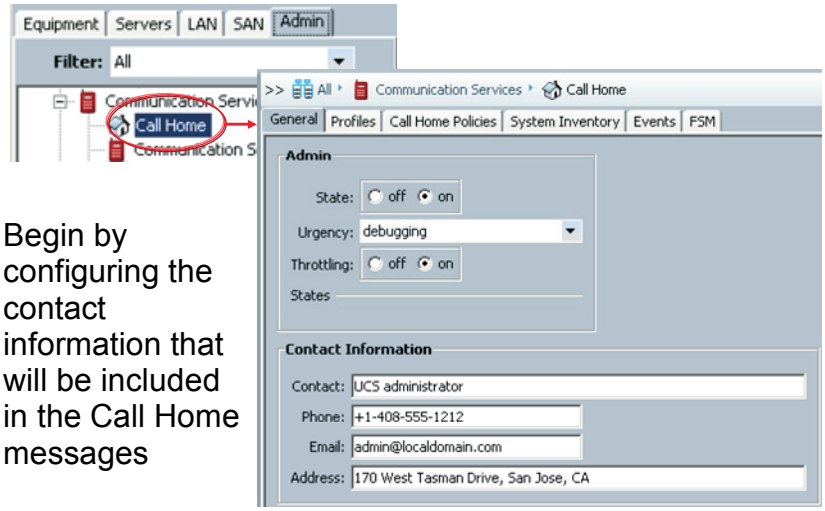

- Begin by configuring the contact information that will be included in the Call Home messages

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows Call Home setup from the Cisco UCS Manager GUI.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

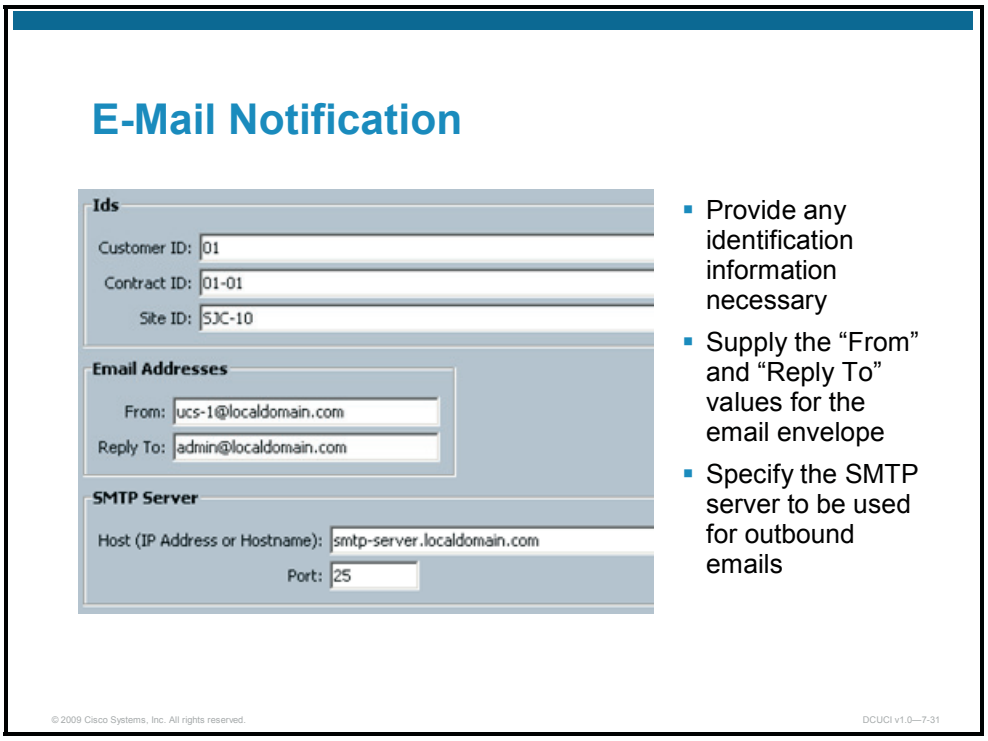

This figure shows Call Home setup from the Cisco UCS Manager GUI (continued).

https://t.me/learningnets

Call Home Profiles

- Three default profiles exist:
  - CiscoTAC-1 – useful for sending XML data to Cisco TAC:
    - Uses special “ciscoTac” alert group to collect information needed by Cisco TAC
  - full_txt – for sending full XML data:
    - By default includes all alert groups at warning severity
  - short_txt – for sending short, plain text data:
    - By default includes all alert groups at warning severity
- Alert groups allow profiles to send only data related to specific functional areas
- Additional profiles can be created as needed:
  - Default profiles can be modified but not removed

© 2009 Cisco Systems, Inc. All rights reserved.

Call Home profiles determine the data, and the format of that data, that is sent upon generation of a Call Home event.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Profile Configuration

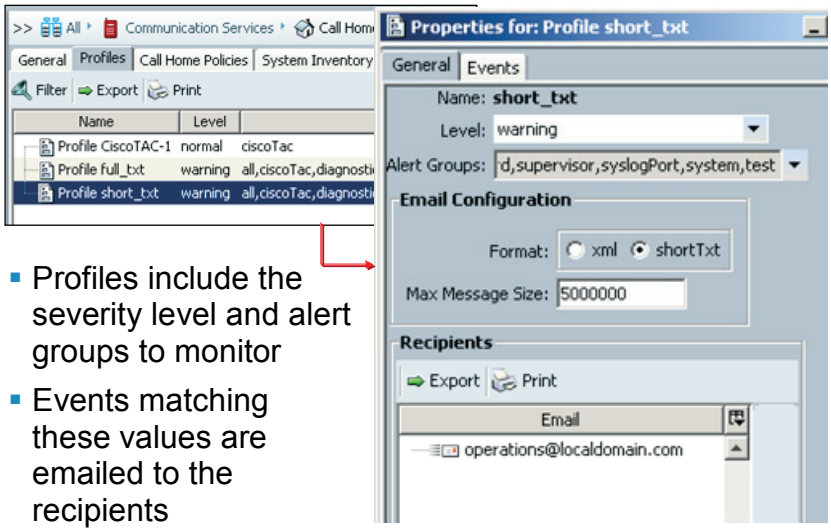

- Profiles include the severity level and alert groups to monitor
- Events matching these values are emailed to the recipients

© 2009 Cisco Systems, Inc. All rights reserved.

Call Home profiles define the alert groups and notification levels, and the users to receive the emails. The figure shows the default profile “short_txt”, as it is configured to send emails to one recipient.

https://t.me/learningnets

Call Home Policies

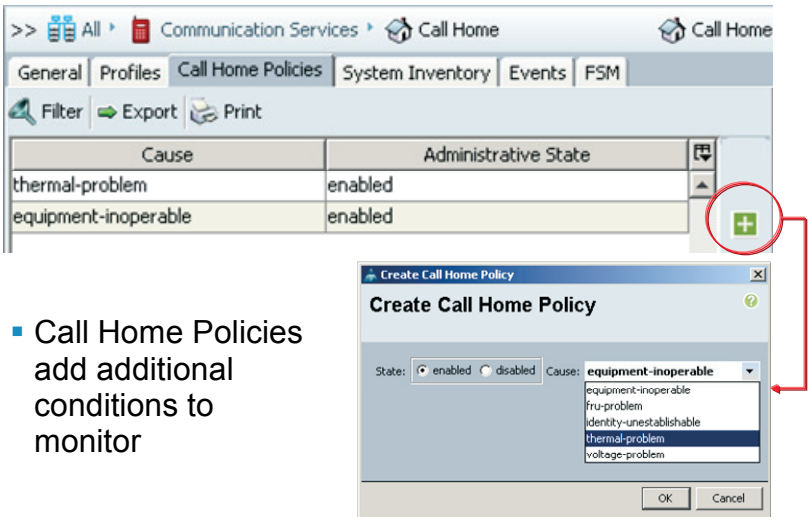

- Call Home Policies add additional conditions to monitor

© 2009 Cisco Systems, Inc. All rights reserved.

Call Home Policies can also be created to add additional conditions to be monitored.

https://t.me/learningnets

System Inventory

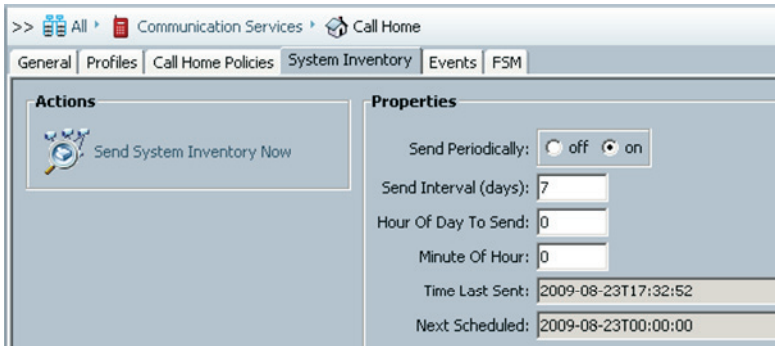

- System Inventory information can be sent manually or scheduled periodically.
- This aids support organizations to track changes to installed equipment.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows System Inventory notification setup.

https://t.me/learningnets

Exporting Data

This topic describes the data export capabilities of the Cisco UCS Manager.

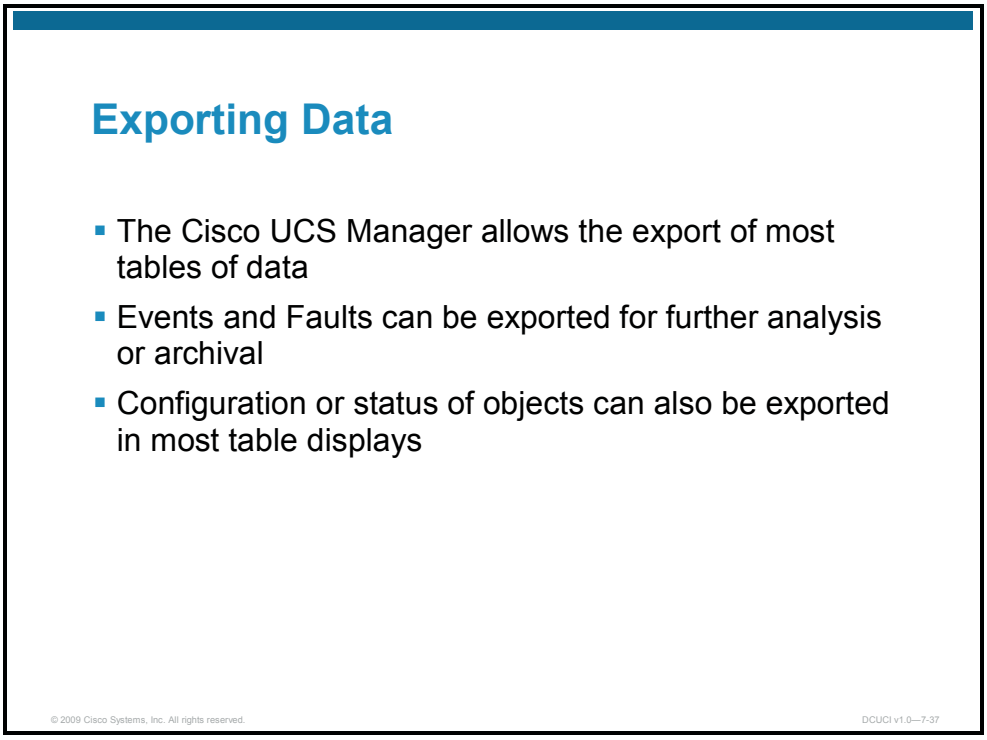

The Cisco UCS Manager provides a facility to export most table data to comma separated value (CSV) files.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

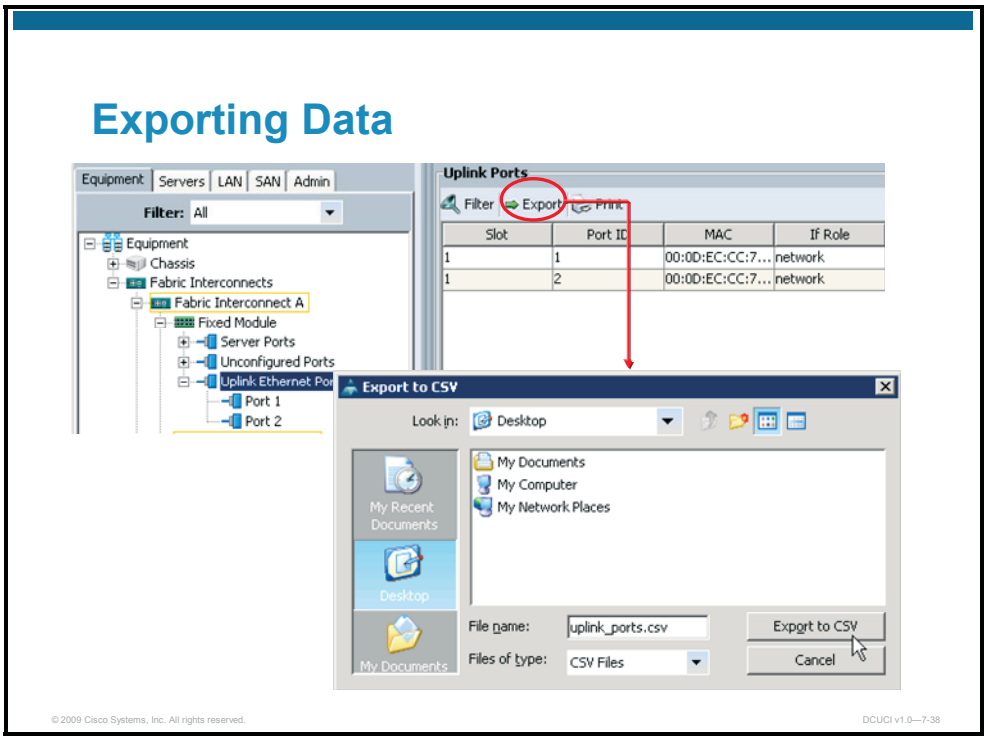

The figure shows the exporting of a table of data to a comma separated file.

https://t.me/learningnets

Viewing Exported Data

- CSV files exported from the Cisco UCS Manager can be used by any application that supports the format
- The example from the previous slide is shown here, after loading into Excel

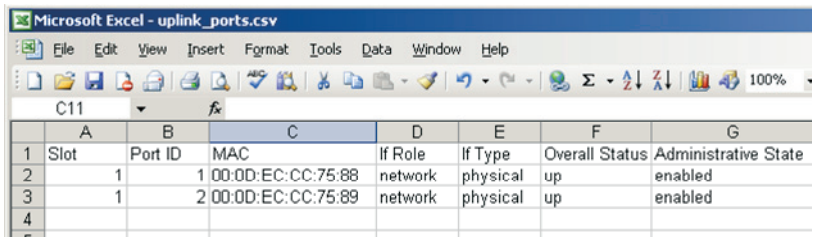

	A	B	C	D	E	F	G
1	Slot	Port ID	MAC	If Role	If Type	Overall Status	Administrative State
2	1	1	00:0D:EC:CC:75:88	network	physical	up	enabled
3	1	2	00:0D:EC:CC:75:89	network	physical	up	enabled
4							

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows an exported report viewed from Microsoft Excel.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Ethanalyzer

This topic describes the Cisco NX-OS Ethanalyzer protocol analyzer.

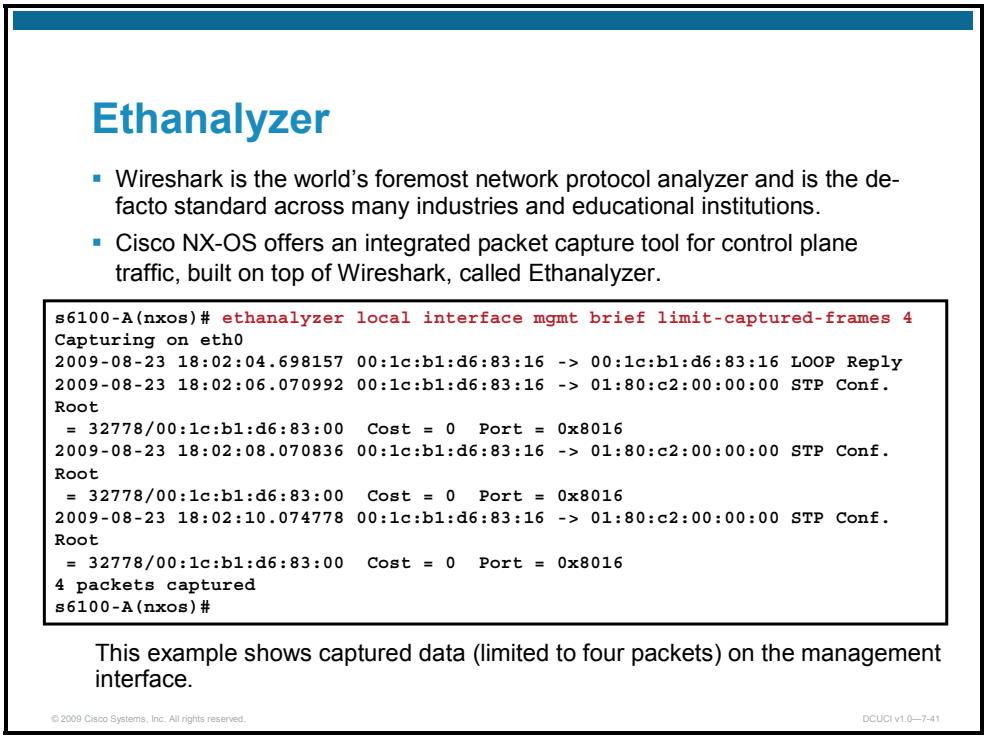

Ethanalyzer is part of the Cisco NX-OS integrated manageability toolset:

- Ethanalyzer improves troubleshooting, thus reducing time-to-resolution.
- Ethanalyzer helps preserve and improve the operational continuity of the network infrastructure.
- Ethanalyzer does this by allowing network administrators to learn about the nature of the control plane traffic flowing through their network.

https://t.me/learningnets

Relationship to Wireshark

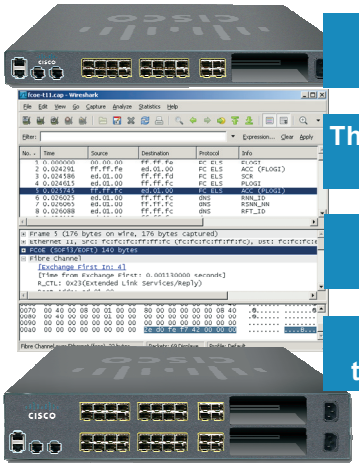

NX-OS runs on top of a Linux Kernel.

The Linux Kernel supports packet capturing using the libpcap library.

Wireshark decodes packets captured through the libpcap library.

Ethanalyzer is a wrapper over TShark, the terminal-based version of Wireshark.

© 2009 Cisco Systems, Inc. All rights reserved.

Features supported in Cisco NX-OS Release 4.0 include:

- Capture packets sent or received by the Supervisor
- Set the number of packets to be captured
- Set the length of the packets to be captured
- Display packets with very detailed protocol information, as well as with a summarized single line of information
- Open and Save packet data captured
- Filter packets captured on many criteria
- Filter packets to be displayed on many criteria

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

Summary

- The Cisco UCS platform offers various tools for system monitoring and reporting.
- Notification can be set up up in various formats based on user defined thresholds as well as system default thresholds.
- Cisco UCS provides export ability for configuration, status, and event data.
- Threshold Policies allow administrators to configure alerts based on the observed statistics of various system components.
- Call Home generates an email to notify administrators or support organization of failures or events.
- The Cisco UCS Manager provides a facility to export most table data to comma separated value (CSV) files.
- Ethanalyzer improves troubleshooting, thus reducing time-to-resolution.

© 2009 Cisco Systems, Inc. All rights reserved.

https://t.me/learningnets

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

Module Summary

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

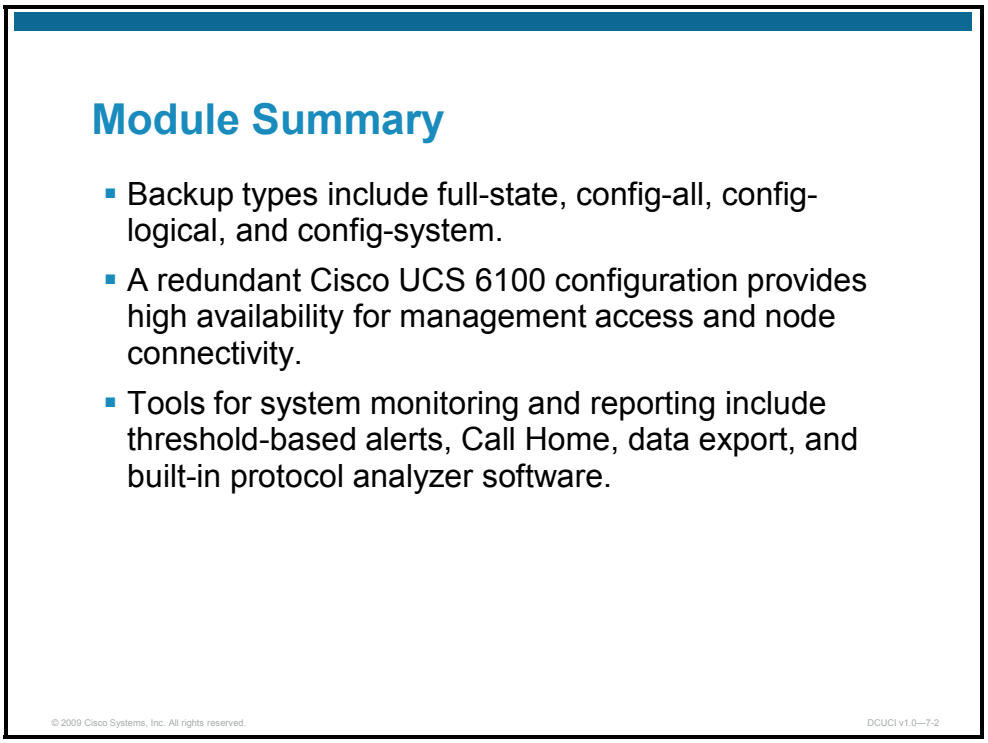

https://t.me/learningnets

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

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) Which of the following is a valid backup type? (Source: Backup and Restore)

A) config-lan
B) config-iom
C) full-config
D) full-state

Q2) Which command, when run from active node A, will cause the cluster to fail over to the subordinate node B? (Source: Managing High Availability)

A) cluster force
B) cluster failover
C) cluster lead b
D) cluster resign

Q3) Cisco UCS Manager allows data to be exported in what format? (Source: Monitoring, Diagnostics, and Troubleshooting)

A) Excel (XLS)
B) Rich-Text Format (RTF)
C) Comma-Separated Values (CSV)
D) Tab-Delimited File (TDF)

https://t.me/learningnets

Module Self-Check Answer Key

Q1) D

Q2) C

Q3) C

 © 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 8

Cisco Unified Computing System Deployments

Overview

This module describes use cases and methods of deploying the Cisco Unified Computing System (Cisco UCS).

Module Objectives

Upon completing this module, you will be able to describe use cases and methods of deploying the Cisco Unified Computing System. This includes being able to meet these objectives:

- Describe various use cases for the Cisco Unified Computing System
- Describe operating system installation methods

https://t.me/learningnets

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

Lesson 1

UCS Use Cases

Overview

In this lesson, you will explore various use cases for the Cisco Unified Computing System (Cisco UCS).

Objectives

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

- Describe the server operations that Cisco UCS can enhance
- Describe the changes in VMware virtualization and the four different use case scenarios with Cisco UCS and VMware
- Describe the different types of database solutions architectures and explain how the Cisco UCS system is ideal for these types of solutions

https://t.me/learningnets

Use Case: Server Operations

In this topic you will learn what types of server operations the Cisco UCS system can enhance.

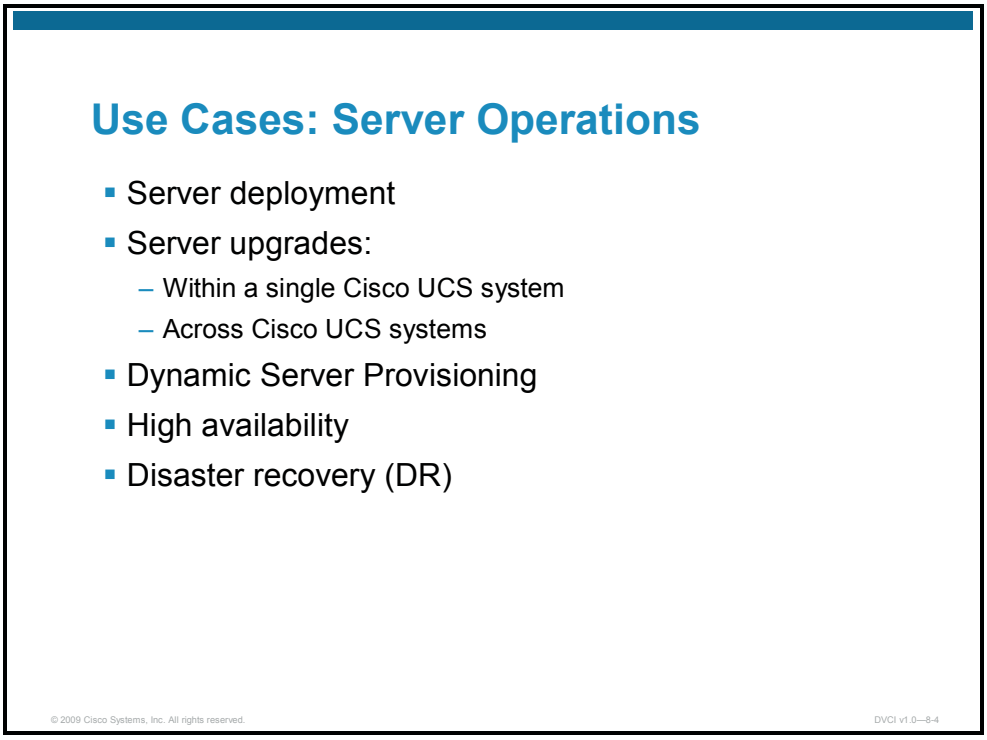

Server operations can be categorized into two distinct areas:

- Server deployments, which include new servers, upgrading, and repurposing
- Business continuity, which includes high availability and disaster recovery

These functions are not unique to blade server markets, and many of these solutions are faced by any company with larger server deployments. On a smaller scale (tens to hundreds of servers), the challenges faced in deployment and business continuity can still take advantage of the features provided by Cisco UCS system. However, the true advantage of the Cisco UCS system can be seen more easily on larger server deployments (thousands of servers).

Also keep in mind that the life cycle for servers in any organization is between 18 months to 3 years. In contrast to networking equipment, this is largely based on technology shifts; for example, the advent of Fibre Channel over Ethernet (FCoE). These technology shifts can happen with frequencies between 3-5 years. The changes in server technologies often follow the lifecycle of solution changes more closely than changes to the technology.

With changes to servers, also be aware that servers are part of a solution. In a way, they are the central technology for which a solution is designed. As the central part of a solution, servers have links and ties to many other technologies, like Ethernet and SAN. Because of this significant relationship, changes to servers are not trivial and often take a lot of planning and require scheduled down time.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

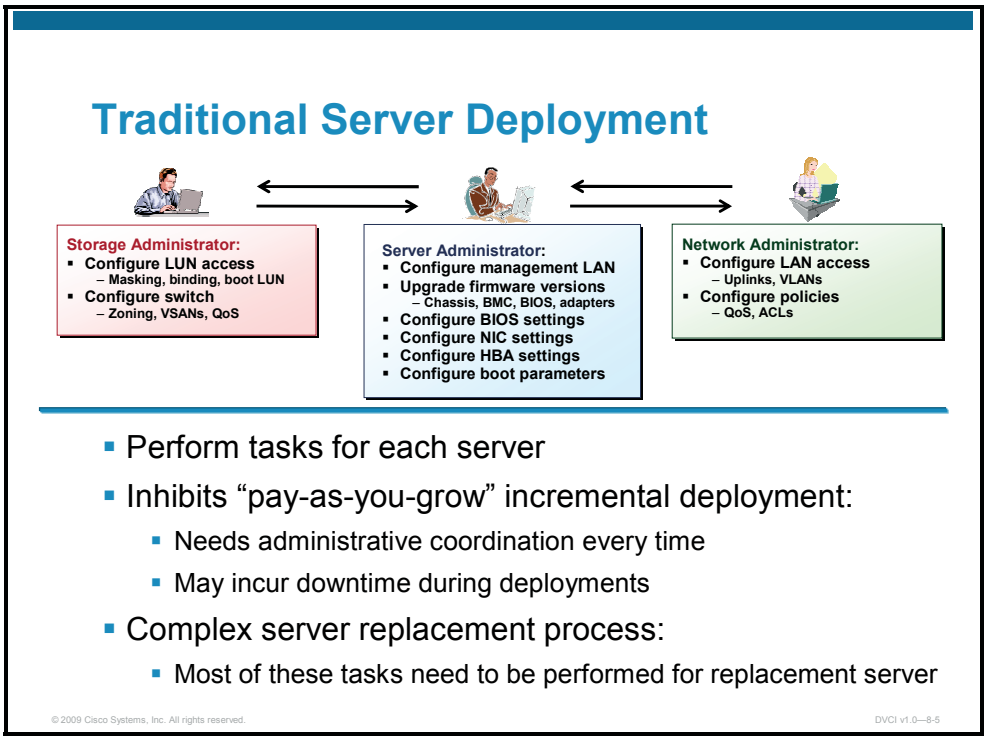

To better understand how the Cisco UCS platform can make server operations more efficient and easier, we need to first examine how servers are traditionally deployed.

The deployment of server solutions is not an easy task because it requires the coordination of tasks between different groups in an organization. In the figure, we see the segmentation of these groups by the tasks they typically provide in a server deployment.

As you can see, each job is focused solely on the technologies to which they relate. Each would require that someone coordinate the activities between these jobs to ensure successful server deployment. This situation impacts dynamic server growth and server replacement, forcing interaction between groups every time you deploy a new server, or replace an existing one. Now, imagine the complexity for organizations that need to deploy thousands of servers.

https://t.me/learningnets

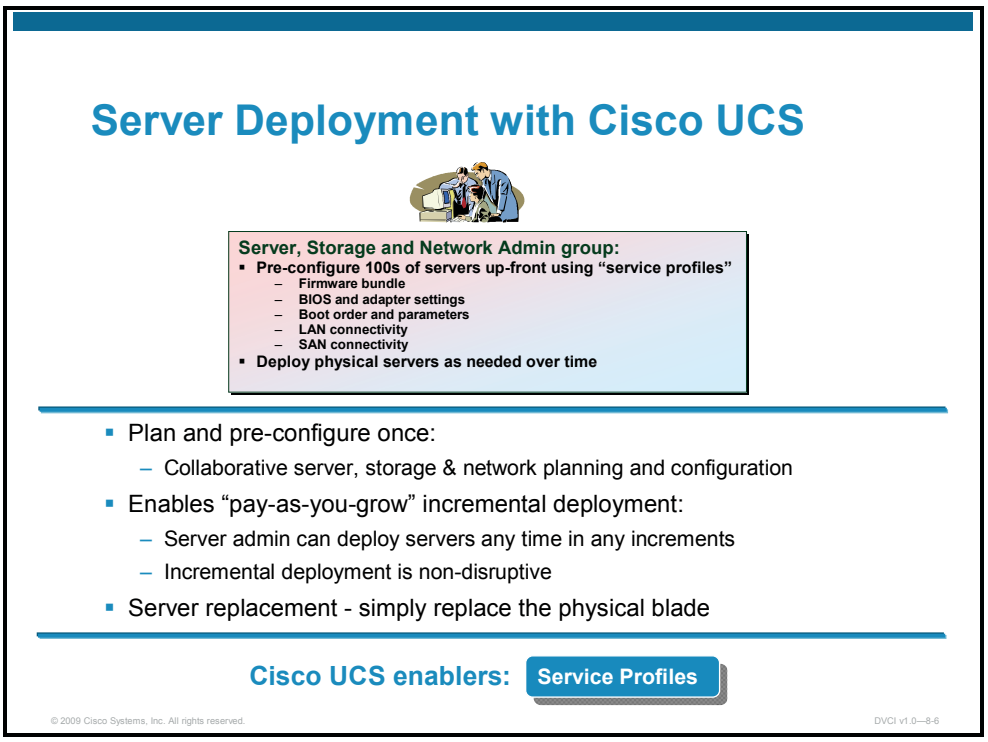

The Cisco UCS platform makes server deployments easy through the use of Service Profiles. The Service Profile is created by a server administrator, and will possibly use resources created by the storage and network administrator like MAC and WWN pools.

Service Profiles allow for the following capabilities:

- **Plan and pre-configure once:** By using resource pools, Service Profiles, and service templates, you can plan and configure for groups of servers at one time and then, using the profiles, provision new servers at any time
- **Incremental deployments:** After Service Profiles are created, adding additional servers is as simple as using a service template to create a new server
- **Server replacement:** This task is as simple as starting a profile on a new blade after it has been physically replaced

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

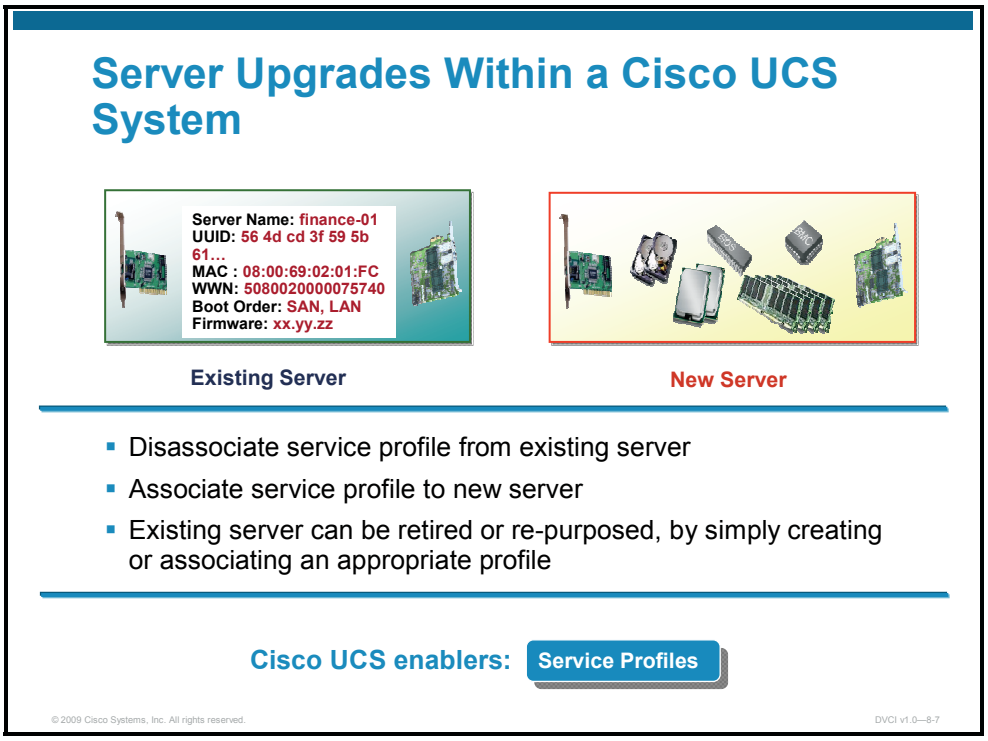

The Cisco UCS system enhances server upgrades through the use of the Service Profiles value proposition. Simply put, the blade in the Cisco UCS system is agnostic and can have its hardware defined through the use of the Service Profile.

A server in a traditional deployment will be built specifically for a solution that contains adapters, firmware, and boot definitions. If you wish to upgrade, you must again go through the process of provisioning the server for new hardware and coordinate that with other administrators. Then, you have to schedule extended time to make the move.

With the Cisco UCS platform you have a choice of two options:

- Change the Service Profile to boot the server on another upgraded blade.
- Copy the Service Profile and assign it to the other blade. Then, you can test the new server before pointing users to it.

https://t.me/learningnets

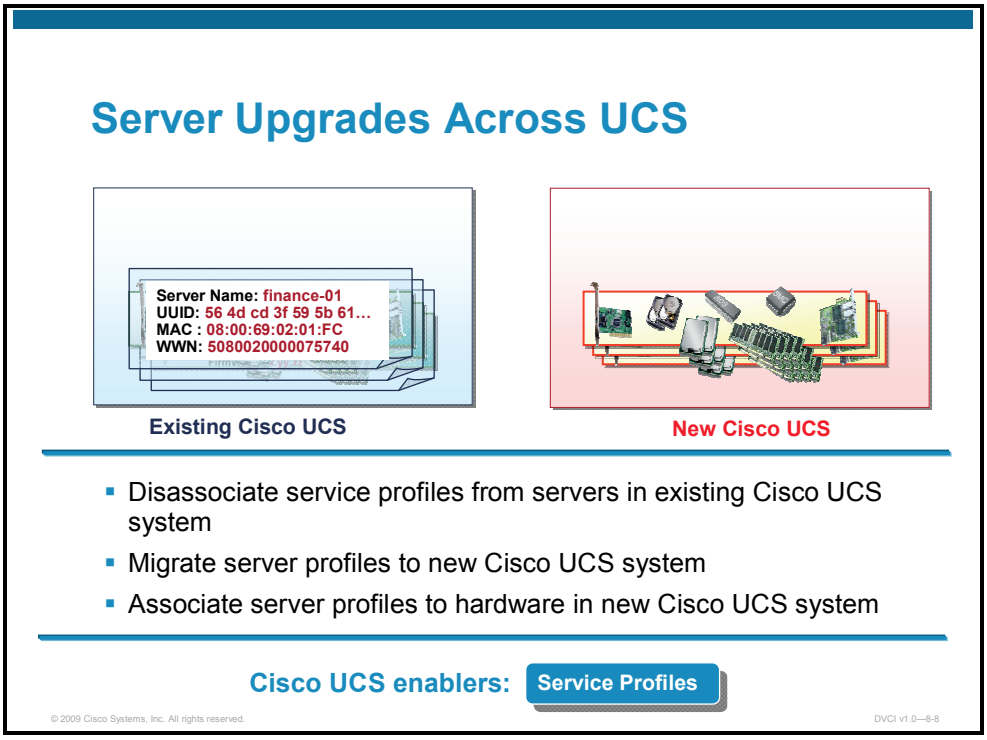

The process would be:

- Install and setup the Cisco UCS system so that it has the same software and firmware versions
- Ensure that SAN and networking resources are configured the same
- Copy over all service profiles
- Re-associate the profiles with the new server hardware

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

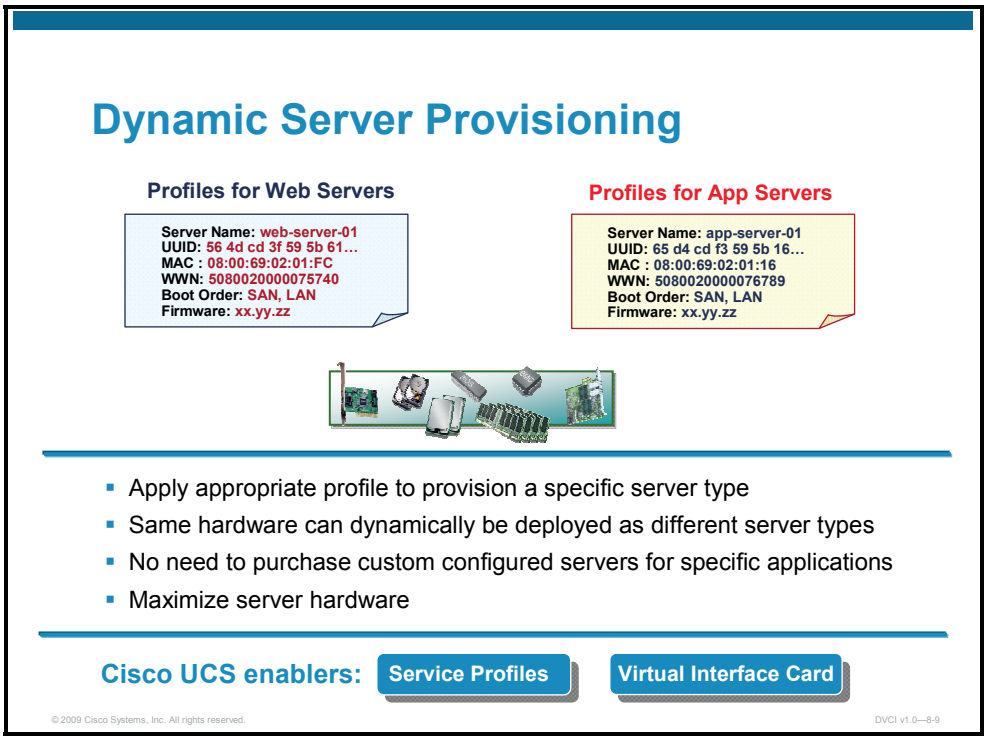

The challenges faced by companies with large numbers of servers include provisioning and server utilization. Typically, when designing a solution, a customer will need to design it for different levels of utilization. This could mean that in certain solutions they may need to purchase and have available additional server hardware and software for burst capacity. This hardware, while being under-utilized is also not able to be dynamically repurposed; if it is not needed, it must simply occupy space, power, and cooling resources.

Because servers are virtual and are configured using profiles, repurposing resources (server hardware) is as simple as disassociating a profile and re-associating a new one. Less-used servers could be shut down and repurposed for more highly used servers.

https://t.me/learningnets

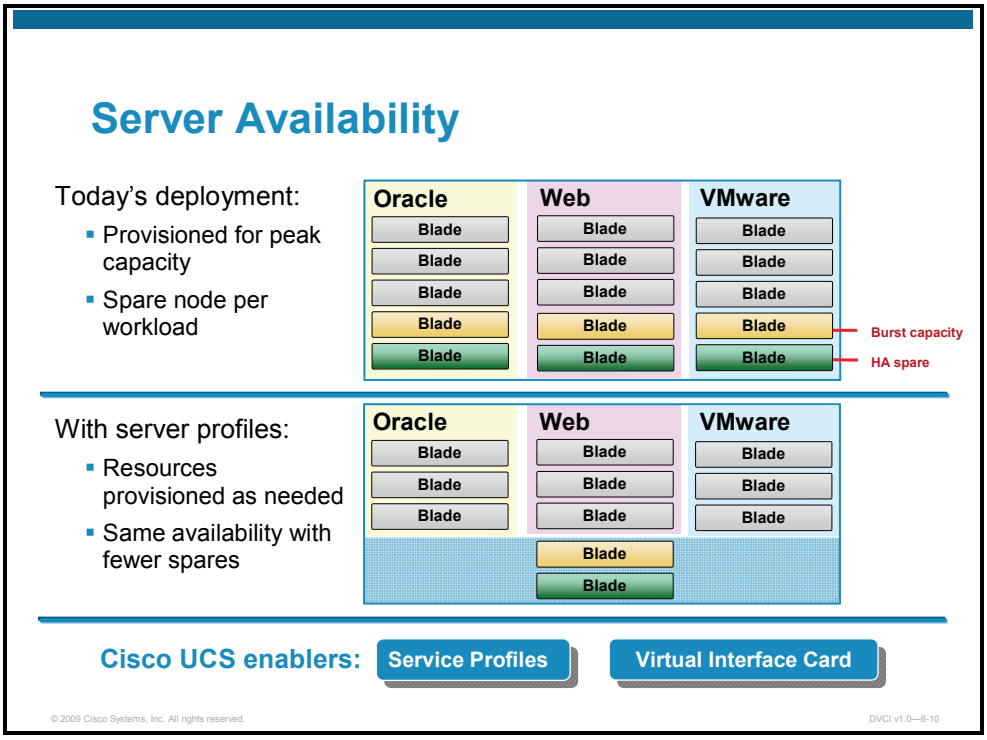

The same example of dynamic server repurposing can be applied to server availability. In a typical server deployment for a solution, it may require as part of the deployment that additional resources will have to account for burst capacity and high availability (HA). This, then, would possibly be applied across other solutions, as well. You can see an example of this in the above picture. In this example, you have six blades that may be only partially utilized, if used at all, to accommodate the burst and high availability needs.

By using Service Profiles, we have reduced this to two blades that can be shared across the three-server solutions. At the time of need, the appropriate profile can be applied to a blade, thus having it fulfill the appropriate role as needed.

You should be careful not to confuse this with high availability, which is provided by software vendors in the above solution. Instead, you should see this as a way to augment what high availability may be supplied through software.

This is also ideal for Cisco UCS systems in a multi-tenancy configuration.

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

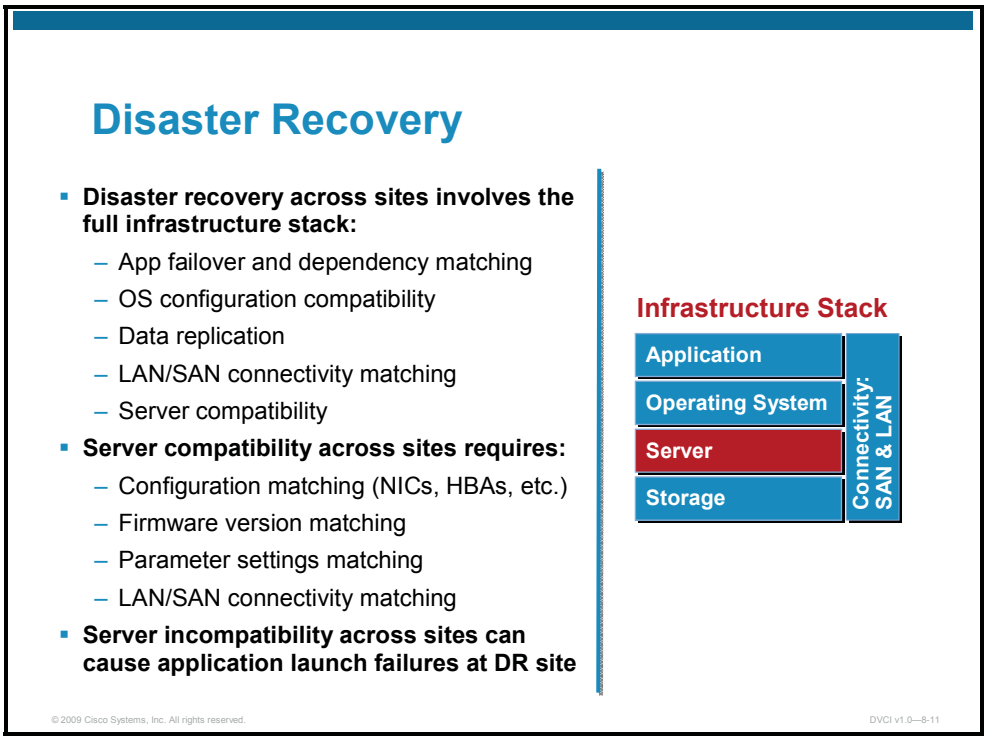

While we do not specifically provide a complete solution for disaster recovery, the features of Cisco UCS do aid in a disaster recovery solution. When designing a disaster recovery solution, you must consider the whole infrastructure stack. This includes:

- Application failover and dependency matching
- OS configuration and compatibility
- Data replication and integrity
- Connectivity matching (LAN and SAN)
- Server compatibility

Of these, the area in which the Cisco UCS system can aid significantly is server compatibility. Server compatibility requirements include:

- Configuration matching (adapters)
- Firmware matching
- Server parameters
- LAN or SAN specifics (vLANS and vSANS)

All of these items are configurable in the Service Profile for a given server, such that when recovering, the Cisco UCS can quickly ensure server compatibility on the recovery system. Without this ability, server compatibility can completely derail a disaster recovery solution.

https://t.me/learningnets

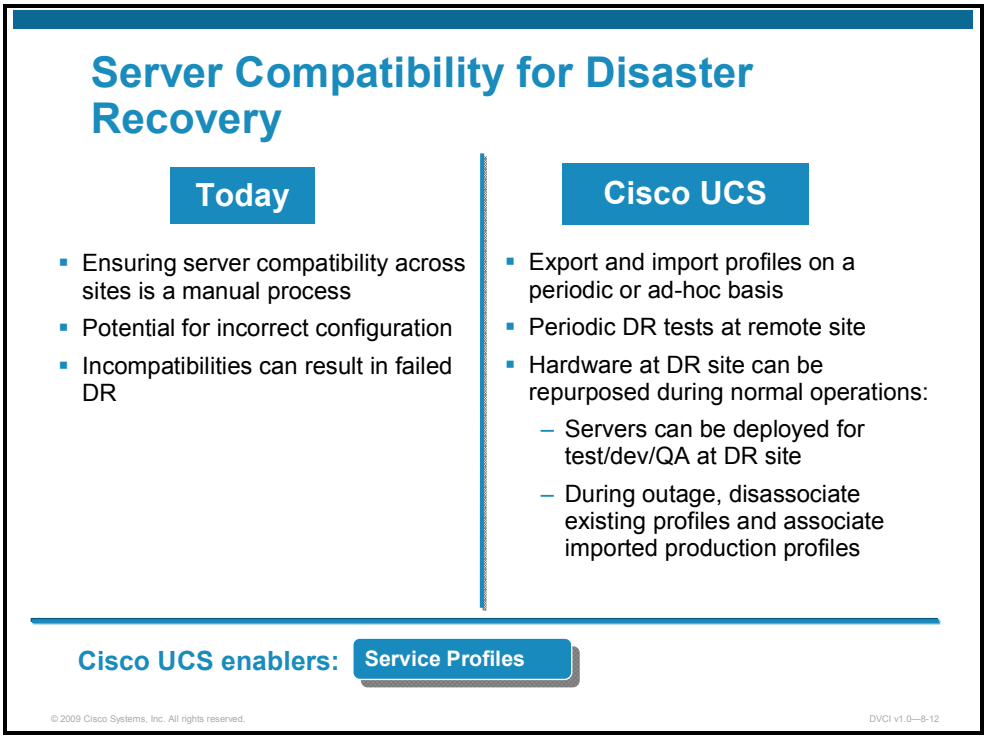

Without using the Cisco UCS system, much of server compatibility is a manual process that leaves a lot of potential for incorrect configuration. Other vendor products provide some of this compatibility, but they are limited in the scope of statelessness or in the ability to transparently use profiles over redundant systems.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Use Case: Virtualization

This topic describes changes in VMware virtualization and the four different use case scenarios with Cisco UCS and VMware.

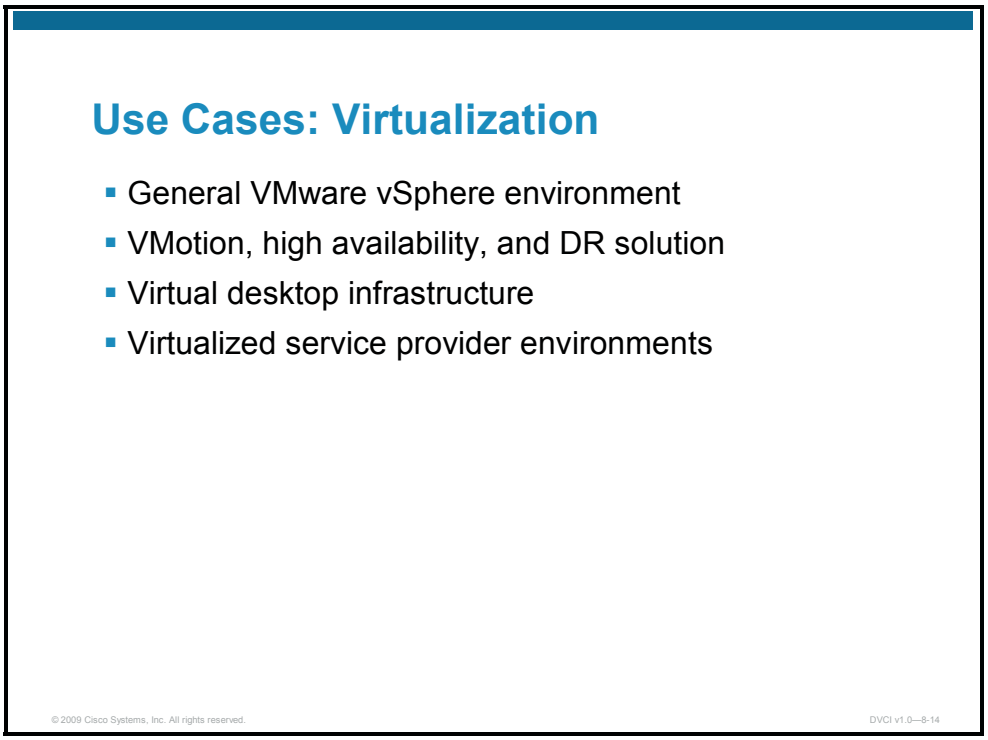

We will break down the use cases of virtualization into four separate cases:

- General vSphere environment—server running vSphere, hosting multiple VM hosts
- VMotion as a high availability and DR solution
- Providing a virtual desktop infrastructure
- Virtualized service provider environment

https://t.me/learningnets

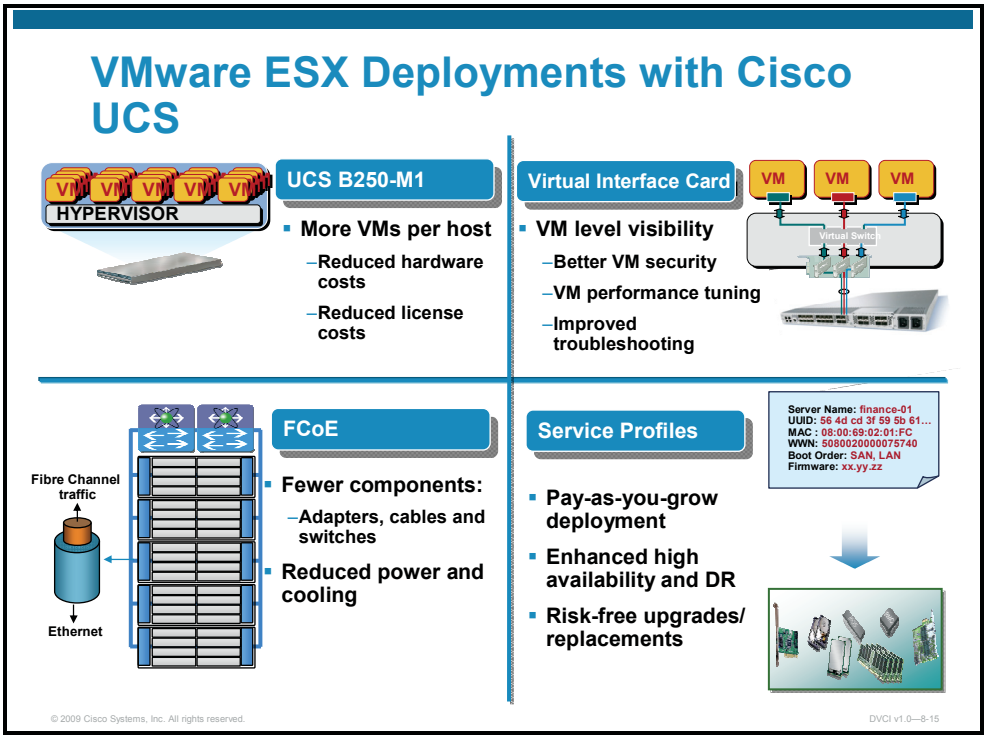

In a general VMware ESX deployment, we will have one or several blades running an ESX environment, and will host multiple virtual machines (VMs) on the blades. In this example, many of our value propositions apply, making the Cisco UCS ideal for ESX deployments.

Expanded Memory Blade Servers (Cisco UCS B250-M1)

The more memory that is available, the more VMs you can host per blade, allowing our blade to have a very cost-effective approach on a per-VM basis. However, there are other advantages, as well:

- Reduced hardware costs
- Reduced licensing (based on the number of cores)
- Larger memory handles vMotion better

Virtualization Integration

The Cisco UCS VIC M81KR adapter provides not only visibility into the VM networking, but also provides:

- Better security
- Performance—virtual adapters moved out to be handled at the hardware level, rather than the hypervisor
- Improved troubleshooting

Unified Fabric

The unified fabric provides a single network for connectivity, simplifying your configuration while reducing power and cooling requirements.

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

Server Statelessness

Service Profiles allow for the rapid replication and deployment of servers, which is good for a pay-as-you-grow environment. In addition, Server statelessness also provides for:

- Enhanced disaster recovery—eliminates worry in server compatibility DR
- Easy upgrades and replacement—server profiles, combined with vMotion

https://t.me/learningnets

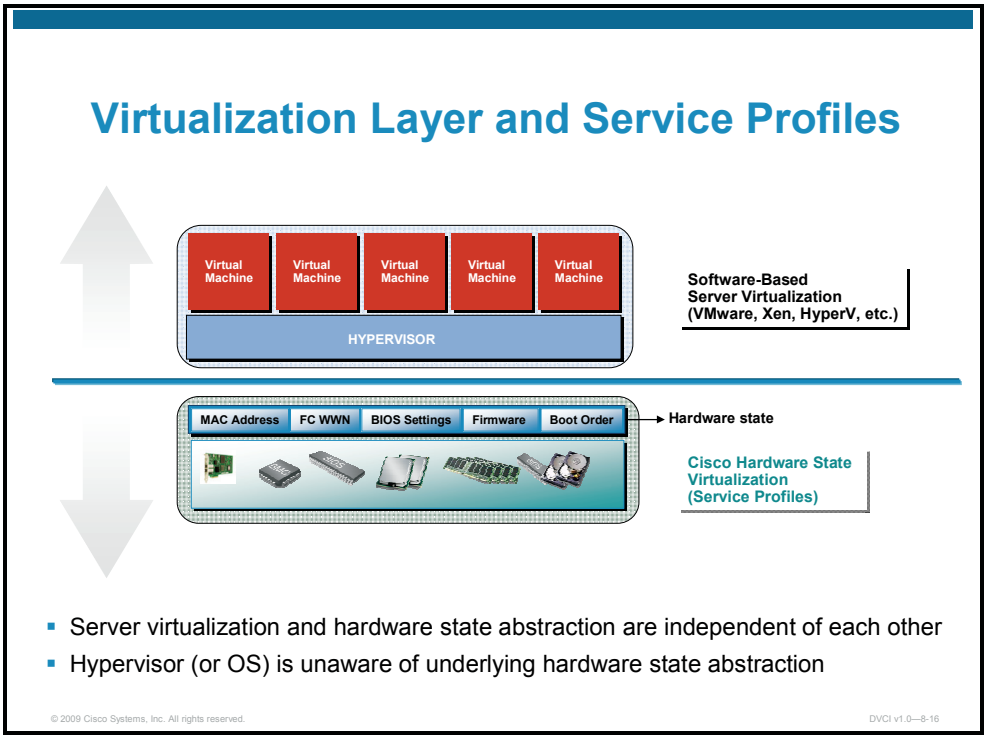

We want to be clear that what we do is not at all like VMware. VMware provides a software-based solution that is enabling server virtualization at a software level. VMware provides a hypervisor that allows for VMs to think they are accessing a common set of expected hardware.

Cisco UCS ensures server hardware compatibility across the entire Cisco UCS system. What this means is that the hypervisor will run on any blade in the system, and our software ensures that the hardware is identically configured, regardless of the blade from which the VMware ESX server is chosen to run.

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

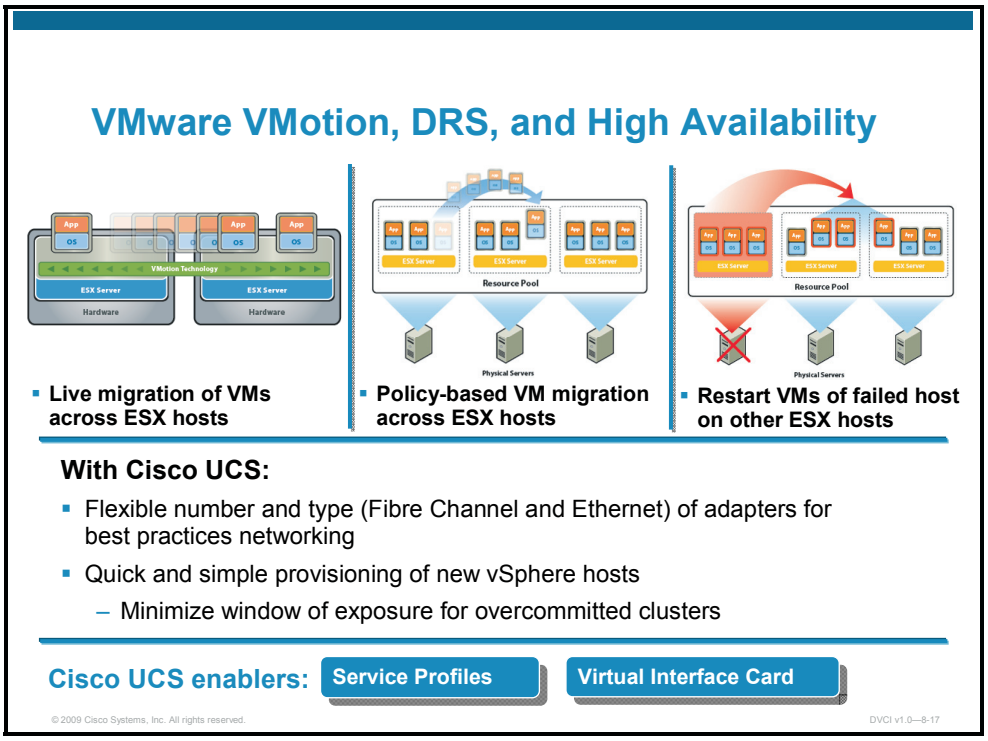

VMware provides a suite of software to enable high availability and DR for the VMs that are being hosted. This migration can occur by the following means:

- Live migration across ESX hosts: A manual intervention that can be used for upgrading servers. With California, we can easily use Service Profiles to create identical ESX servers, as needed, to move VM between hosts. Network and SAN access would be identical, saving time.
- Policy-based migrations: Moving a VM from one to another, based on some type of policy, like server utilization statistics. All access can be replicated simply through virtual NICs in Service Profiles.
- VM restarts due to failed hosts: Hosts can be quickly replicated with Service Profiles, thus relieving failover hosts. Larger memory footprints reduce impact of increasing the number of VMs during failure.

https://t.me/learningnets

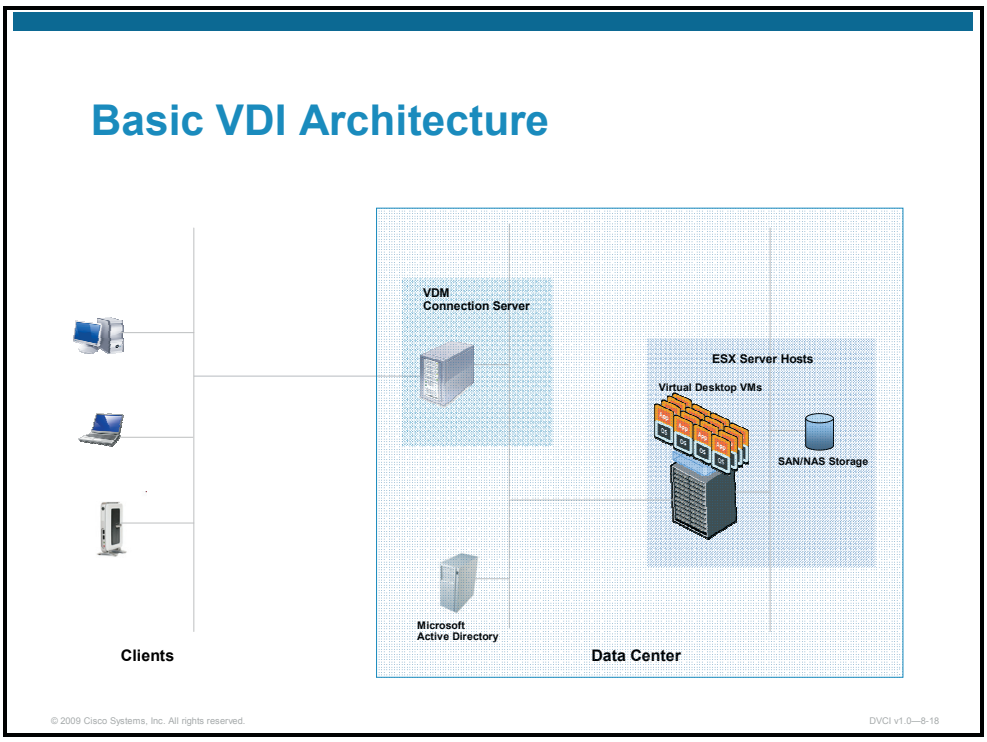

This is a basic architecture for a Virtual desktop infrastructure (VDI) environment. Clients access virtual desktops through a VPN Device Manager (VDM) server and are authenticated by—in this case—an Active Directory server. Virtual desktops are then hosted by an ESX server with network and SAN connectivity.

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

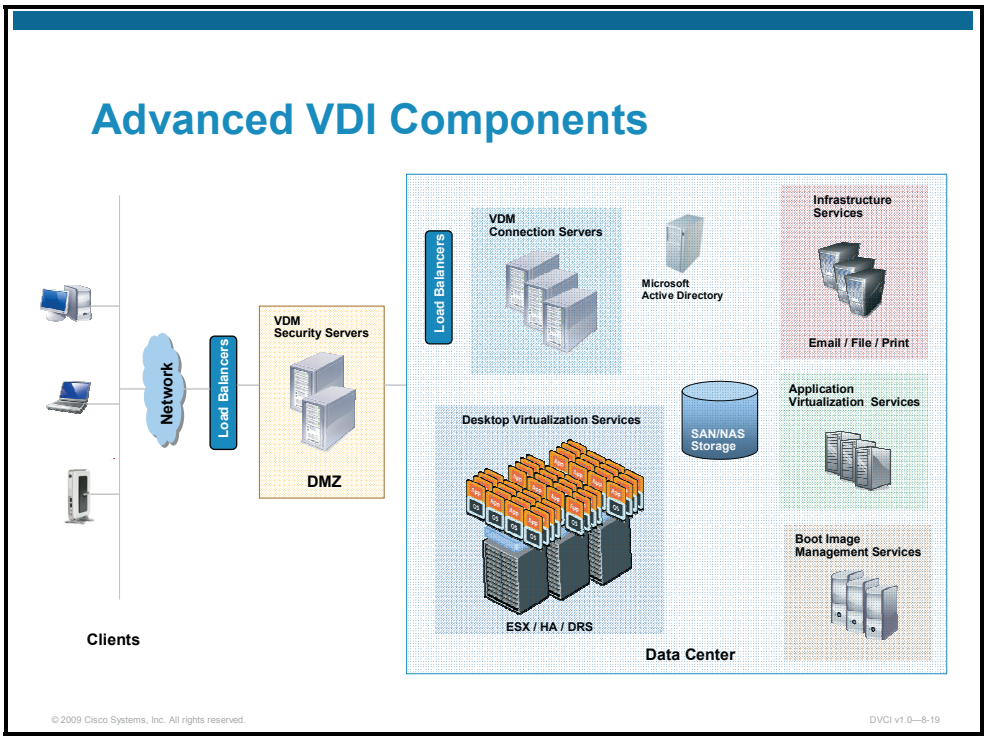

The figure shows a more advanced version of the previous example. In addition to providing the virtual desktops, we are also providing applications and other services. The results are identical; however, there are additional players that must be added to provide the applications and services.

https://t.me/learningnets

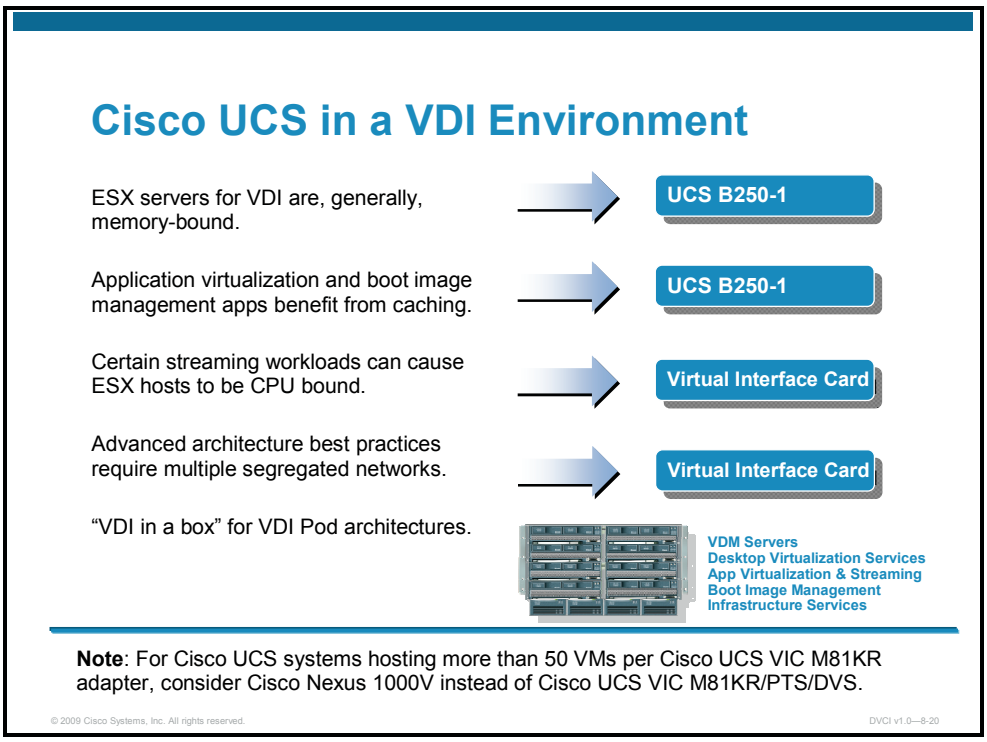

Based on the two previous examples, how does the Cisco UCS system enhance VDI deployments?

It does so in the following ways:

- VDI are very CPU and memory-intensive. The UCS B250-1 blade offering is ideal in that the memory is scalable to very large amounts, and the Nehalem CPU will be expanding the numbers of cores.
- Some customers may not want to put too many VMs "in one basket". Our ability to rapidly deploy servers using Service Profiles provides better high availability than traditional servers.
- Application Servers and boot servers can benefit from the greater memory available for caching.
- Streaming workloads can cause VMs to be server-bound. Cisco UCS virtualized adapters and adapter profiles provide ease of migration.
- Some architectures require specific network segregation. Cisco UCS virtualized adapter and vLANs can easily provide this segregation.

Lastly, the Cisco UCS system is ideal, because even in an advanced VDI deployment, it can provide for all of the server platforms that are needed, allowing a VDI deployment in a box or "pod".

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

Use Case: Databases

This topic describes the different types of database solutions architectures and explains how the Cisco UCS system is ideal for these types of solutions.

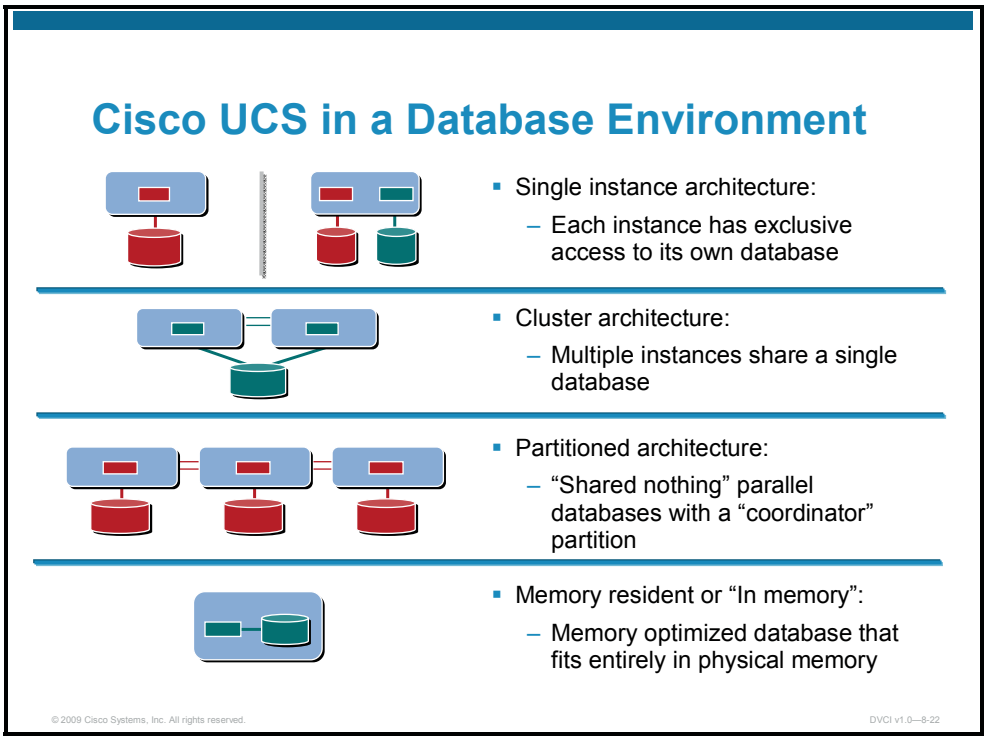

Before we look at how the Cisco UCS system can fit into a database solution, we need to quickly examine the different types of database architectures you may encounter. The four main types of databases that can benefit from using the Cisco UCS solution are:

- Single instance
- Clustered
- Partitioned
- Memory resident

https://t.me/learningnets

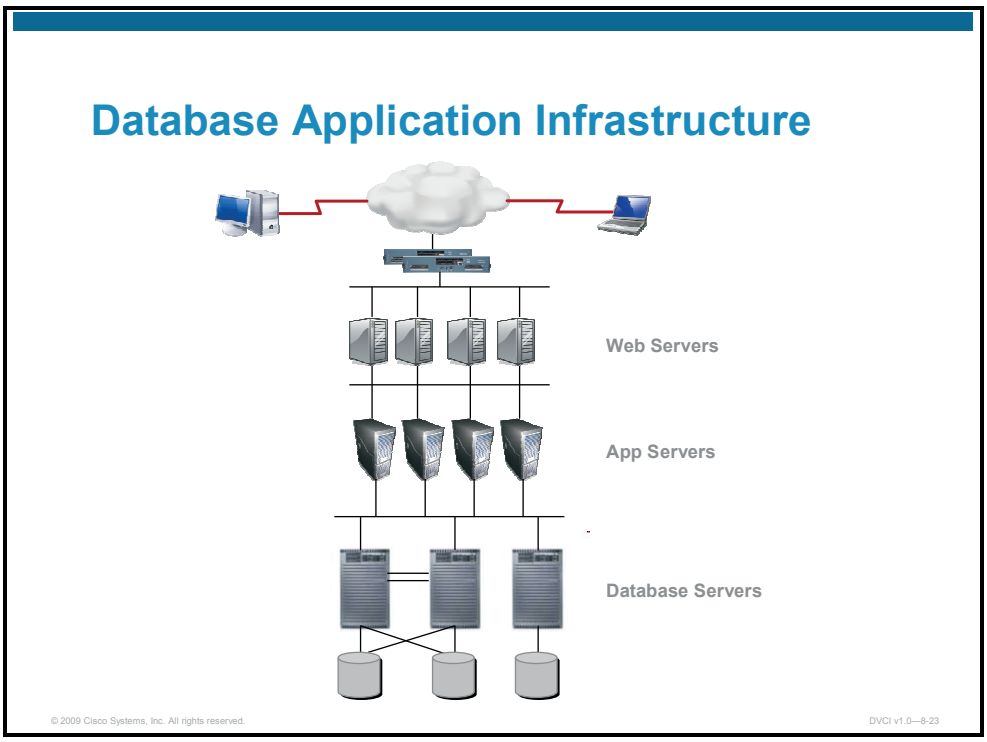

In this example, UCS could fit in at each level. Using the large memory footprint blades, UCS could be a good fit for the database servers. Our one-half width blades would be a good fit for the application server level. Lastly, using VMware, UCS might be able to provide a cost-effective approach to the web server access layer.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

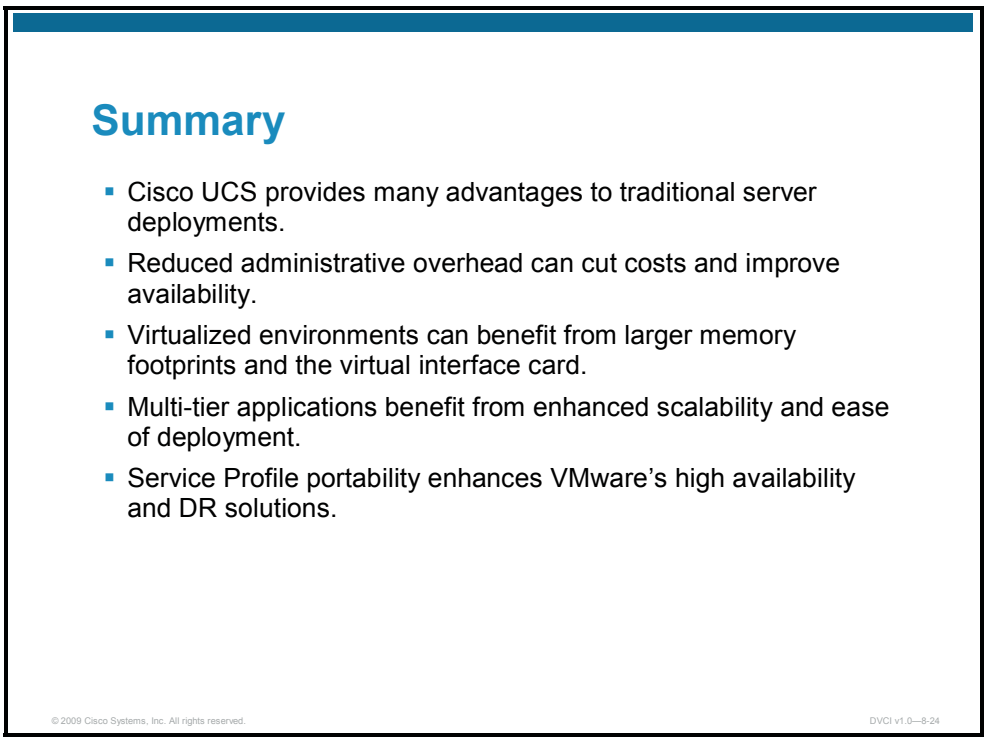

https://t.me/learningnets

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

Lesson 2

Operating System Provisioning

Overview

This topic describes operating system installation methods.

Objectives

Upon completing this lesson, you will be able to describe operating system installation methods. This includes being able to meet these objectives:

- Describe SAN-based operating system installation
- Describe local disk operating system installation
- Describe network-based operating system installation

https://t.me/learningnets

SAN Based installation

This topic describes SAN-based operating system installation.

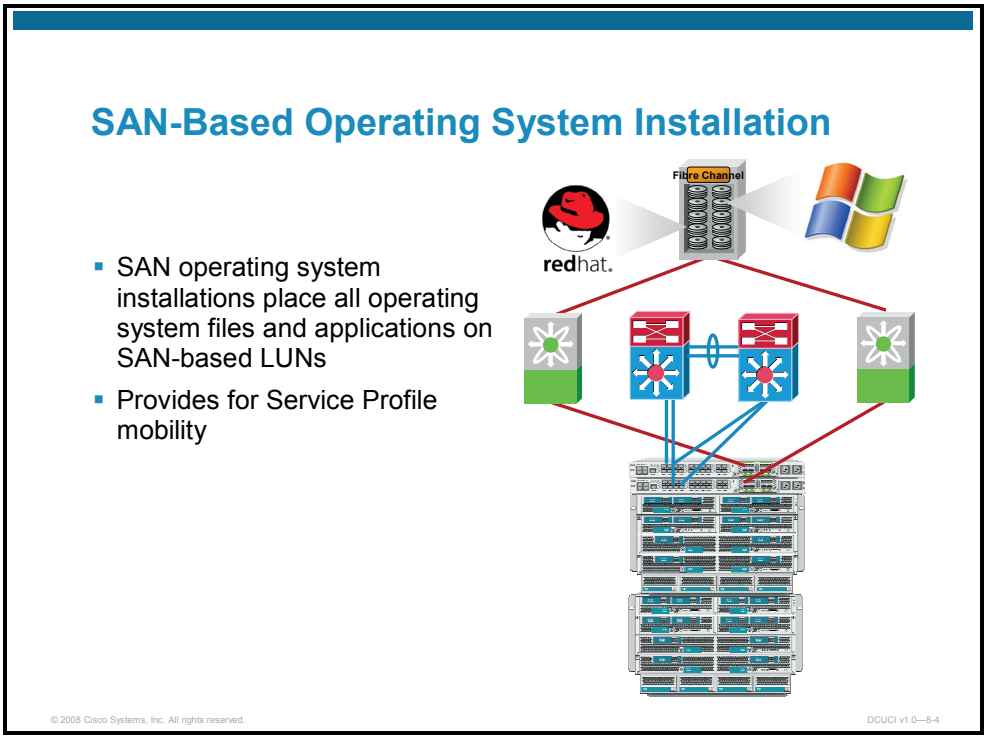

SAN boot can be used to centralize disk storage. This can reduce administrative overhead by minimizing points of configuration and replacement for local disks which tend to have a moderate to high failure rate.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Portability with Boot from SAN on UCS

- Due to virtualization of SAN identifiers, moving Service Profile between physical compute nodes doesn't change the boot method
- With proper configuration, operating system doesn't know move has occurred

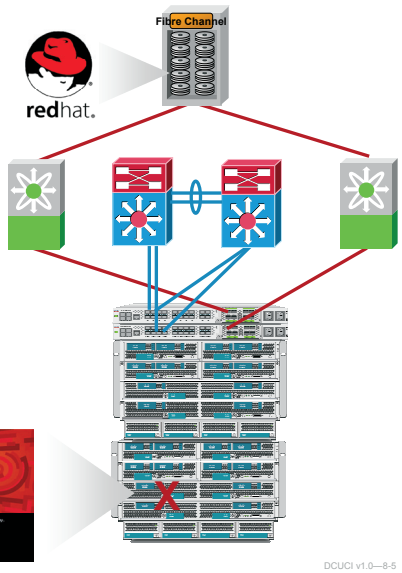

© 2008 Cisco Systems, Inc. All rights reserved.

In Cisco Unified Computing System (Cisco UCS) environments, servers can be stateless. This means none of the servers physical identifiers such as MAC address and world wide name (WWN) are tied to the physical server. Instead, these identifiers are stored in the Service Profile. This statelessness allows profiles to be migrated quickly to new compute nodes in boot from SAN environments.

https://t.me/learningnets

Considerations for SAN Boot

- Local disks could still be used for temporary, volatile data:
  - Not recommended for swap files, as lack of file after Service Profile movement may cause operating system problems
- Local disks not required in SAN boot scenario
- If using a SAN logical unit number (LUN) for boot device, ensure that profiles are moved between compute nodes with similar hardware:
  - Operating systems may fail to load properly if moved between nodes with very different hardware (Oplin to Menlo, for example)

© 2008 Cisco Systems, Inc. All rights reserved.

It is not recommended to use the local disks, if any, for swap files or partitions. Many operating systems will experience errors after moving the Service Profile from one node to another, as the swap file and signature may be missing or different.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Local Disk Installations

This topic describes operating system installation on a compute node's local disk.

Local Disk Installations

- Most closely matches traditional server deployment schemes
- Still may use SAN storage for application data
- Capacity is limited, portability is reduced
- External provisioning systems may be able to migrate operating system and applications
- May not properly migrate in all failure modes

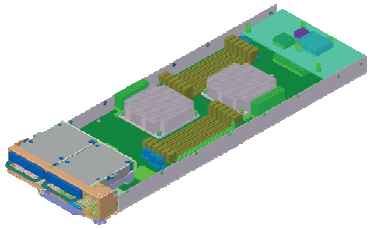

© 2008 Cisco Systems, Inc. All rights reserved.

When installing operating systems on local disk, Service Profile portability is significantly reduced. When moving a Service Profile, the disks could physically be moved to the new node, or an external provisioning system could be used to move or re-provision the necessary data.

https://t.me/learningnets

Network-Based Installation

This topic describes network-based operating system installation.

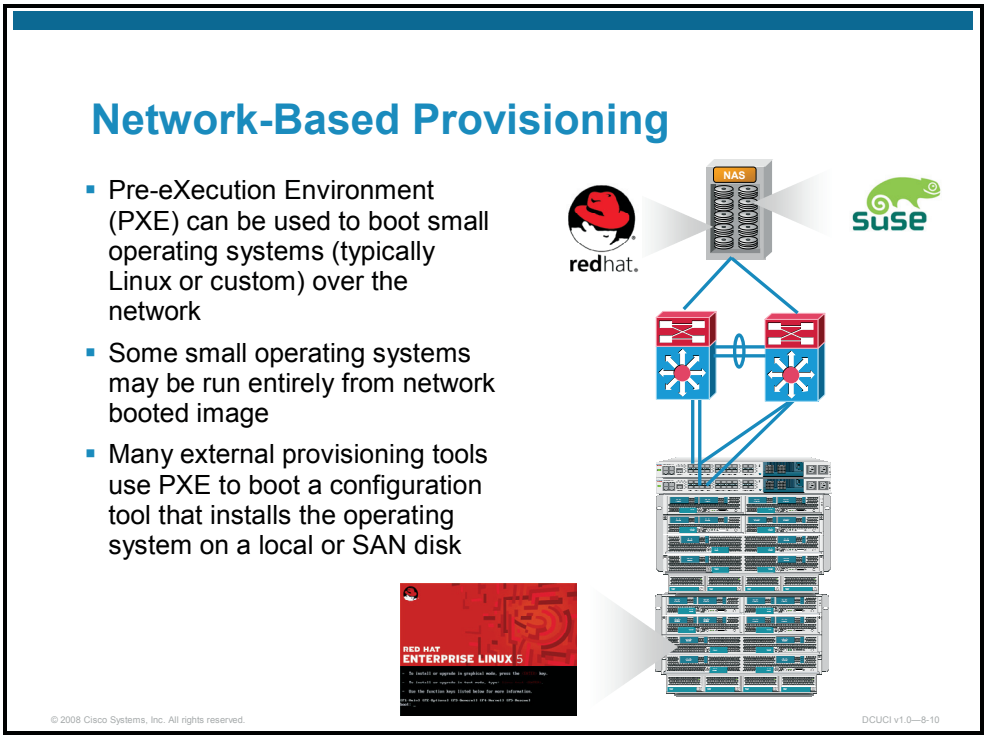

PXE boot can be used to boot devices from network accessible storage. In PXE boot environments, the operating system is downloaded from a network location and run in memory. This could be the target operating system (such as an appliance-type Linux image) or a configuration tool intended to install the final operating system image on local disk or SAN storage.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

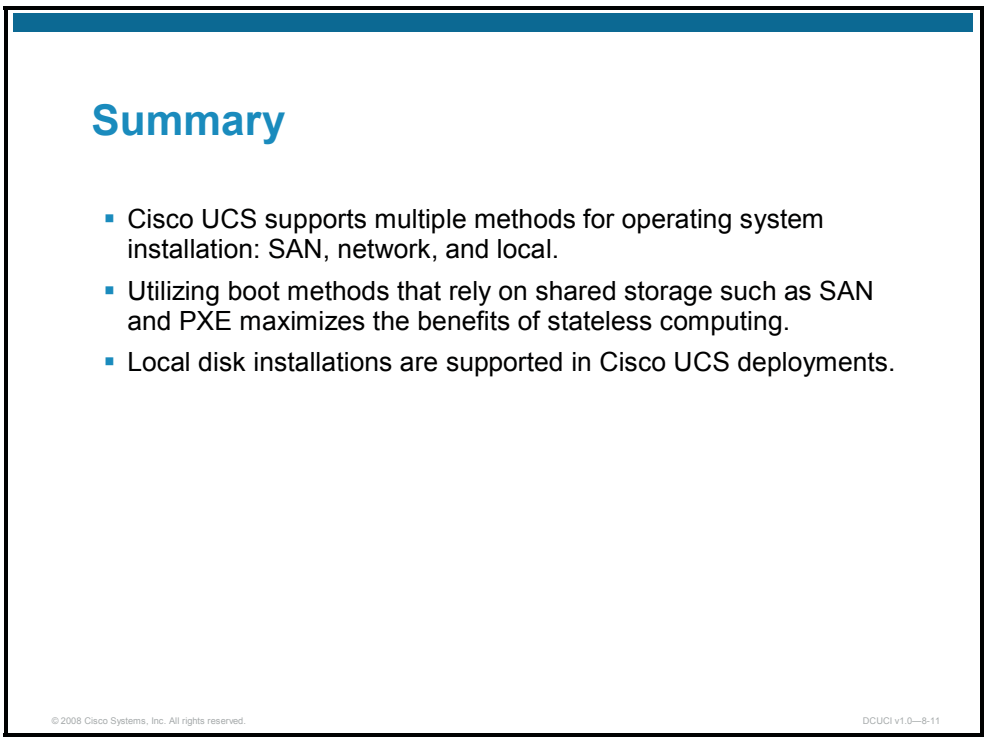

https://t.me/learningnets

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

Module Summary

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

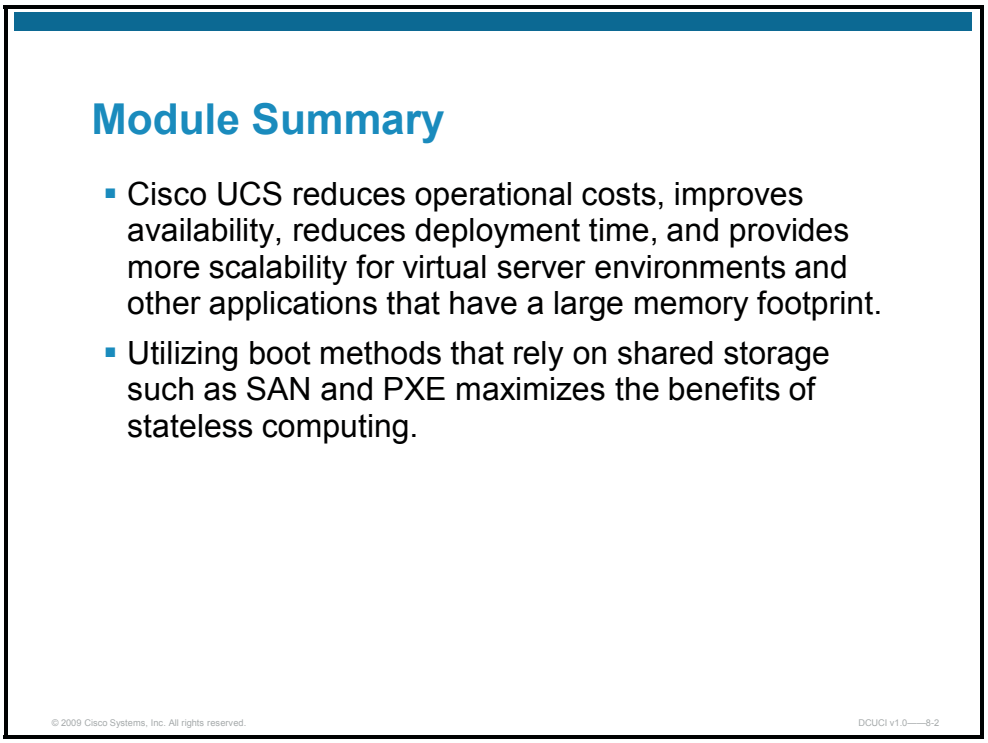

https://t.me/learningnets

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

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) In which case should you consider the Cisco Nexus 1000V instead of the Cisco UCS Virtual Interface Card (VIC)? (Source: UCS Use Cases)

A) More than 50 virtual machines on an ESX host
B) Heavy streaming network traffic
C) Use of VMotion
D) Database servers

Q2) What UCS advantage can be lost when using local disk operating system installations? (Source: OS Provisioning)

A) I/O Consolidation
B) Service Profiles
C) KVM-over-IP
D) Profile Mobility

https://t.me/learningnets

Module Self-Check Answer Key

Q1) A

Q2) D

© 2009 Firefly Communications, LLC

https://t.me/learningnets

Module 9

Networks in a Virtual Server Environment

Overview

This module explains how the Cisco Nexus 1000V fits in to a VMware server virtualization environment.

Module Objectives

Upon completing this module, you will be able to explain how the Cisco Nexus 1000V fits in to a VMware server virtualization environment. This includes being able to meet these objectives:

- Describe the Cisco Nexus 1000V switch and its role in a virtual server network environment
- Explain the VMware server virtualization environment
- Describe the VMware networking environment
- Describe the Cisco Nexus 1000V network architecture
- Explain how to configure VMware switching
- Install and configure a Cisco Nexus 1000V Switch
- Configure basic Cisco Nexus 1000V networking

https://t.me/learningnets

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

Lesson 1

Introduction to the Cisco Nexus 1000V

Overview

This lesson describes the Cisco Nexus 1000V switch and its role in a virtual server network environment.

Objectives

Upon completing this lesson, you will be able to describe the Cisco Nexus 1000V switch and its role in a virtual server network environment. This includes being able to meet these objectives:

- Describe data center access layer networks
- Explain the current virtualization environment
- Explain Cisco Virtual Switching
- Introduce the features of the Cisco Nexus 1000V

https://t.me/learningnets

Data Center Access Layer Networks

This topic describes data center access layer networks.

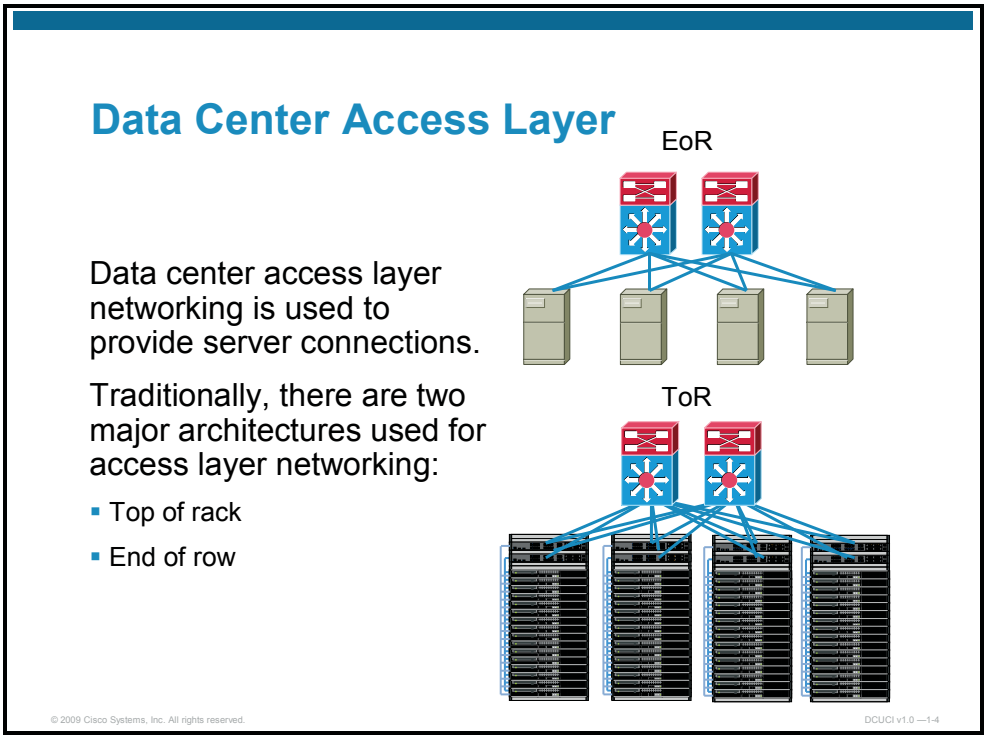

Typical data center access layers are built in one of two ways. Top of rack (ToR) architectures place switches at the top of each rack for server access. End of row (EoR) architectures place switches at the end of server rack rows for server connections.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

ToR Access Layer

ToR architecture:

- Flexible and scalable pod design:
  - Design once deploy many server rack design
  - Simplified rack replication
- Reduced cabling:
  - Servers cabled to ToR
  - Fewer cross-rack cables:
    - Switch uplinks only (ISL)
  - Lower cabling costs

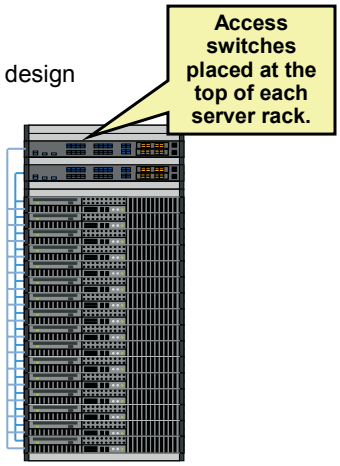

© 2009 Cisco Systems, Inc. All rights reserved.

In a ToR network design, redundant switches are placed at the top of each server rack allowing top down cabling with fewer uplinks to aggregation switches. This design minimizes cabling complexity but adds management complexity and power and cooling costs.

https://t.me/learningnets

EoR Access Layer

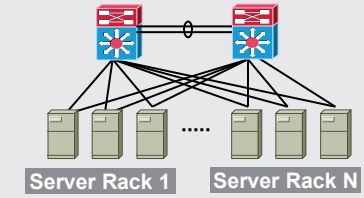

EoR architecture:

- Fewer network management points
- Fewer switching devices:
  - Less power
  - Less cooling
- Lower CapEx and OpEx
- Simplified service and software upgrade rollout
- Maintains feature consistency (security, QoS, multicast, and so on)

© 2009 Cisco Systems, Inc. All rights reserved.

In an EoR design, servers are cabled to access switches placed at the end of each server row. This design increases the complexity of cabling while reducing the management complexity.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Host Network Configuration

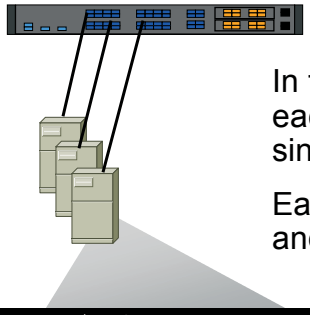

In traditional access layer networks, each end node or server ties to a single physical switch port.

Each port maintains its own statistics and can be configured manually.

```
Ethernet1/25 is up
    Address is 000c.2901.1a7b
    Rx
    1889 Input Packets 1624 Unicast Packets
    1 Multicast Packets 264 Broadcast Packets
    302538 Bytes
    Tx
    2025 Output Packets 1829 Unicast Packets
    0 Multicast Packets 196 Broadcast Packets 204 Flood Packets
    313526 Bytes
    15 Input Packet Drops 0 Output Packet Drops
```

© 2009 Cisco Systems, Inc. All rights reserved.

In standard physical server deployments, each physical server is attached to a switch port. Each physical server's networking configuration can be configured individually on the attached switch port.

https://t.me/learningnets

The Physical Server Environment (Cont.)

When troubleshooting network issues, the network administration team has full visibility to the server level. The one-to-one server to switch port configuration allows for fluid end-to-end analysis.

A compromised server can be quickly identified and locked down through standard network tools.

```
switch(config)# int e1/7
switch(config-if)# shutdown
```

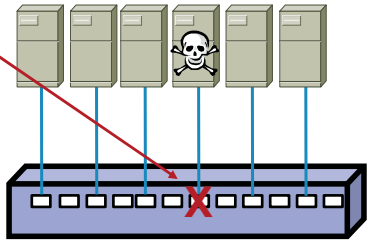

© 2009 Cisco Systems, Inc. All rights reserved.

Troubleshooting and analyzing network issues within a physical server network environment is accomplished through standard network tools and provides visibility down to the server level. Problems can be easily isolated and individual servers can be quickly isolated from the network until corrective actions can be taken. In the case of a security compromise, the offending server can be quickly isolated in order to protect the production network.

https://t.me/learningnets

Virtual Environment Overview

This topic explains the current virtualization environment.

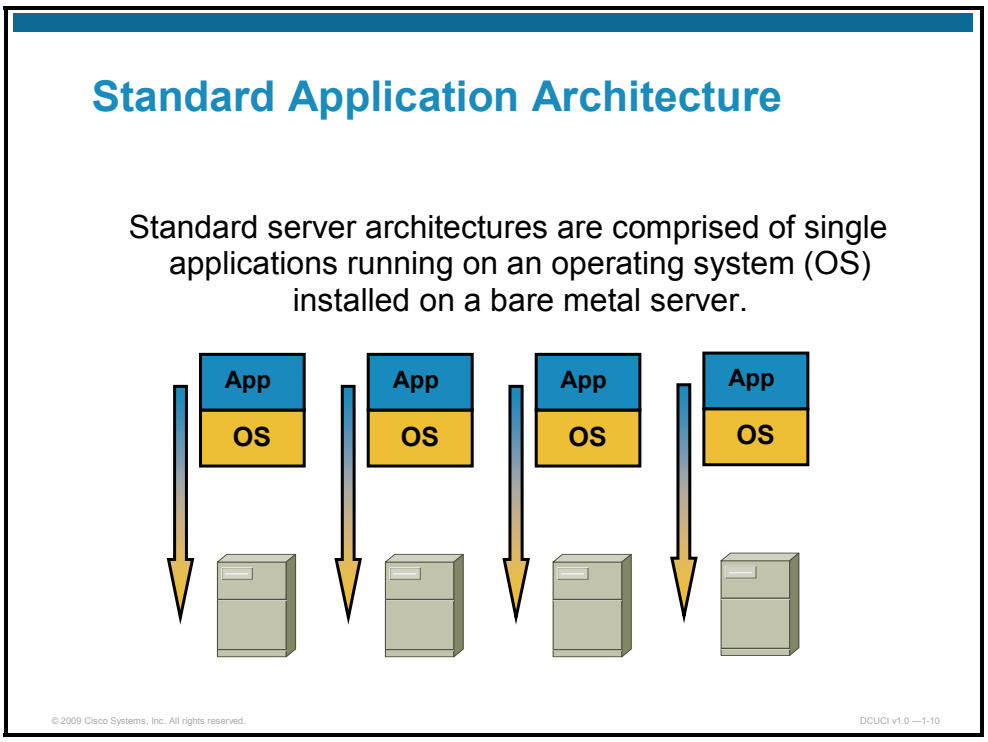

The standard application architecture on x86-based servers ties one operating system to a physical piece of hardware (server), and runs one application on that operating system. A common example of this would be an active directory server. The application is Active Directory, running on Windows Operating System on an individual server.

https://t.me/learningnets

Basic Virtualized Server Environment (Cont.)

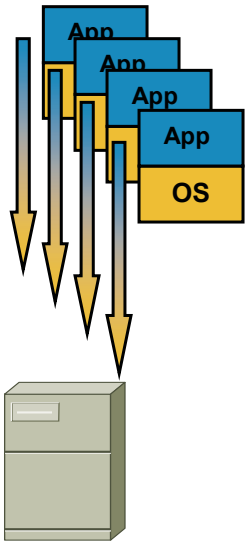

VMware environments allow multiple operating systems to be installed on a single physical server.

This method of server consolidation provides reduced costs and increased utilization.

© 2009 Cisco Systems, Inc. All rights reserved.

In the most basic virtualized environment, multiple operating systems run as guests on a single physical server. Each guest operating system runs an individual application.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

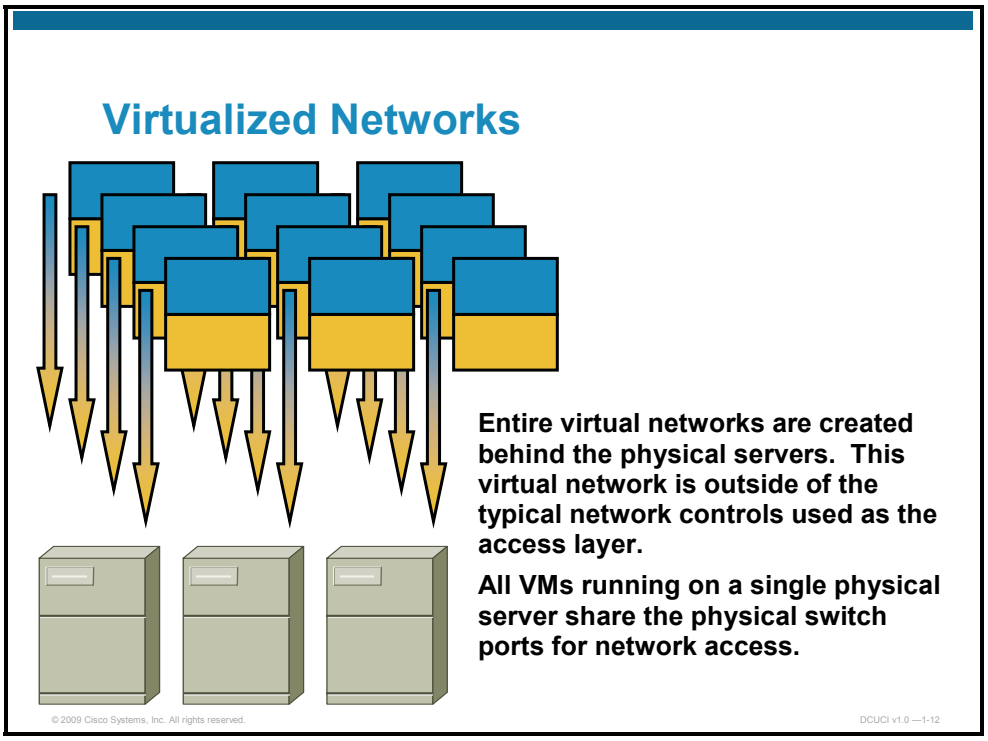

As multiple Virtual Machines (VMs) are created on each physical server, virtual networks are also constructed to support the I/O needs of the VMs. These networks sit outside the boundary of standard networking controls and best practices.

https://t.me/learningnets

Physical Switches in Virtualized Environments

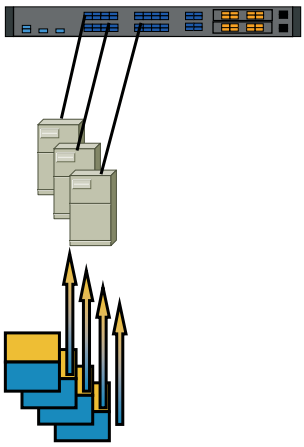

Multiple VMs running on a server connect through the same physical ports.

Physical port configuration effects all attached VMs instead of a single server.

Port configuration is static and not transportable with the VM.

© 2009 Cisco Systems, Inc. All rights reserved.

In virtualized environments, multiple virtual servers reside behind a single physical switch port. Any configuration done on the switch port applies to each virtual machine.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

The Virtual Server Environment (Cont.)

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 only be traced as far as the physical port and further diagnosis will need to occur.

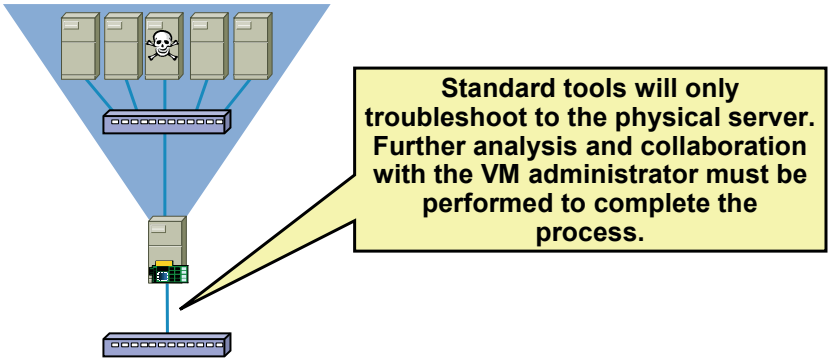

© 2009 Cisco Systems, Inc. All rights reserved.

Troubleshooting virtual networking issues becomes a more intensive task because multiple servers reside on a single physical piece of hardware. Standard network visibility is only to the physical level. Collaboration and/or additional software is required to troubleshoot network issues down to the problematic virtual machine.

https://t.me/learningnets

Network 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

Virtual server environments require new sets of network policies and configuration practices, and may involve a loss of security granularity at the server level.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows a physical network to virtual network feature comparison.

https://t.me/learningnets

Cisco Virtual Switching Overview

This topic explains Cisco Virtual Switching.

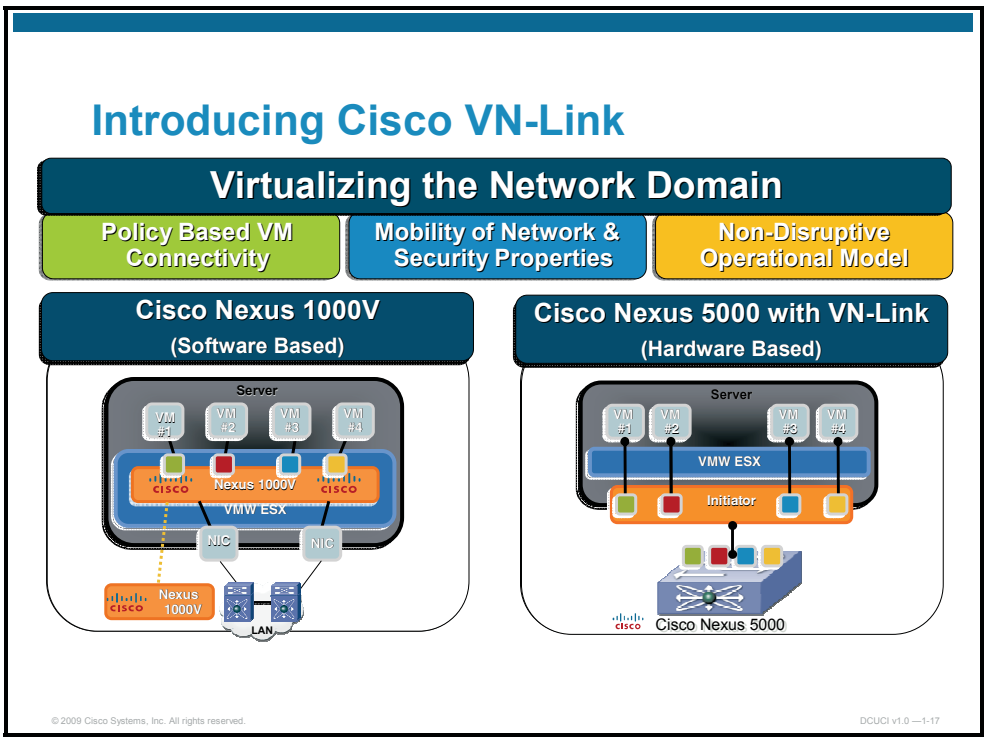

The Cisco Nexus 1000V is a software-based solution providing VM level network configurability and management. The Cisco Nexus 1000V works with any upstream switching system to provide standard networking controls to the virtual environment.

VN_Link technology was jointly developed by Cisco and VMware, and has been proposed to the IEEE for standardization. The technology is designed to move the network access layer into the virtual environment in order to provide enhanced network functionality at the VM level. VN_Link is deployed as a hardware-based solution offering VM visibility, policy based VM connectivity, policy mobility and a non-disruptive operational model.

https://t.me/learningnets

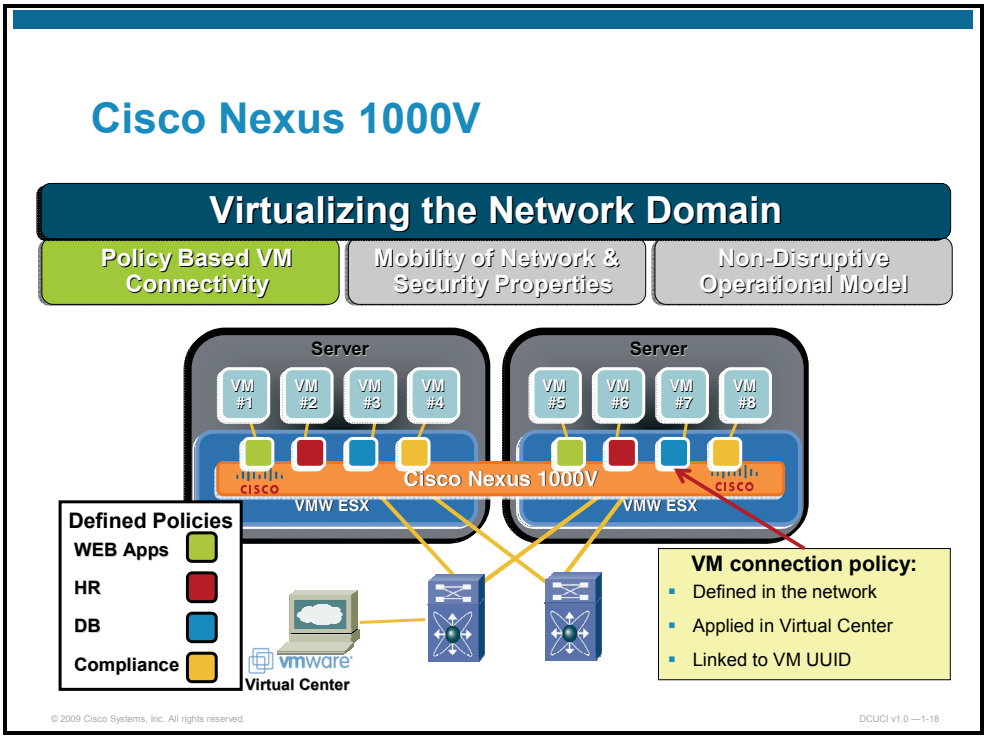

VM connection policies are defined in the network and applied to individual VMs from within Virtual Center. These policies are linked to the universally unique identifier (UUID) of the VM and are not based on physical or virtual ports.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

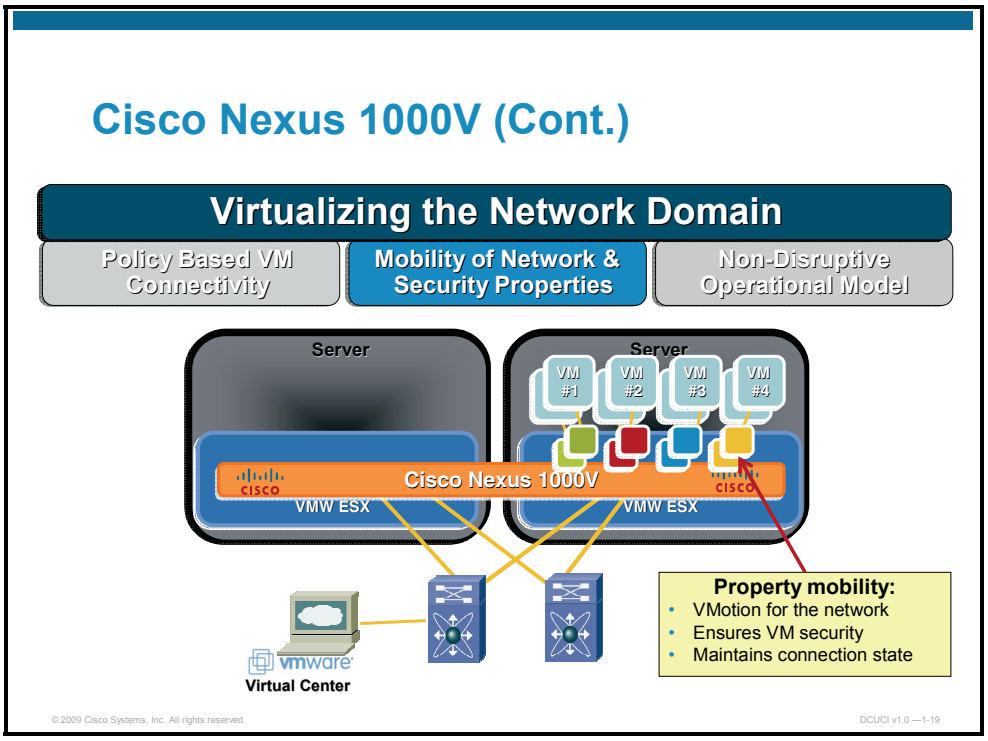

Through the VMware Virtual Center APIs, the Cisco Nexus 1000V monitors VM migration and ensures policy enforcement as machines transition between physical ports. Security policies are applied and enforced as VMs migrate through automatic or manual processes.

https://t.me/learningnets

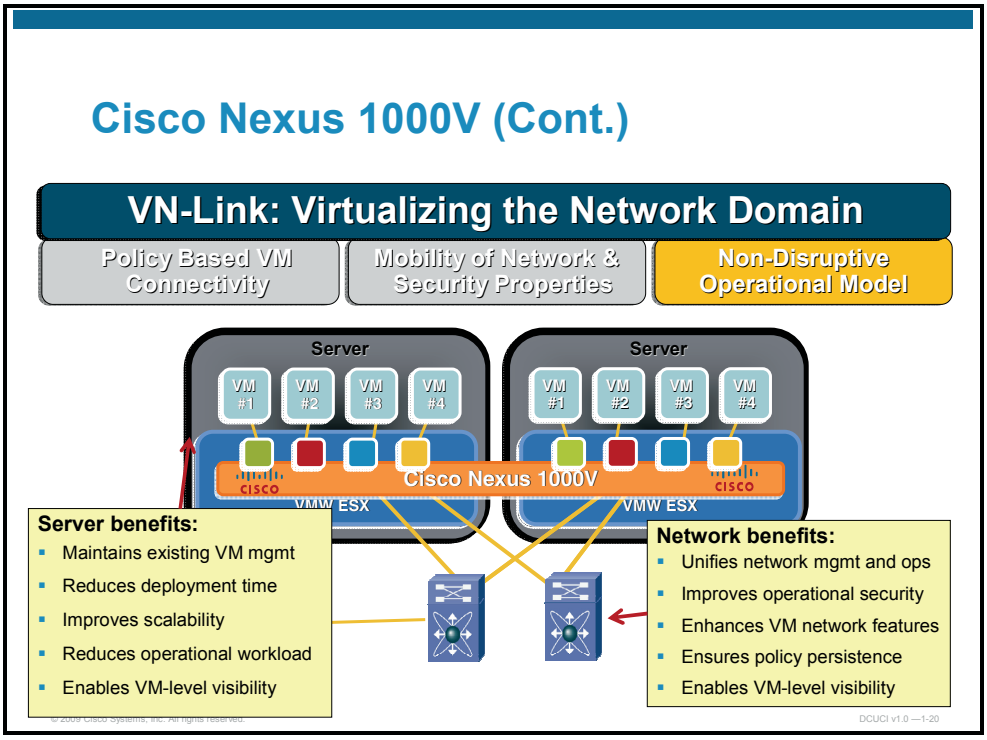

When using the Cisco Nexus 1000V, the management model for VMs stays the same and is handled by the VM administrator. Network administrators create security profiles and these policies are applied to individual VMs by the VM admin. Deployment time is reduced through pre-configured repeatable processes, reducing the operational workload. The benefits include unified network management and operations, enhanced network features at the VM level, and VM level visibility.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Feature Overview

This topic introduces the features of the Cisco Nexus 1000V.

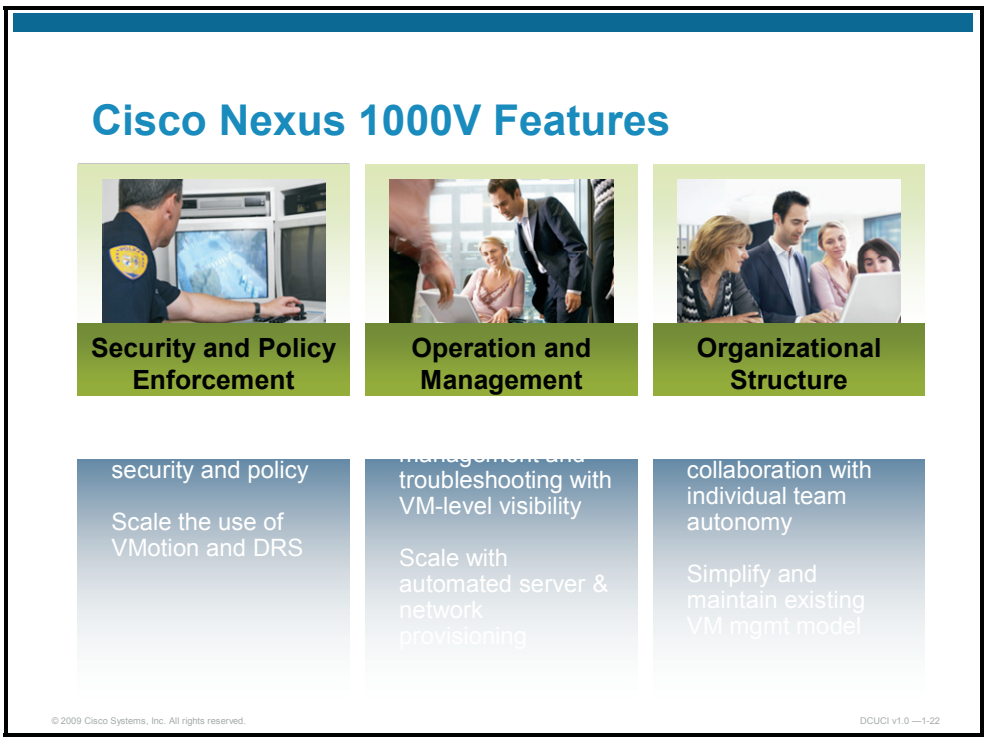

The Cisco Nexus 1000V offers many new features that were previously unavailable within the virtual access layer. These features fall into three major categories of Security and Policy Enforcement, Operation and Management, and Organizational Structure, as described in the figure. These three feature categories come together to provide the tools available to the physical access layer at the VM level.

https://t.me/learningnets

Cisco Nexus 1000V Features (Cont.)

Security and Policy Enforcement

TACACS	Granular access control lists (ACLs) provide tighter VM level security.
Radius	Centralized authentication (AAA, TACACS, RADIUS) provides enhanced management access control of virtual switching.
PVLAN	Private VLANs (PVLAN) provide Layer 2 segregation for multiple internal or external customers within the same subnet.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes the Security and Policy Enforcement features of the Cisco Nexus 1000V.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 1000V Features (Cont.)

Operation and Management

Plane	Single Management Point for Multiple ESX Hosts
DCNM	Network Wide Management Integration
VMotion Awareness	Automated Policy Migration
Port Profiles	Configuration Consistency
ISSU	In-Service Software Upgrades (ISSU)
NetFlow	Network Monitoring and Accounting

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes the operation and management features of the Cisco Nexus 1000V. The Nexus 1000V supports NetFlow version 9 with NetFlow Data Export (NDE).

https://t.me/learningnets

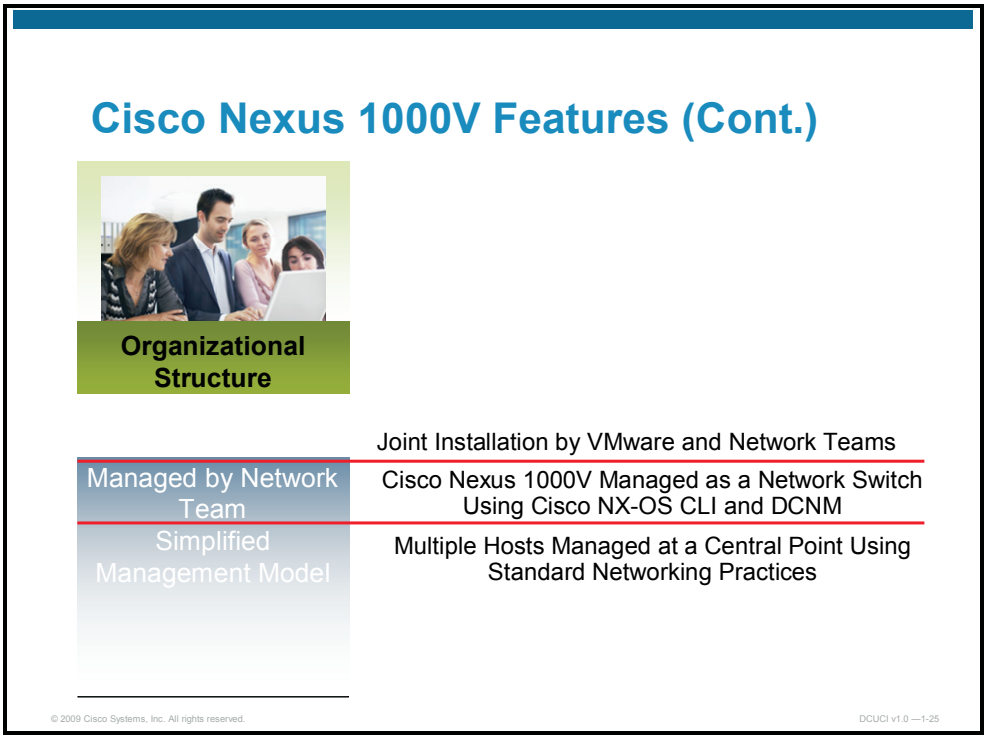

The figure describes the organizational structure benefits of the Cisco Nexus 1000V.

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
VM Creation	vCenter Based	N/A
Security	N/A	N/A
Visibility	vCenter	N/A
Management	vCenter	N/A

© 2009 Cisco Systems, Inc. All rights reserved.

The table describes the admin roles in standard VMware environments.

https://t.me/learningnets

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
VM 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

© 2009 Cisco Systems, Inc. All rights reserved.

The table describes the benefits provided to the Network and Server admin using the Cisco Nexus 1000V.

https://t.me/learningnets

Summary

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

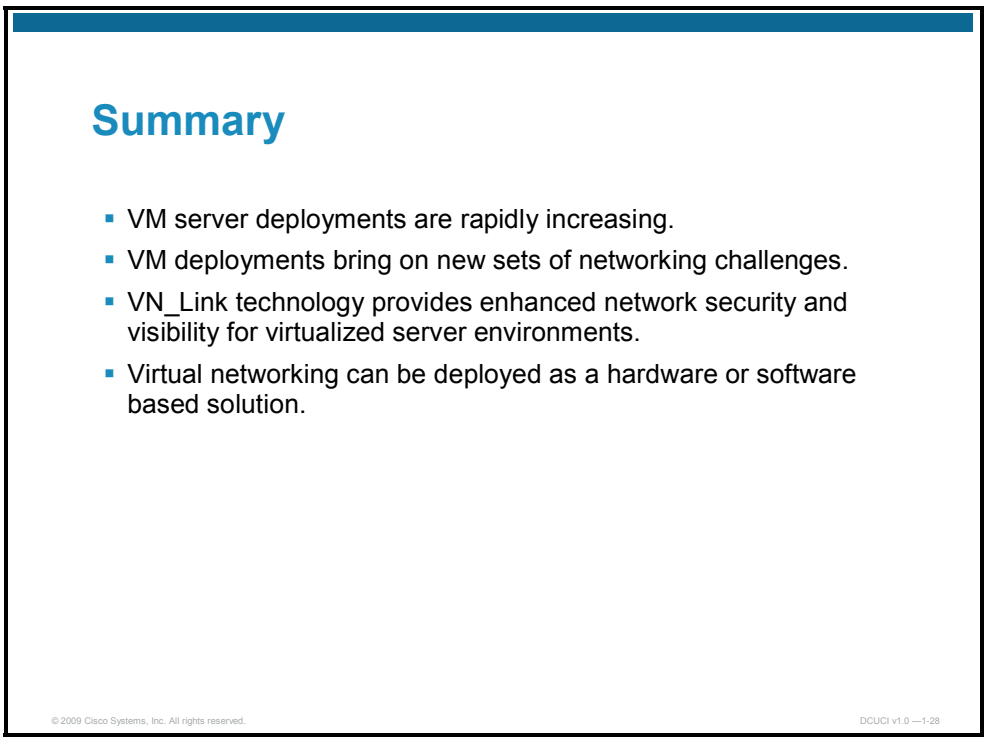

https://t.me/learningnets

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

Lesson 2

Server Virtualization Primer

Overview

This lesson explains VMware server virtualization.

Objectives

Upon completing this lesson, you will be able to explain the VMware server virtualization environment. This includes being able to meet these objectives:

- Describe the virtualized server model
- Describe the VMware virtualization environment
- Describe the virtual network architecture of VMware
- Describe the virtual storage network architecture of VMware
- Describe VMware environment management

https://t.me/learningnets

Server Virtualization

This topic introduces server virtualization.

Standard Application Architecture

- In a standard server architecture, each application and operating system is tied to a physical server.
- This is the common architecture used for x86-based servers.

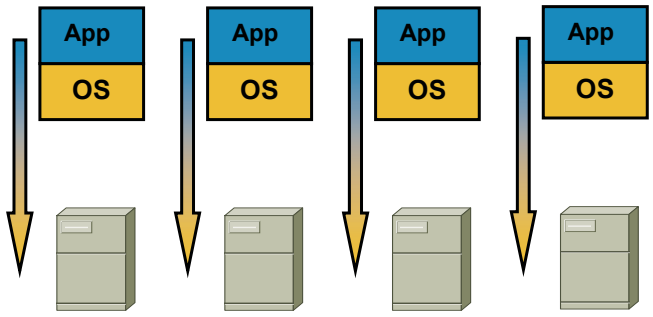

© 2009 Cisco Systems, Inc. All rights reserved.

The standard application architecture on x86-based servers ties one operating system to a physical piece of hardware (server) and runs one application on that operating system. A common example of this would be a Microsoft Active Directory server. The application is Active Directory running within Windows Operating System on a dedicated individual server.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Standard Application Architecture (Cont.)

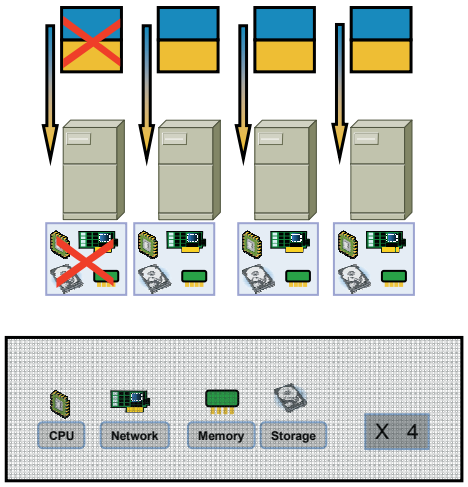

Benefits:

- Independent application environments
- Administration and security segregation
- Failure isolation:
  - One hardware failure causes only one application failure

Drawbacks:

- Low server utilization:
  - Typically 10–20% to allow for peak demand
- High administrative and environment costs

© 2009 Cisco Systems, Inc. All rights reserved.

The purpose of the standard application architecture is to provide fault tolerance between applications. If one piece of hardware or software crashes, only one application is effected at a time.

https://t.me/learningnets

Basic Virtualized Server Environment

- Multiple applications and operating systems run independently on the same physical server
- Virtual machines (VMs) participate within the data center much like a physical server:
  - Independent IP addressing and virtual MAC addresses
  - Separate operating systems
  - Individual virtual hardware

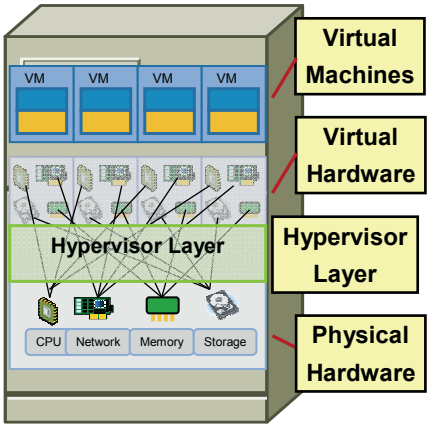

© 2009 Cisco Systems, Inc. All rights reserved.

In the most basic virtualized environment, multiple operating systems run as guests on a single physical server. Each guest operating system runs an individual application.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Basic Virtualized Server Environment (Cont.)

Benefits:

- Reduction of physical servers
- Reduced power and cooling
- Reduced cabling
- Centralized administration
- Ease of server deployment:
  - Physical servers can take hours to deploy while VMs can be deployed in minutes

Challenges:

- Lack of administration separation
- Shared architecture security concerns
- Lack of VM level LAN and SAN management

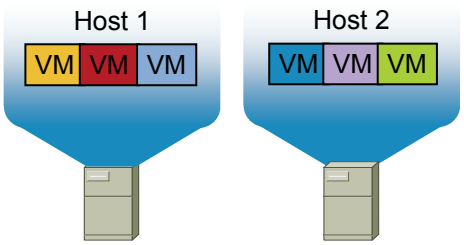

© 2009 Cisco Systems, Inc. All rights reserved.

Virtualization offers great benefits, including reduced power and cooling needs. The basic virtualization environment also brings with it new challenges due to the radical changes it requires of the datacenter.

Applications are isolated in individual VMs for protection. VMs can be clustered in the box or across boxes for increased application uptime.

https://t.me/learningnets

VMware Virtualization

This topic introduces the VMware virtualization environment.

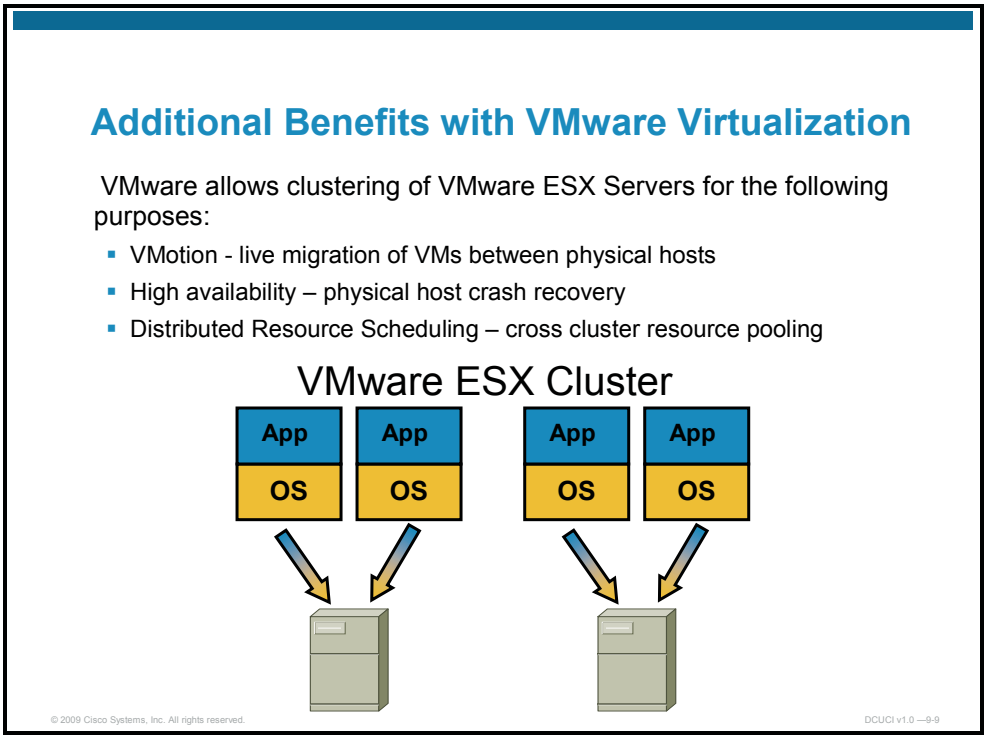

VMware offers cluster services to increase the value received from virtualization. These cluster services include high availability and Distributed Resource Scheduling (DRS). High availability and DRS can be utilized to increase availability and to maximize utilization.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

High Availability

- High availability monitors for physical server failures and boots crashed VMs on available hardware
- Servers will experience downtime approximately equivalent to a system cold boot

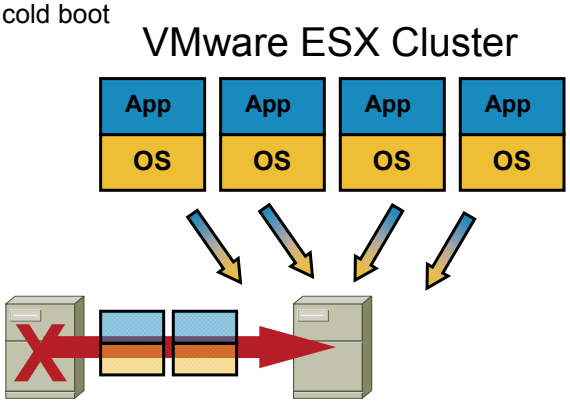

© 2009 Cisco Systems, Inc. All rights reserved.

High availability is used to monitor physical VMware ESX Hosts and to react in the event of a failure. If a physical host fails, the VMs running on that host are brought back up on available hosts in the cluster.

https://t.me/learningnets

Distributed Resource Scheduler

When user set thresholds are exceeded, the DRS agent will either make recommendations for migration or automatically migrate the VMs (as defined by user).

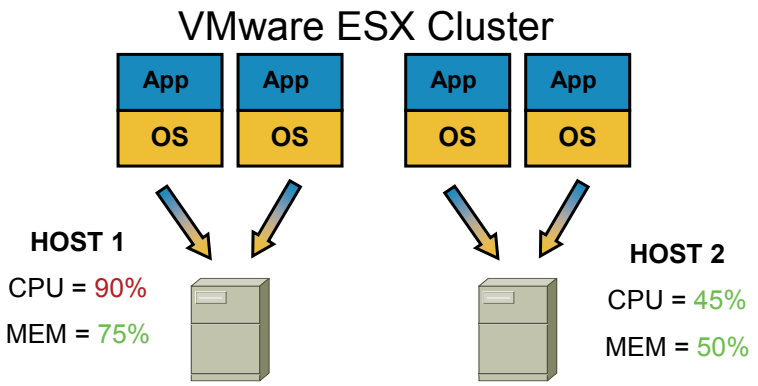

© 2009 Cisco Systems, Inc. All rights reserved.

DRS monitors the available resources of the physical servers within the cluster. Users set thresholds for migration. Migration can be handled automatically, or DRS can be set to make recommendations only.

https://t.me/learningnets

Distributed Resource Scheduler (Cont.)

Using VMotion, the server stays running throughout the process and the application remains unaware.

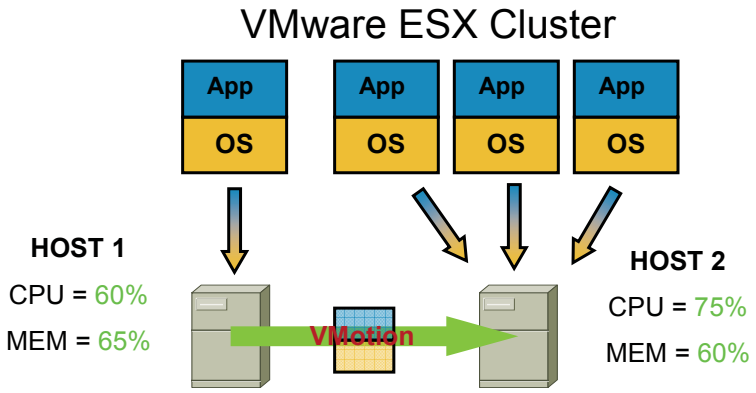

© 2009 Cisco Systems, Inc. All rights reserved.

DRS is used to evenly distribute the workload across available physical resources. DRS allows the maximization of resource utilization during fluctuating performance peeks.

https://t.me/learningnets

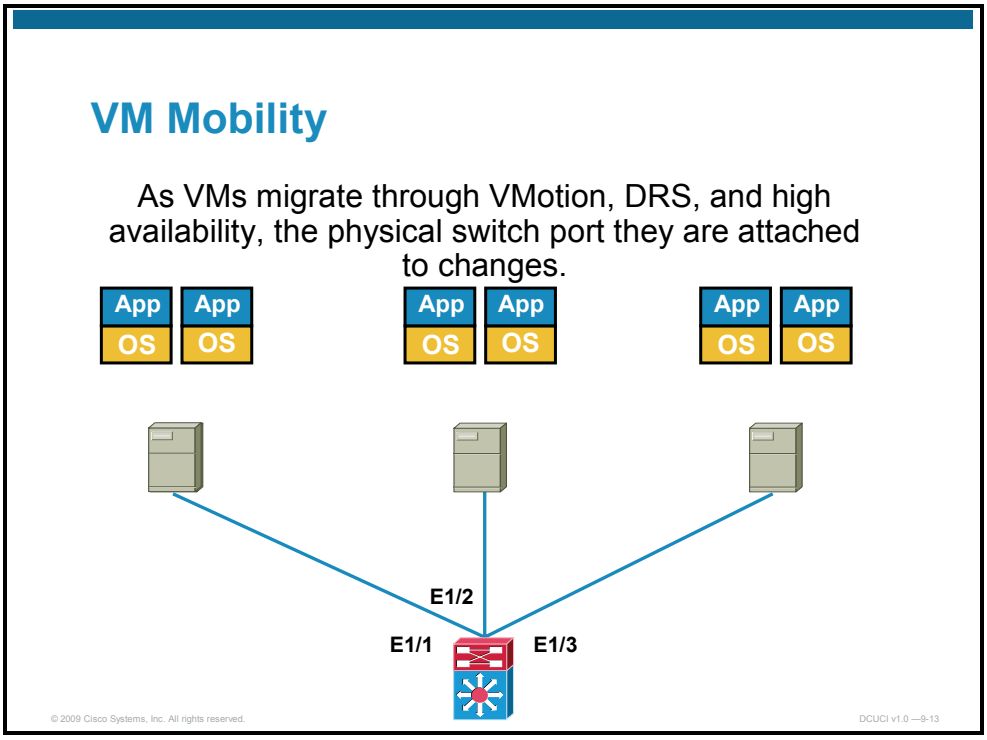

VMware offers cluster services to increase the value received from virtualization. These cluster services include high availability and DRS. high availability and DRS can be utilized to increase availability and to maximize utilization. Additionally, by using VMotion, administrators can manually migrate VMs between hosts.

The mobility of VMs places additional requirements on the physical network configuration. With a standalone host running multiple VMs, server ports connected to virtual hosts must be configured in a generic enough fashion to provide appropriate network connectivity for each individual VMs running on the host. When migration is brought into play, each network port that a given VM may migrate to will also need to have an identical configuration.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Virtual Networking with VMware

This topic introduces VMware network virtualization.

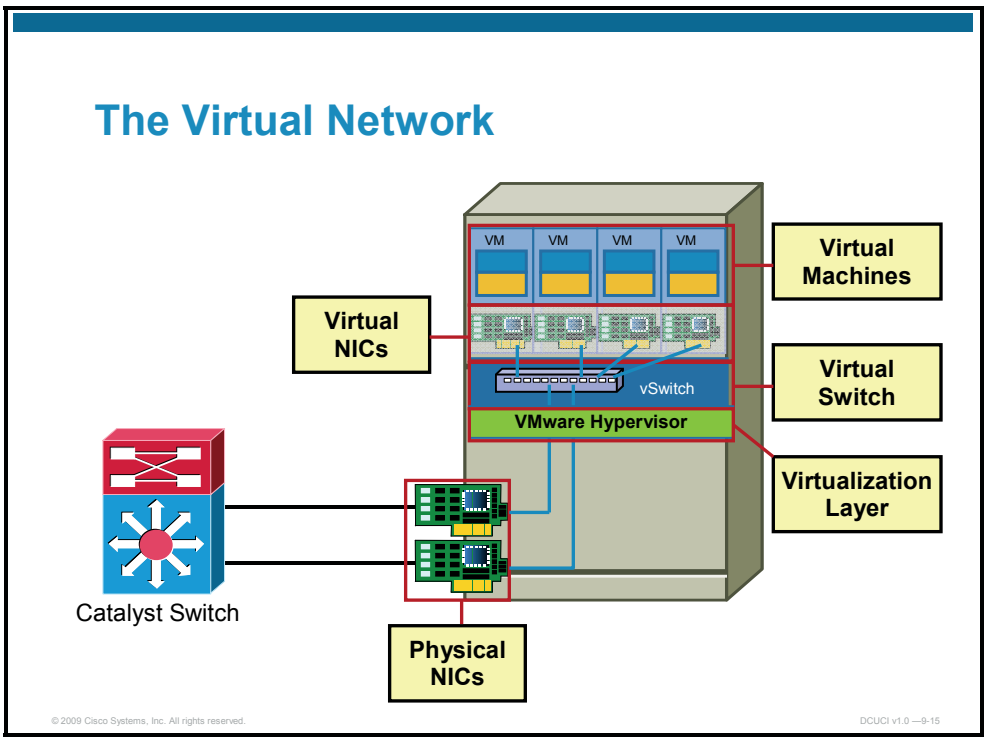

- **Physical networks:** Collections of devices for the sharing of resources.
- **Virtual networks:** Collections of virtual devices running on the same system for the sharing of resources.
- **Physical Ethernet Switch:** Used to manage network traffic between nodes. Contains multiple ports that connect to individual devices or uplink to other switches. Maintains a table of connected hosts in order to better facilitate network traffic flow.
- **Virtual Ethernet Switch (vSwitch):** Similar to a physical switch. Maintains a table of connected devices to assist with passing frames. Capable of being uplinked to physical switch using physical network interface cards (NICs). Does not contain the advanced functionality of a physical switch.
- **Host:** The physical machine that VMs run on.
- **VMware ESX Host:** A VMware host running VMware's operating system.
- **Guest:** A VM running on a host machine.
- **Virtual Center (vCenter):** External server running specialized software used for configuration and monitoring of one or more virtual servers.
- **vMotion:** VMware's live migration ability for VMs. This allows a working server to be transferred from one physical device to another without interruption (typically when testing, one ping packet will be dropped during vMotion).
- **VMWare cluster:** Two or more VMware ESX Hosts clustered together for the purpose of enabling high availability or DRS.
- **DRS:** Allows Virtual Center to monitor physical resources and migrate VMs using vMotion to spread utilization.

https://t.me/learningnets

- **High availability:** Allows Virtual Center to monitor VMware ESX Hosts for hardware failure. Virtual machines running on a failed host are restarted on available host servers.
- **Port Group:** A subset of ports on a vSwitch.
- **vmKernal:** The VMware operating system which controls and virtualizes the physical resources. Can either be installed on disk (internal or boot from SAN) or is embedded in hardware as in the case of VMware ESX 3i.

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

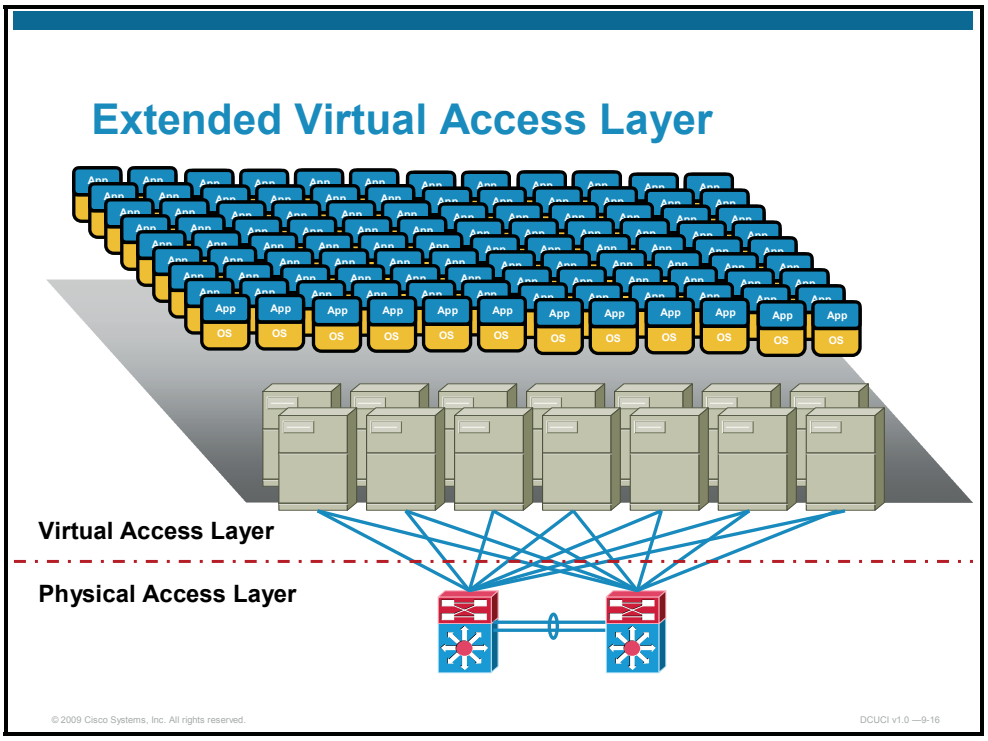

VMware extends the network access layer into the physical server. The virtual access layer can consist of hundreds or thousands of nodes networked through software switches and can be totally invisible to standard networking tools. These extended networks do not have the same controls (such as ACL and QoS) as physical networks.

https://t.me/learningnets

Virtual Storage Networking with VMware

This topic introduces VMware storage network virtualization.

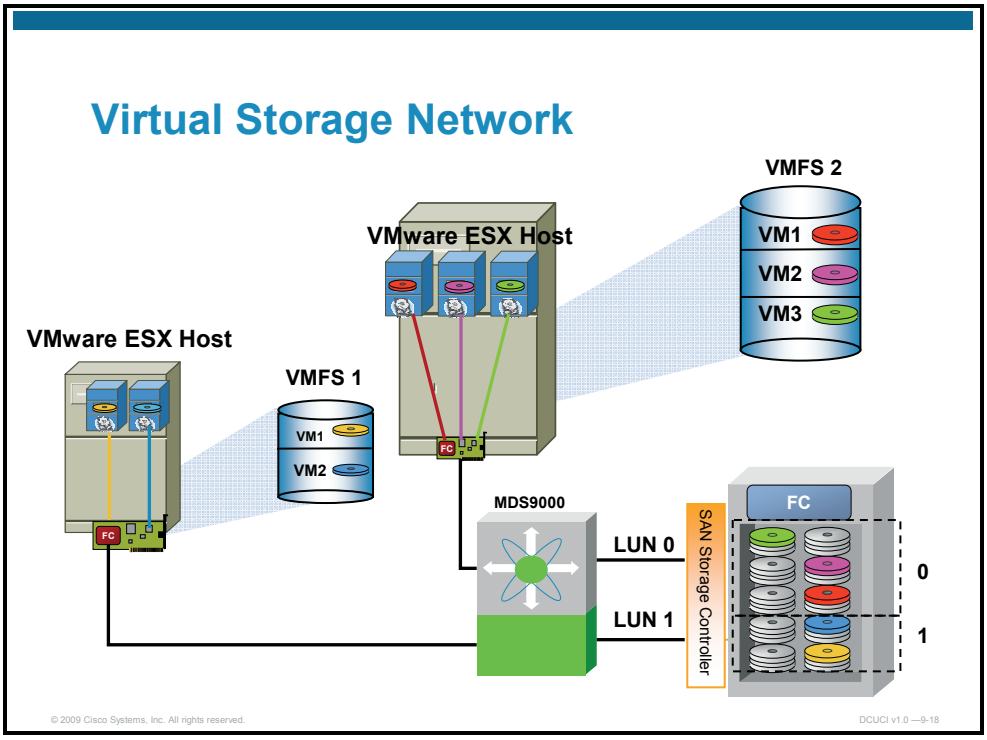

Traditionally, storage logical unit numbers (LUNs) are presented to the VMware ESX Host and then formatted as Virtual Machine File System (VMFS) volumes. Each volume can contain one or more VMs, which are stored as files on the VMFS volume:

- LUNs are masked and zoned to the VMware ESX Host, not the VM.
- LUNs are formatted by VMware ESX as Virtual Machine File System volumes.
- VMs (operating system and data) are stored as files on VMFS volumes.
- Virtual disks are presented to the VMs as Small Computer Systems Interface (SCSI) LUNs using a virtual SCSI hardware adapter. This is common for all types of storage presented to the VMware ESX Host.
- While there are other methods, this is the most typical.

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

Virtual Storage Network Management

- Different management from traditional SAN environments.
- LUN to server assignment by VMware admin not storage admin:
  - VMware admin may assign one or more VMs to a single physical LUN
  - Different from SAN single-initiator zoning
- VSANs and SAN QoS can't be applied to individual VMs

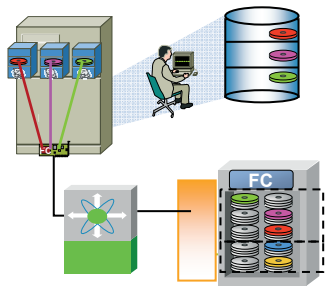

© 2009 Cisco Systems, Inc. All rights reserved.

This method of control is different from traditional SAN environments. Traditionally, the storage admin assigns LUNs and the SAN admin zones individual devices to their LUN in a one-to-one fashion. In a virtualized environment, the VMware admin may assign one or more VMs to a single physical LUN.

VMware ESX Server storage virtualization allows VMs to access underlying physical storage as though it were "just a bunch of disks" (JBOD) SCSI within the VM–regardless of the physical storage topology or protocol. In other words, a VM accesses physical storage by issuing read and write commands to what appears to be a local SCSI controller with a locally-attached SCSI drive. Either an LSILogic or BusLogic SCSI controller driver is loaded in the VM so that the guest operating system can access storage exactly as if this were a physical environment.

When an application within the VM issues a file read or write request to the operating system, the operating system performs a file-to-block conversion and passes the request to the driver. However, the driver in a VMware ESX Server environment does not talk directly to the hardware; instead, the driver passes the block read/write request to the VMkernel where the physical device driver resides and then the read/write request is forwarded to the actual physical hardware device and forwarded to the storage controller. In previous versions of VMware ESX, the physical device drivers were not loaded in the kernel, which created an extra leg in the journey from the VM to the physical storage. The integration of the drivers into the kernel in VMware ESX Server 3 thereby removes an extra translation layer and improves I/O performance.

The storage controller may be a locally-attached RAID controller or a remote, multi-pathed SAN device. The physical storage infrastructure is completely hidden from the VM. To the SAN, however, the converse is true: VMs are completely hidden from the physical storage infrastructure. The storage controller sees I/O requests that appear to originate from a VMware ESX Server; all storage bus traffic from VMs on a particular physical host appears to originate from a single source.

https://t.me/learningnets

There are two ways to make blocks of storage accessible to a VM:

- Using an encapsulated, VMware File System (VMFS)-hosted VM disk file
- Using a raw LUN formatted with the operating system's native file system

VMFS

The vast majority of (unclustered) VMs use encapsulated disk files stored on a VMFS volume.

Note VMFS is a high-performance file system that stores large, monolithic virtual disk files and is tuned for this task alone.

To understand why VMFS is used requires an understanding of VM disk files. Perhaps the closest analogy to a VM disk file is an .ISO image of a CD-ROM disk, which is a single, large file containing a file system with many individual files. Through the virtualization layer, the storage blocks within this single, large file are presented to the VM as a SCSI disk drive, made possible by the file and block translations described above. To the VM, this file is a hard disk, with physical geometry, files, and a file system; to the storage controller, this is a range of blocks.

A VM disk file is, for all intents and purposes, the hard drive of a VM. This file contains the operating system, applications, data, and all the settings associated with a typical conventional hard drive. If an administrator were to delete a VM disk file, it would be analogous to throwing a physical hard drive in the trash – the data, the operating system, the applications, the settings, and even blocks of storage would be lost. By the same token, if an administrator were to copy a VM disk file, an exact duplicate of the VM's hard drive would be created for use as a backup or for cloning the particular configuration.

Unlike Windows and Linux operating systems, VMware ESX Server does not lock a LUN when it is mounted–a simple fact that is the source of both power and potential confusion in a VMware ESX Server environment. When configuring a switched SAN topology, it is critical to use zoning, selective storage presentation, or LUN masking to limit the number of physical servers (non-VMware ESX Server) that can see a particular LUN. Without limiting which physical–Windows or Linux–servers can see a LUN, locking and LUN contention will quickly cause data to become inaccessible or inconsistent between nodes.

VMFS is inherently a distributed file system, allowing more than one VMware ESX Server to view the same LUN. Unlike Windows/New Technology File System (NTFS) or Linux/ext3, VMware ESX Server/VMFS supports simultaneous access by multiple hosts. This means that, while numerous VMware ESX Server instances may view the contents of a VMFS LUN, only one VMware ESX Server may open a file at any given moment. To a VMware ESX Server and VMFS, when a VM is powered on, the VM disk file is locked.

While VMotion is described in detail later, it might be helpful to explain that in a VMotion operation, the VM disk file remains in place on the SAN, in the same LUN; file ownership is simply transferred between VMware ESX Server hosts that have access to the same LUN.

The distributed nature of VMFS means that when configuring the SAN to which VMware ESX Server is attached, zoning should be configured to allow multiple VMware ESX Servers to access the same LUN where the VMFS partition resides. This may be out of the ordinary for the SAN administrator.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

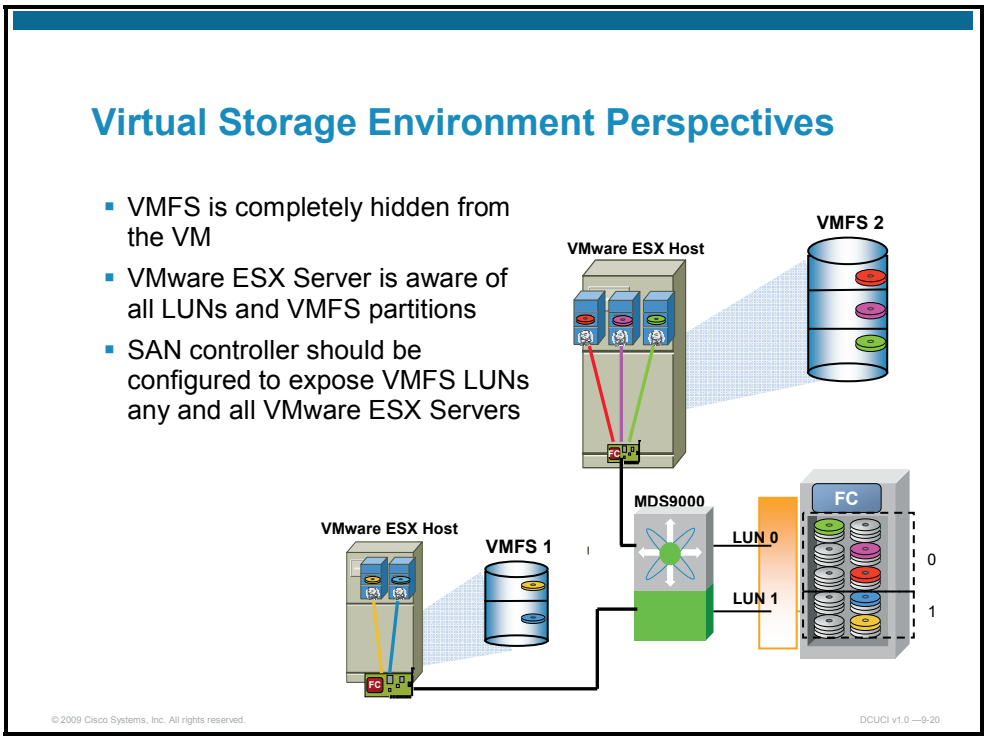

There are many perspectives to a virtualized storage environment:

- To VMs, VMFS is completely hidden. A VM is not aware that the storage that it sees is encapsulated in a larger file within a VMFS volume.
- To VMware ESX Server, multiple LUNs and multiple VMFS partitions may be visible. VMware ESX Server can run multiple VMs from multiple SAN devices; that is, VMware ESX Server does not require or prefer that all VMs are in the same VMFS volumes.
- To the SAN, the controller should be configured to expose the LUN containing the VMFS volume to any and all VMware ESX Servers that might be involved in VMotion operations for a particular VM.

LUN Performance Considerations

When constructing a LUN for VMFS volumes, you should follow some basic storage rules that apply to VMFS and a few that require further consideration.

As with any LUN, more spindles mean more concurrent I/O operations. When planning a storage configuration that maximizes storage performance, you should deploy as many spindles as practical to create your VMFS LUN.

Remember that the VMFS volume will host multiple VMs, which has two effects on LUN performance:

- Since a single VMFS volume may have multiple VMware ESX Servers and each VMware ESX Server may have multiple VMs within the same partition, the I/O loads on a VMFS-formatted LUN can be significantly higher than the loads on a single-host, single-operating system LUN.

https://t.me/learningnets

- Since many VM disk files are likely to be stored within a single VMFS volume, the importance for fault tolerance on this LUN is amplified. Always employ at least the level of fault tolerance used for physical machines. Fault tolerance becomes even more of a concern if a larger VMFS volume is created from multiple, smaller VMFS extents within VMware ESX Server. Should any one extent fail, all data within that extent would be lost, whereas information on the remaining extents would remain available. Therefore, measures like Redundant Array of Independent Disks (RAID) technology and stand-by drives should be considered standard as part of any VMFS LUN.

From a performance perspective, tuning an array for a particular application may not be as effective with VMs as with physical machines. Since VM storage is abstracted from the VM and, typically, encapsulated in a VM disk file within a VMFS volume, it is probable that the same parameters that enhanced database performance in an NTFS partition will not deliver the same gains in a virtualized environment. As a result, at this time there are no recommended application-specific tuning parameters for a VMFS formatted LUN.

Using Raw Device Mapping

Raw device mapping can allow a VM to access a LUN in much the same way as a non-virtualized machine. In this scenario, where LUNs are created on a per-machine basis, the strategy of tuning a LUN for the specific application within a VM may be more appropriate.

Since raw device mappings do not encapsulate the VM disk as a file within the VMFS file system, LUN access more closely resembles the native application access for which the LUN is tuned.

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

VMware Environment Management

This topic introduces VMware environment management.

VMware Components

VMware environments consist of physical servers (running the VMware operating system) known as VMware ESX Hosts.

Optionally, a vCenter Server can be used to manage multiple hosts.

The vSphere Client is a GUI tool used to manage VMware ESX Hosts and vCenter Servers.

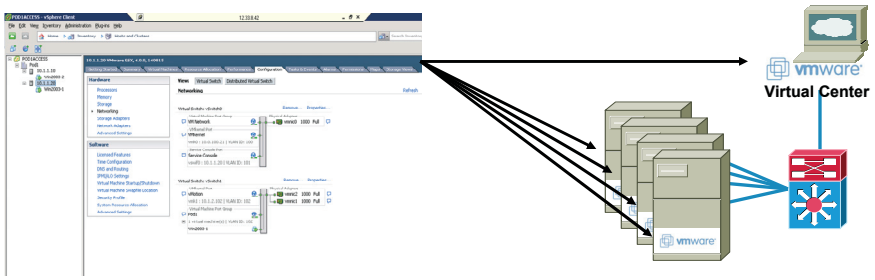

© 2009 Cisco Systems, Inc. All rights reserved.

VMware environments consist of physical hosts running the VMware operating system which can be used to deploy VMs. These hosts can be independent or clustered for advanced services. Host management can be performed using a Linux-based management interface through SSH or through the vSphere GUI.

VMware offers a vCenter server as an optional component for virtualized environments. This server allows configuration and management of multiple VMware hosts and provides the capability to utilize advanced functionality such as high availability and DRS. The vCenter server can run as a physical or VM and is managed using the vSphere client.

https://t.me/learningnets

VMware Hosts

A VMware Host is a physical server running the VMware operating system on bare metal.

The VMware operating system provides functionality for the creation of virtual servers, known as VMs.

These hosts can be managed utilizing a Linux-based command-line interface (CLI) interface or the vSphere GUI Client.

© 2009 Cisco Systems, Inc. All rights reserved.

VMware ESX Hosts are the main component of VMware environments. These servers run VMware on bare metal. This VMware operating system provides the functionality to create virtual servers with virtual hardware that shares the pool of available physical hardware.

https://t.me/learningnets

vCenter Server

vCenter Server provides a single point of management for multiple VMware hosts as well as the tools used for advanced VMware functionality.

vCenter Server can run as a physical or VM.

vCenter Server is required for features that utilize more than one host such as:

- VMotion
- High availability
- DRS

© 2009 Cisco Systems, Inc. All rights reserved.

vCenter Server is a central configuration point for multiple VMware hosts. vCenter Server is licensed separately from the VMware Hosts and runs as a physical or VM. vCenter Server is required for advanced functionality such as: high availability, DRS, and VMotion.

https://t.me/learningnets

vSphere Client

The vSphere Client is the preferred management tool for VMware ESX Servers as well as the management tool for the vCenter Server.

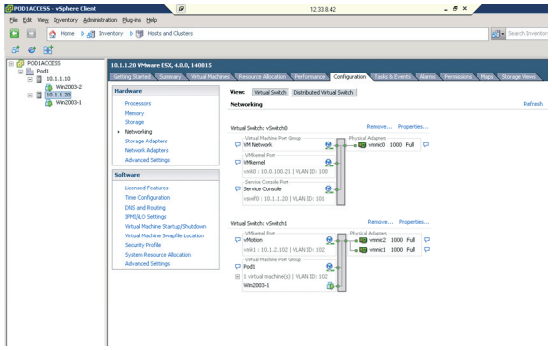

© 2009 Cisco Systems, Inc. All rights reserved.

The vSphere Client is the preferred method for configuring VMware environments. The client is a windows based GUI tool which allows configuration of single hosts via direct connection or multiple hosts via connection to a licensed vCenter Server.

https://t.me/learningnets

vSphere Client (Cont.)

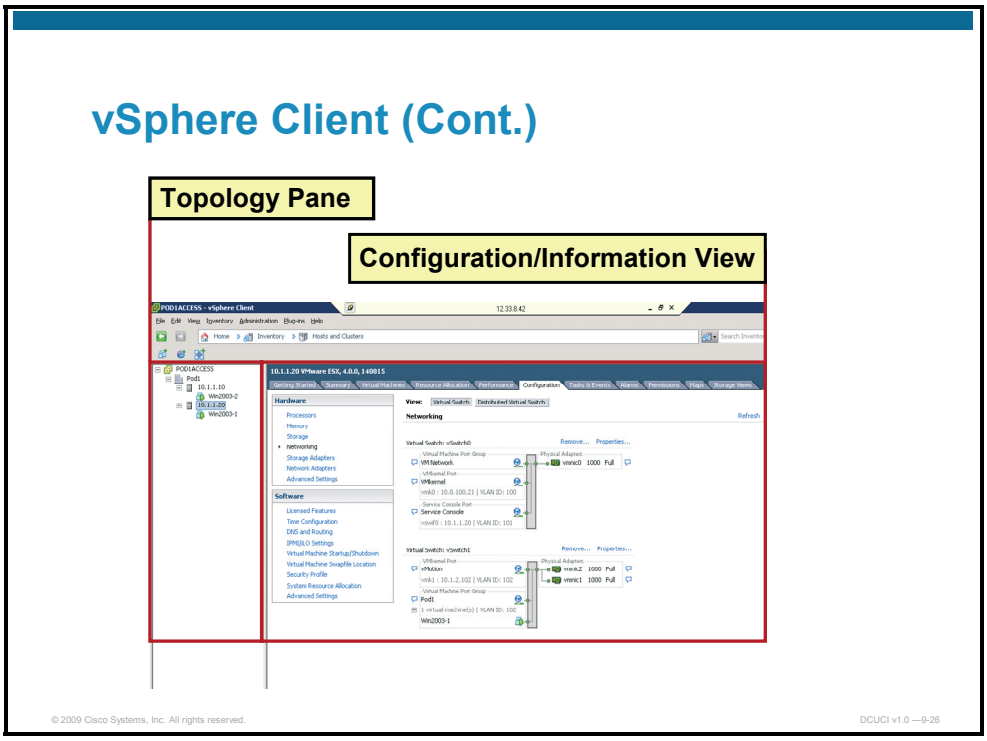

© 2009 Cisco Systems, Inc. All rights reserved.

The vSphere Client is navigated using a series of separate panes and tabs. The main view is the Hosts and Clusters view, which shows Datacenters, Folders, Hosts, and Virtual Machines. In the right hand pane, information regarding the currently selected object is displayed. Various tabs are available at the top of the Configuration/Information pane depending on the type of object selected.

https://t.me/learningnets

Summary

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

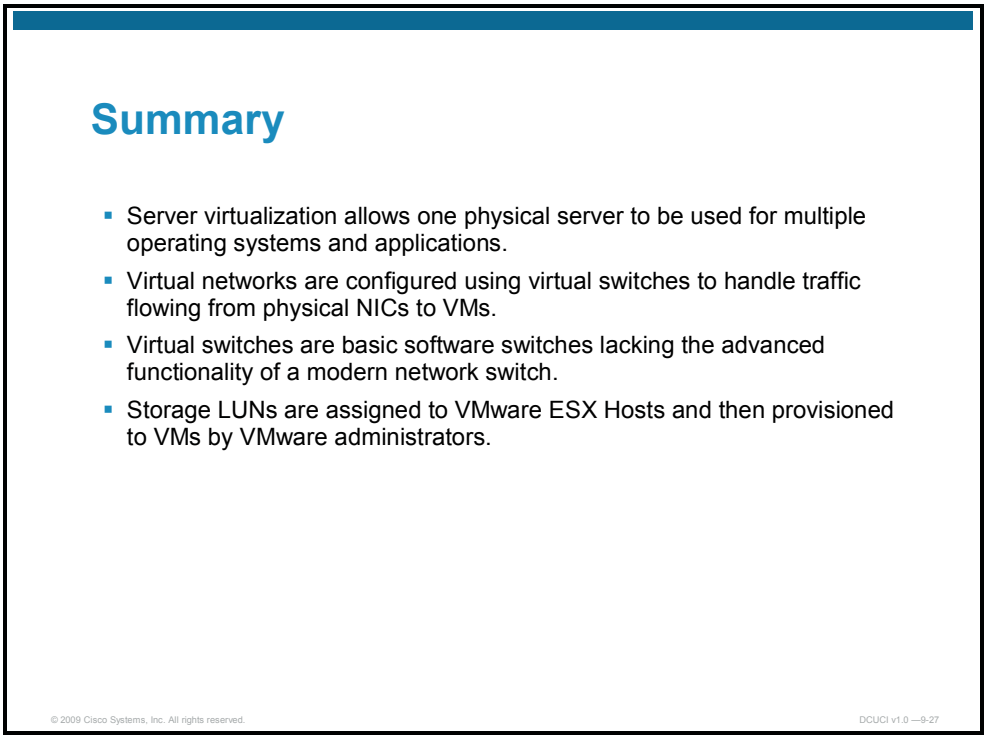

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

VMware Ethernet Networking

Overview

This lesson describes the VMware networking environment.

Objectives

Upon completing this lesson, you will be able to describe the VMware networking environment. This includes being able to meet these objectives:

- Introduce the virtual access layer
- Describe the VMware vSwitch
- Describe VMware network configuration
- Introduce Distributed Virtual Switching

https://t.me/learningnets

Virtual Access Layer

This topic introduces the virtual access layer.

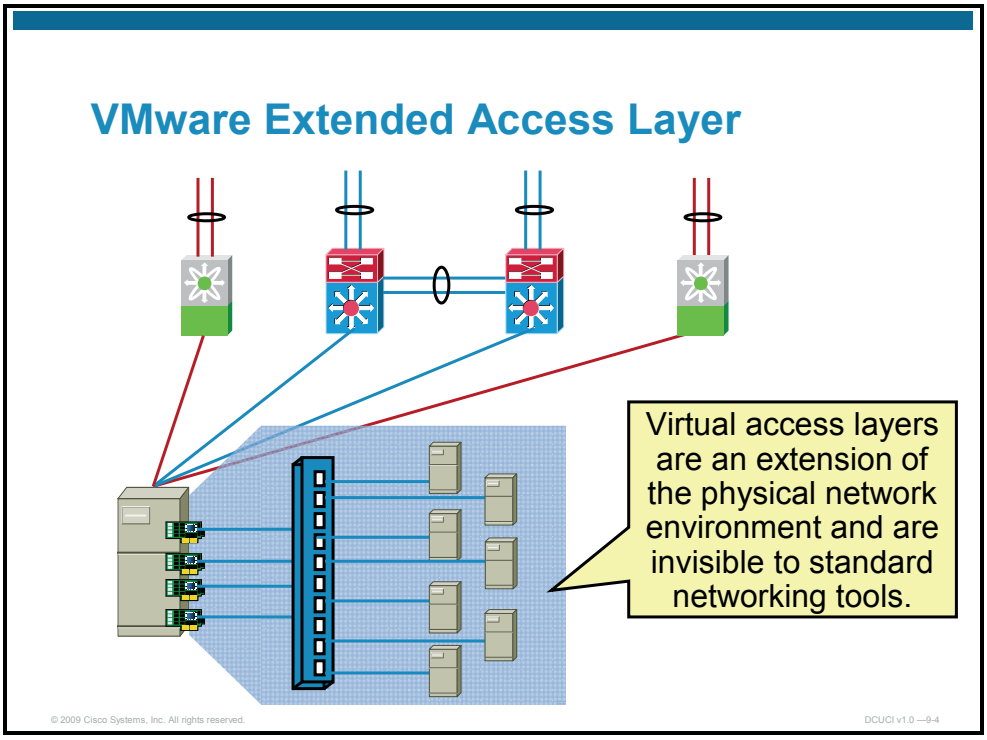

In VMware environments, the VM networking layer extends the access layer into the VMware ESX server. This is the case in both ToR and EoR deployments. This virtual access layer does not have the same security, management, or visibility of physical access networks.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

VMware Virtual Network Overview

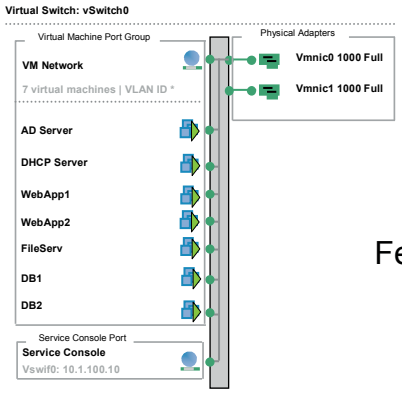

Virtual access layer:

- A series of virtual switches, servers, and network interfaces residing inside a VMware ESX Host
- Allows the sharing of physical network resources between virtual machines (VMs)

Features:

- L2 Ethernet switching
- VLAN segmentation (using port groups)
- Outbound rate limiting

© 2009 Cisco Systems, Inc. All rights reserved.

VMware extends the access layer into the VMware ESX Host using software switches configurable through VMware tools. The figure shows an example of a vSwitch within VMware. The figure shows a vSwitch with two uplink ports, two internal port groups, and several Virtual Machine Ports.

https://t.me/learningnets

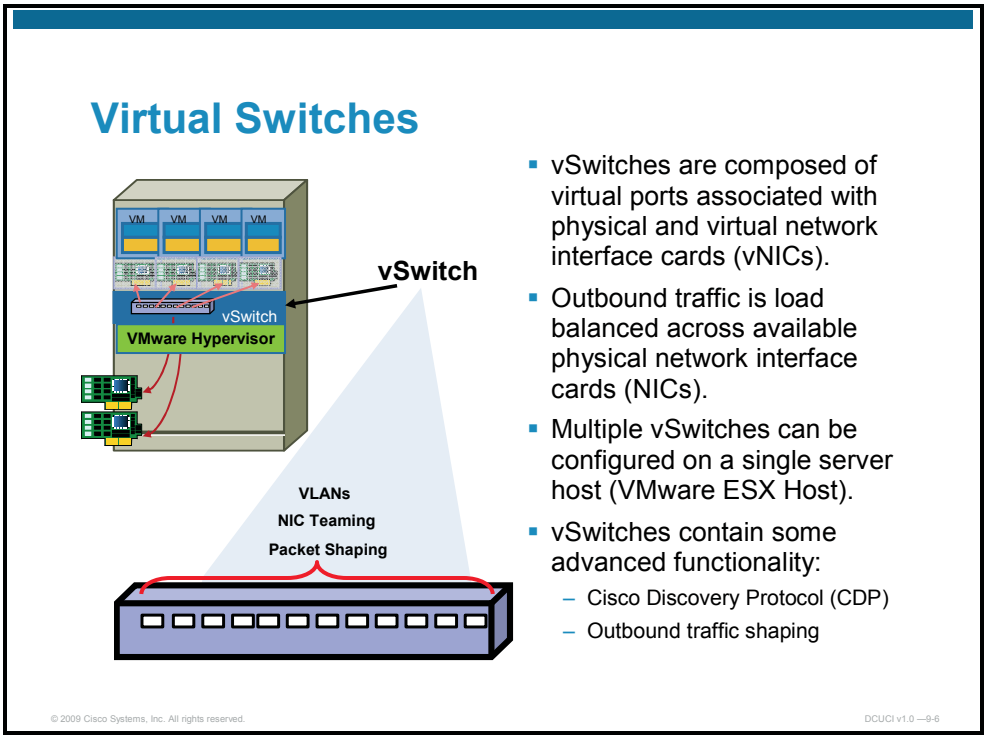

Virtual switches are software devices within the virtualized operating environment that handle the networking aspect of the VMs. VMs have vNICs that connect to vSwitches for LAN communication.

A virtual switch has three types of connections available:

- **Service console port:** Allows access to physical network through a vSwitch.
- **vMotion and IP storage port:** Used by the VMkernel TCP/IP stack for functions such as vMotion and access to network attached storage (NAS).
- **Virtual machine port group:** Allows VMs to connect to the network.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

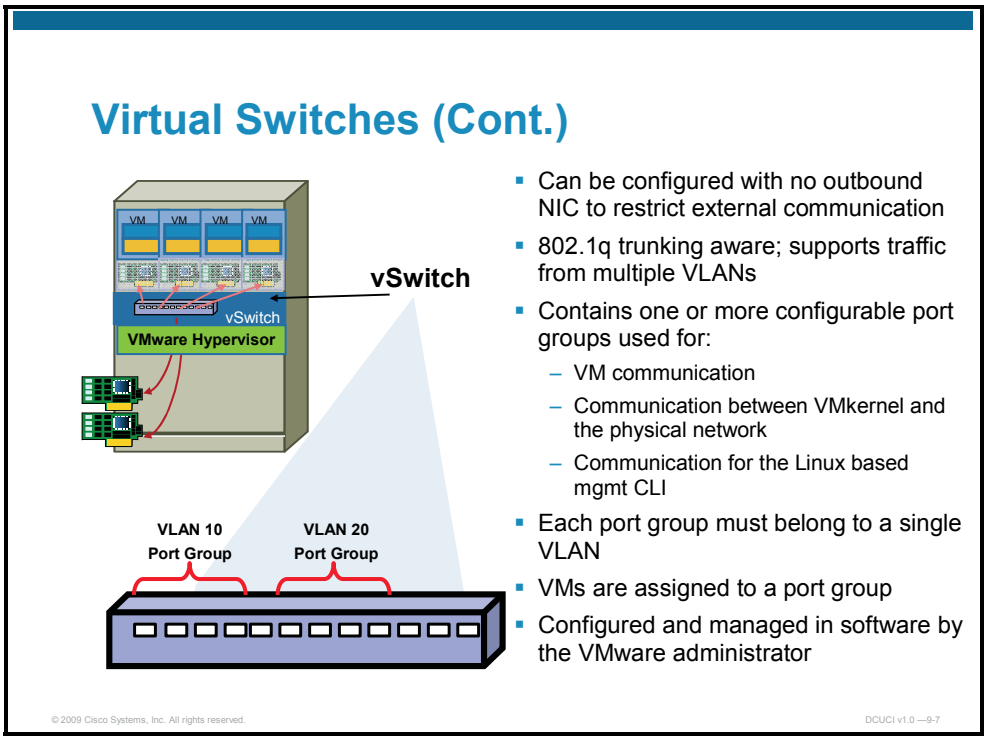

vSwitches contain virtual port groups. These port groups are either used for communication between VMs or for outbound communication from the virtual operating system software kernel. The software kernel (VMkernel) needs direct access to the network for features such as vMotion and network-based storage.

Each virtual switch can have different connection types:

- Service console
- VMkernel
- Virtual machines

Port groups on virtual switches define:

- VLANs
- Security
- Traffic shaping
- NIC teaming

https://t.me/learningnets

Vi3 Networks

vSwitches are configured on each VMware ESX Host.

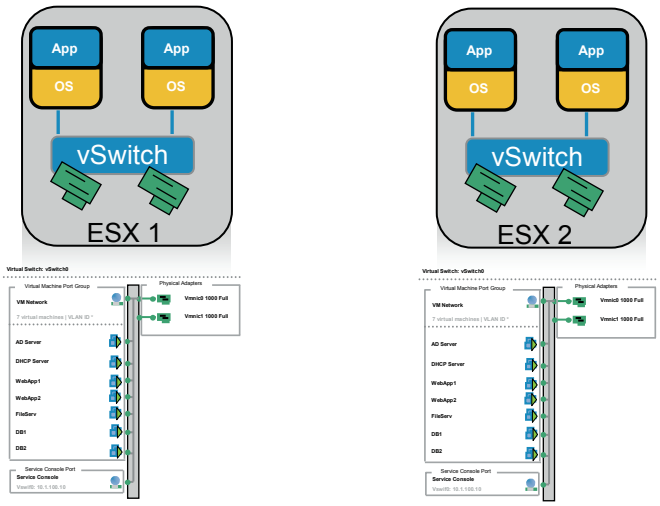

© 2009 Cisco Systems, Inc. All rights reserved.

One or more vSwitches can be configured individually on each VMware ESX Host. If vMotion or clustering is used, each participating hosts vSwitch must be identically configured. This is a manual process that must be performed on each host in the cluster.

https://t.me/learningnets

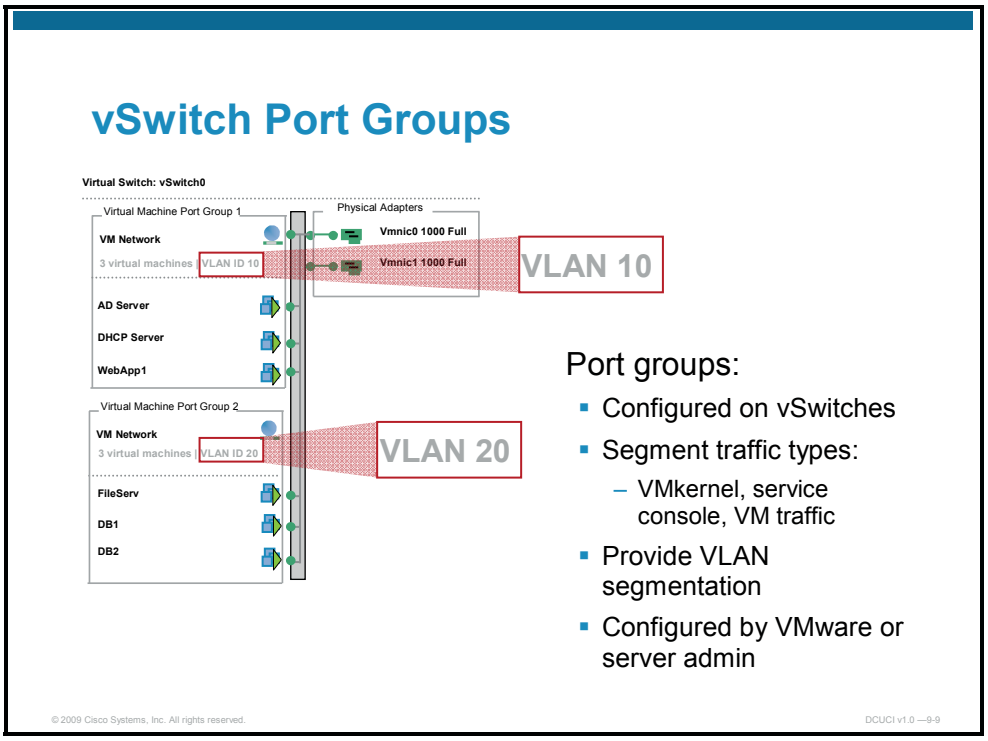

Port groups are configured on vSwitches to segregate traffic based on type or VLAN. The figure shows a vSwitch with two port groups.

https://t.me/learningnets

VMware Network Configuration

Network policies such as VLAN, traffic shaping and Rate Limiting, which are typically designated by network teams, are configured by the VMware/server team.

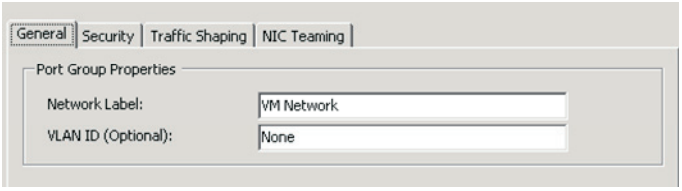

© 2009 Cisco Systems, Inc. All rights reserved.

VMware network settings are configured using VMware tools. The figure shows some of the networking options of VMware ESX.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

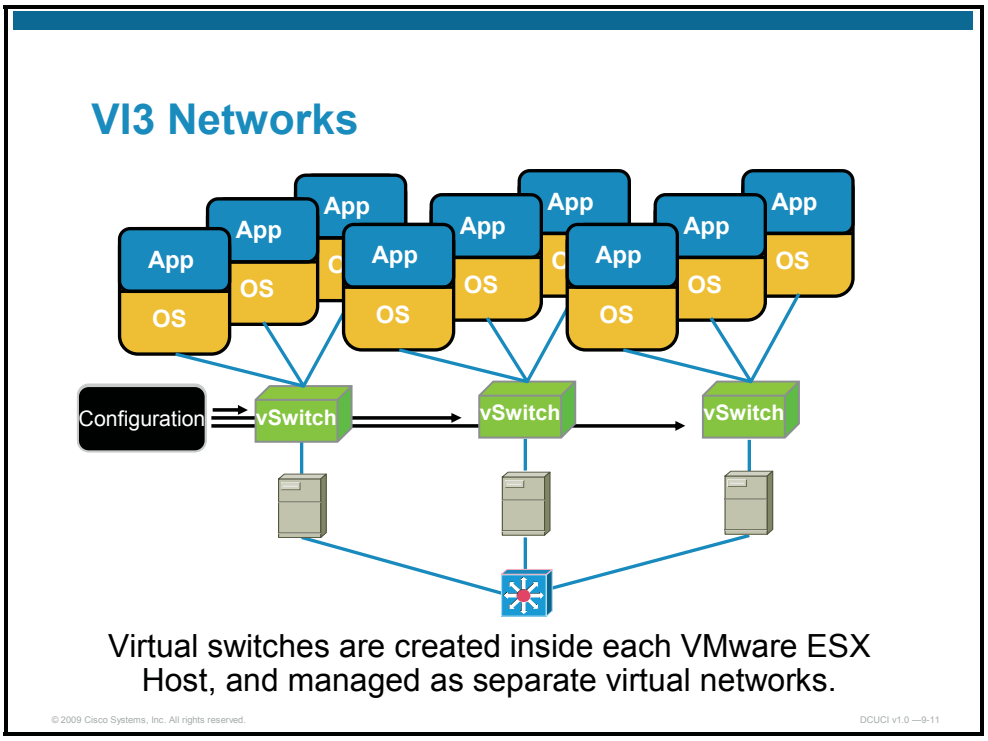

In the Vi3 environment, each vSwitch must be configured individually. Each VMware ESX Host may have one or more vSwitches.

https://t.me/learningnets

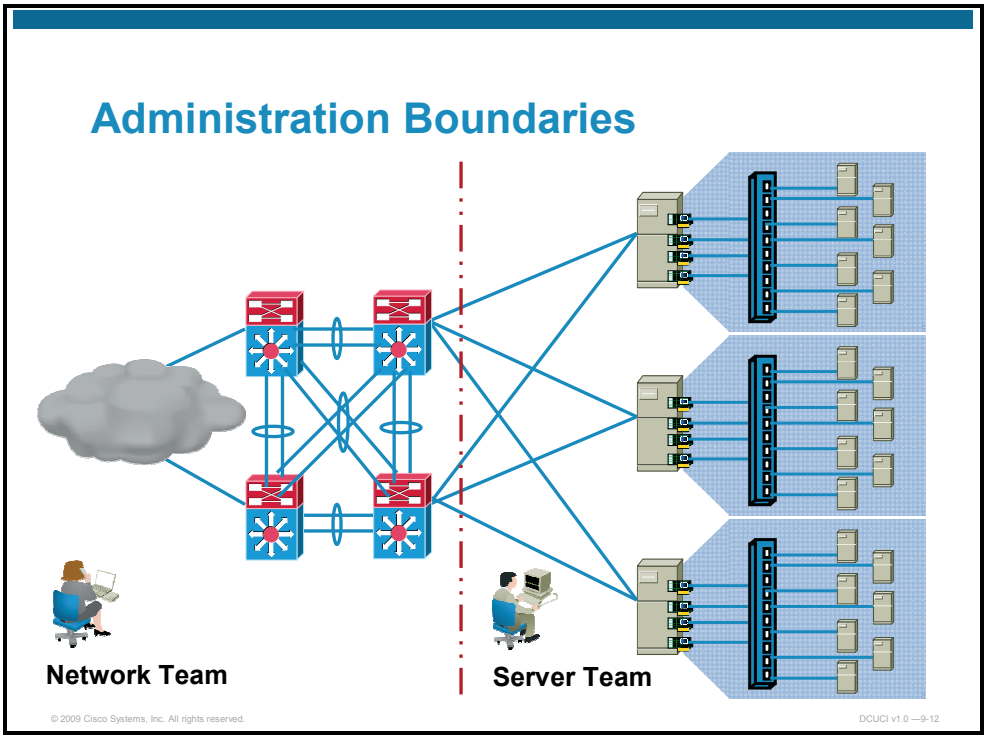

The figure shows the administration boundaries of the VMware environment. These boundaries differ greatly from those that existed before VMware. In non-VMware environments, LAN teams have control of the network from the entrance point to the individual server. In VMware environments, entire virtual networks exist behind the physical server and are managed by VMware or server teams.

https://t.me/learningnets

The VMware vSwitch

This topic explains the VMware vSwitch.

vSwitch Characteristics

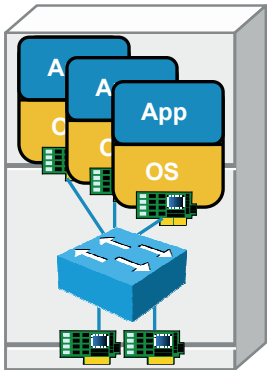

VMware vSwitches operate in a similar fashion to a physical Layer 2 switch:

- Forwards frames based on Layer 2 information (using MAC address)
- Maintains a MAC address table
- Supports internal switching for VMs
- Can tag frames with a VLAN ID using 802.1Q tagging
- Can support port channels (for active/active NIC teaming only)

© 2009 Cisco Systems, Inc. All rights reserved.

A VMware vSwitch is a software construct that provides both internal VM communication and outbound uplinks for VMs. vSwitches operate in a fashion similar to that of a physical switch.

https://t.me/learningnets

Configuration Limits

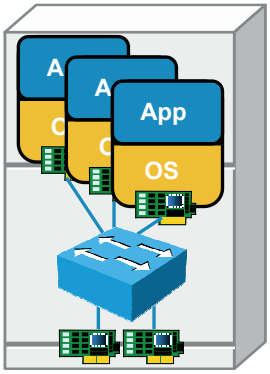

- Multiple vSwitches per VMware ESX Host
- Up to 1016 virtual ports per vSwitch
- Up to 32 vmNICs (outbound adapters) per vSwitch

© 2009 Cisco Systems, Inc. All rights reserved.

Virtual switches can contain up to 1016 virtual ports, a max of 32 of these can be for outbound adapters.

Note	A NIC in VMware is represented by an interface called a vmNIC. The vmNIC number is allocated during VMware installation.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

MAC Address Table

vSwitch MAC address tables record MAC address and port numbers for VMs.

Frames to unknown MAC addresses are forwarded to physical NIC uplinks.

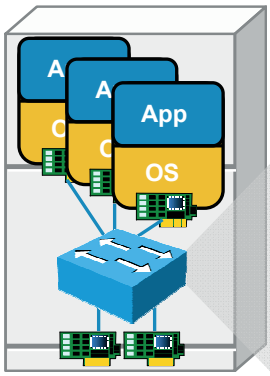

VM1 vPort 3	xx:xx:xx:xx:xx:xx
VM2 vPort 4	xx:xx:xx:xx:xx:xx
VM3 vPort 5	xx:xx:xx:xx:xx:xx

© 2009 Cisco Systems, Inc. All rights reserved.

vSwitches maintain a MAC table of all connected VMs. vSwitches never learn addresses on uplink ports and instead forward all unknown unicast messages using one of the available uplink NICs.

https://t.me/learningnets

Loop Prevention

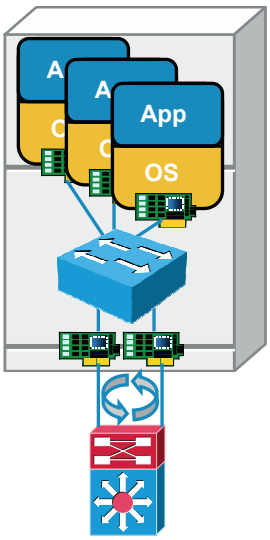

vSwitches do not participate in Spanning-Tree Protocol (STP).

Loops are prevented by passing outbound traffic on only one physical NIC at a time:

- Active/passive NIC teaming only the active link is used
- In active/active NIC teaming, each VM utilizes only one uplink port at any given time
- vSwitches do not flood frames to available uplinks
- vSwitches can not be connected to one another; two port types are supported:
  - vmNIC (outbound physical NIC)
  - vNIC (virtual VM network adapter)

© 2009 Cisco Systems, Inc. All rights reserved.

vSwitches are designed to avoid loops and thus do not participate in STP. Because vSwitches can not be uplinked to one another, internal loops are avoided. Broadcast, floods, and unknown unicast traffic are forwarded to only one available uplink at a time, thus avoiding multiple frame copies to the upstream switch.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

NIC Teaming

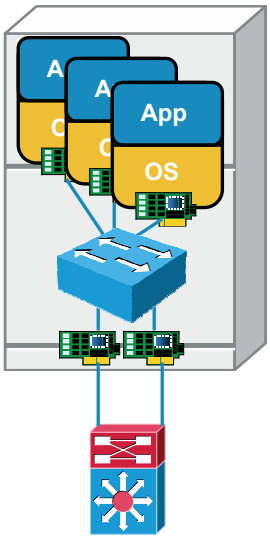

VMware leverages NIC teaming for physical NICs in order to provide fault tolerance and performance benefits:

- VMware drivers and software provide NIC teaming functionality with no vendor software requirements (some third-party drivers are available as options)
- Two teaming configuration types:
  - Active/passive
  - Active/active load balancing:
    - VM Port ID
    - VM MAC hash
    - IP-based hash (source and destination)

© 2009 Cisco Systems, Inc. All rights reserved.

NIC teaming allows multiple NICs to be grouped for the purposes of high availability and load balancing. VMware uses built-in drivers for NIC teaming, which allows both active/active and active/passive configurations. There are three load balancing options for teamed NICs: VM Port ID, VM MAC, and IP-based hash. IP-based hash is the most granular and typically recommended.

https://t.me/learningnets

VLAN Tagging

vSwitches support three 802.1Q VLAN tagging methods:

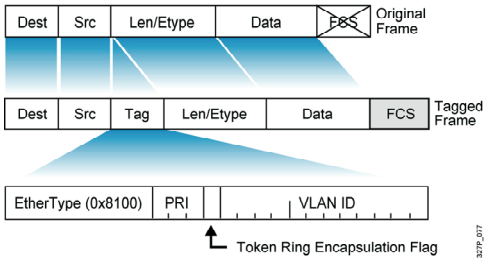

- External Switch Tagging (EST):
  - Frames are left untagged and any tagging occurs in the physical network
- Virtual Guest Tagging (VGT):
  - VMs tag frames and the vSwitch preserves the tag
- Virtual Switch Tagging (VST):
  - The vSwitch tags frames based on port group VLAN
  - This is the preferred method

© 2009 Cisco Systems, Inc. All rights reserved.

vSwitches are capable of supporting VLAN IDs in one of three methods:

- **External Switch Tagging (EST):** This method leaves frames untagged and allows the external switch to handle all tagging operations.
- **Virtual Guest Tagging (VGT):** This method allows software within the guest operating system to tag the frame, and maintains that tag within the virtual network.
- **Virtual Switch Tagging (VST):** This is the preferred and default method. In this mode, the vSwitch will tag each frame based on the assigned port group VLAN.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

vSwitch Security

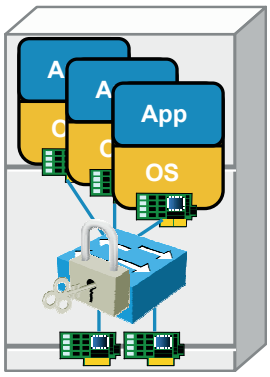

vSwitch security options:

- Promiscuous mode:
  - Enabled by default
  - Prevents VMs from seeing unicast traffic not intended for the ports MAC
- MAC address change lockdown:
  - Prevents VMs from changing their own MAC address
- Forged transmit blocking:
  - Prevents VMs from sending frames that appear to come from other nodes on the network

© 2009 Cisco Systems, Inc. All rights reserved.

VMware vSwitches offer three security features in addition to VLAN segmentation:

- **Promiscuous mode:** Enabled by default, promiscuous mode prevents VMs from seeing unicast traffic not intended for the ports virtual MAC address.
- **MAC address lockdown:** Prevents VMs from changing the assigned virtual MAC address.
- **Forged transmit blocking:** Ensures that VMs are not able to send frames that appear to come from other network nodes.

https://t.me/learningnets

VMware Network Management

VMware/server teams manage the vSwitch and VMware ESX networking using several tools:

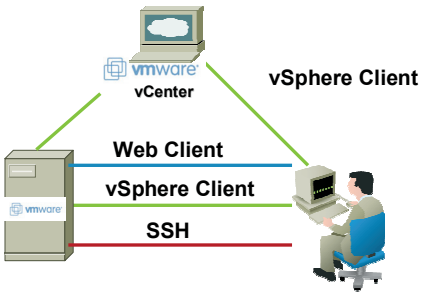

- Virtual Center (vCenter):
  - Licensed feature
  - Manages multiple hosts
  - Supports clustering
- Web Client:
  - Host and VM management via web browser
- Secure Shell (SSH):
  - Secure connection to management CLI
  - CLI based on Red Hat Linux

© 2009 Cisco Systems, Inc. All rights reserved.

VMware offers various options for managing the virtual environment and network:

- **vCenter Server:** Provides management for multiple VMware ESX Hosts and allows VMware clustering and vMotion. Virtual Center is a licensed software package.
- **vSphere Client:** This GUI client is used to attach to and manage VMware ESX Hosts or Virtual Center Management Servers.
- **Web Client:** VMware servers and Virtual Center support web based actions for many management tasks.
- **SSH:** SSH can be used to connect to the Red Hat Linux based management CLI of the VMware ESX Host. This management console is not a direct attachment to the VMware ESX kernel; it is a management shell running within the VMware kernel.

https://t.me/learningnets

VMware Network Configuration

This topic explains VMware network configuration.

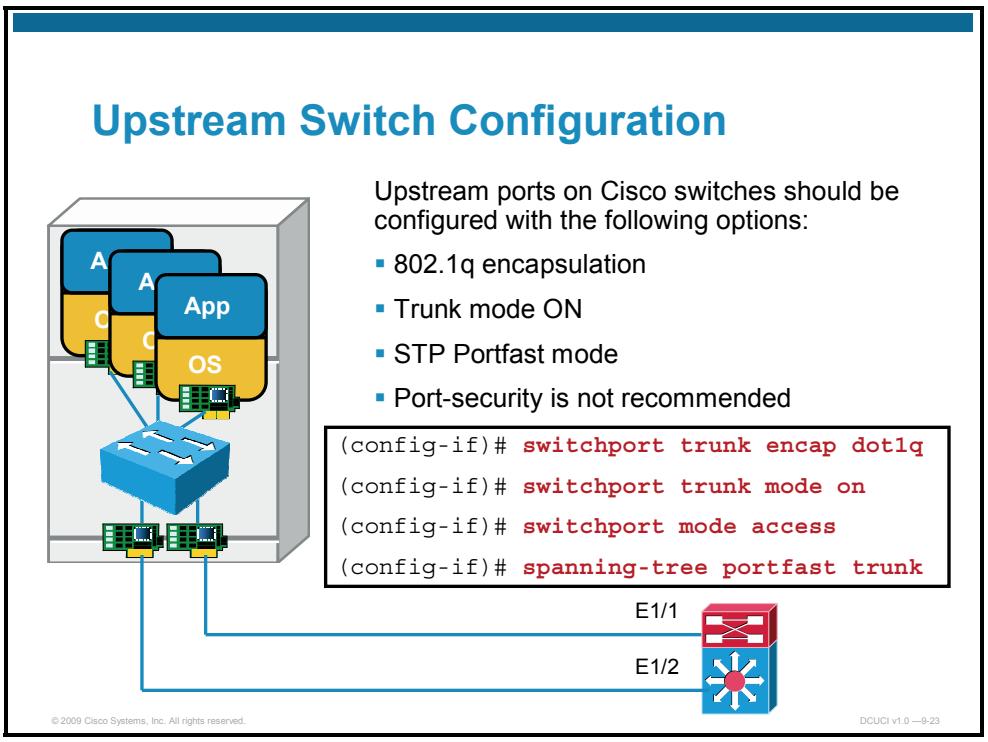

Specific configuration requirements and best practices exist for upstream switch ports. In a Cisco switching environment, the most common recommendations are 802.1Q encapsulation, trunk mode on, Portfast mode, and no port-security configuration. This configuration provides the necessary VLANs to the vSwitch while ensuring that the port activates as quickly as possible in the STP scheme.

Port-security is not recommended due to the fact that multiple VMs with individual virtual MAC addresses will be accessing the same port. For example, if port-security were set to allow a maximum of 10 MAC addresses on the port and nine VMs existed on the server, one available MAC would be consumed by the physical port and the maximum of ten would be reached. If the VMware administrator created a tenth VM and enabled its networking, the port would lock down due to the eleventh MAC, and no traffic would be able to flow through the port.

https://t.me/learningnets

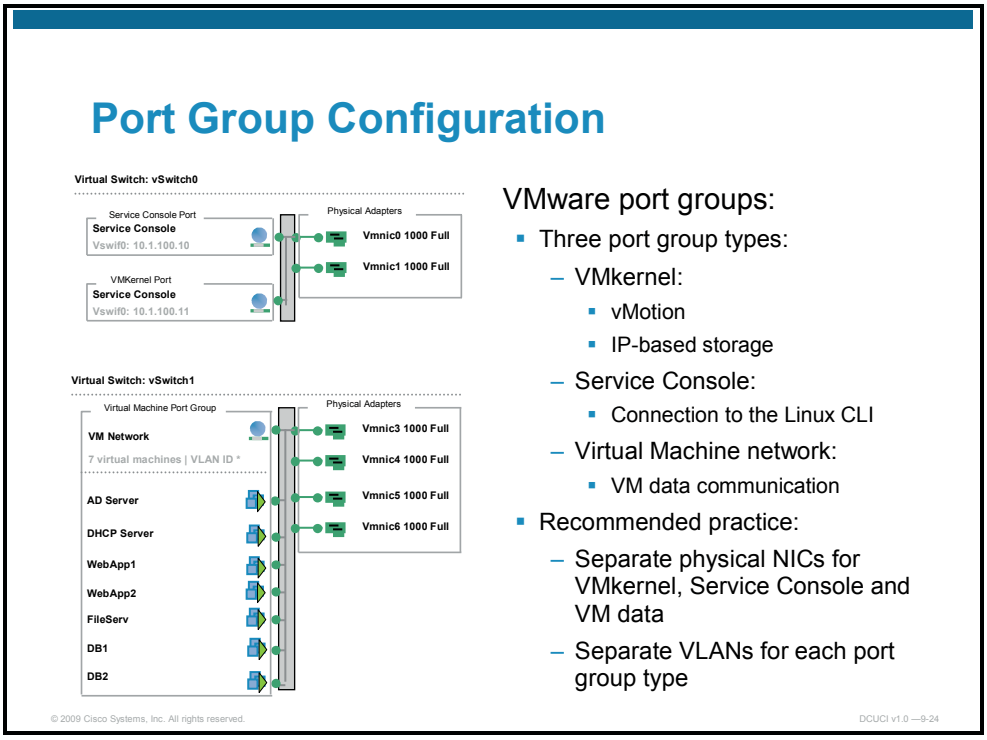

Within the VMware vSwitch, port groups are a method of segmenting traffic by VLAN and type. The three VMware port group types are:

- **VMkernel:** The VMKernel port is used to support vMotion as well as IP-based storage such as Internet Small Computer Systems Interface (iSCSI) and NAS.
- **Service Console:** The Service Console port provides outbound connectivity for the console shell.
- **Virtual Machine Network:** This port group is used to provide network connectivity for virtual machines.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

vSwitch Configuration Options

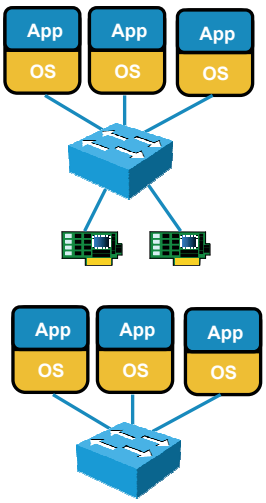

- vSwitches can be configured with 0-32 vmNICs for outbound communication
- VMware NIC teaming is used to load balance traffic across available uplinks
- Internal only vSwitches can be used for testing and traffic isolation:
  - VMs can be remotely accessed through VMware tools even with no outbound uplinks

© 2009 Cisco Systems, Inc. All rights reserved.

vSwitches support up to 32 outbound vNICs. A vSwitch can be created with no outbound vNIC in order to limit communication to the attached VMs. This configuration is very useful in test and development situations where no outbound connection is needed or desired.

https://t.me/learningnets

vSwitch Configuration Options (Cont.)

Software firewall appliances and routers can be used within VMs to create secure internal networks or demilitarized zones (DMZs).

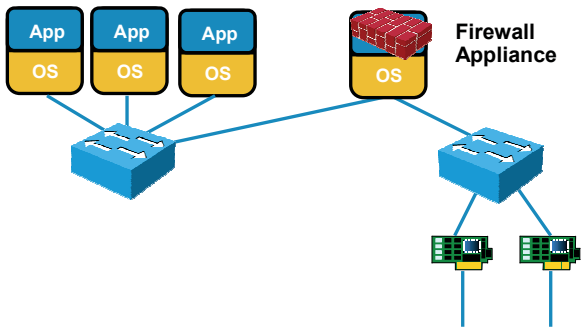

© 2009 Cisco Systems, Inc. All rights reserved.

Many configuration options exist for internal VM networks. As an example, a firewall services application can be run within a VM. This VM can be dual attached to a vSwitch with outbound connections and an internal-only vSwitch. In this configuration example, VM will provide firewall services for all inbound and outbound VM traffic for devices attached to the internal only switch.

https://t.me/learningnets

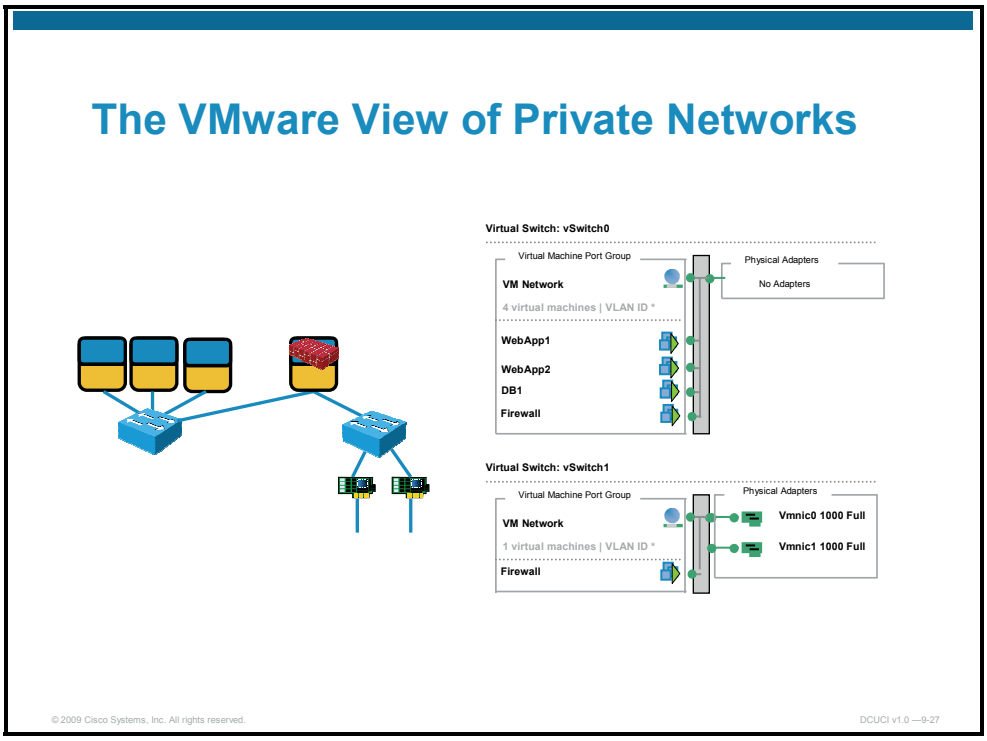

The figure shows the logical view of an internal firewall, and the VMware view of that firewall.

https://t.me/learningnets

Layer 2 Adjacency

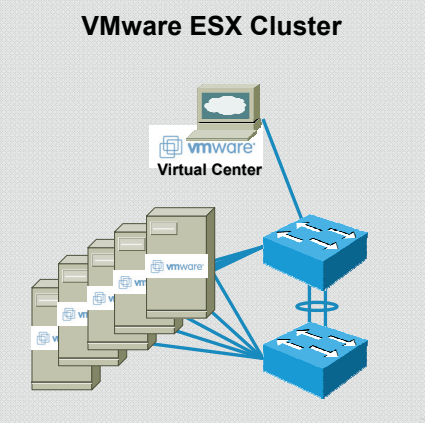

VMware advanced functionality requires Layer 2 adjacency for VMkernel ports:

- vMotion:
  - Live migration of virtual servers
- High availability:
  - For failure recovery
- DRS:
  - For better utilization and performance

© 2009 Cisco Systems, Inc. All rights reserved.

Advanced VMware services require Layer 2 adjacency. These services include high availability, Distributed Resource Scheduling (DRS), and vMotion. All VMware ESX Hosts in a cluster must be Layer 2 adjacent. Due to the fact that VMware advanced services require L2 adjacency, very large L2 domains can exist when many hosts each contain many VMs.

https://t.me/learningnets

Distributed Virtual Switching

This topic introduces Distributed Virtual Switching (DVS).

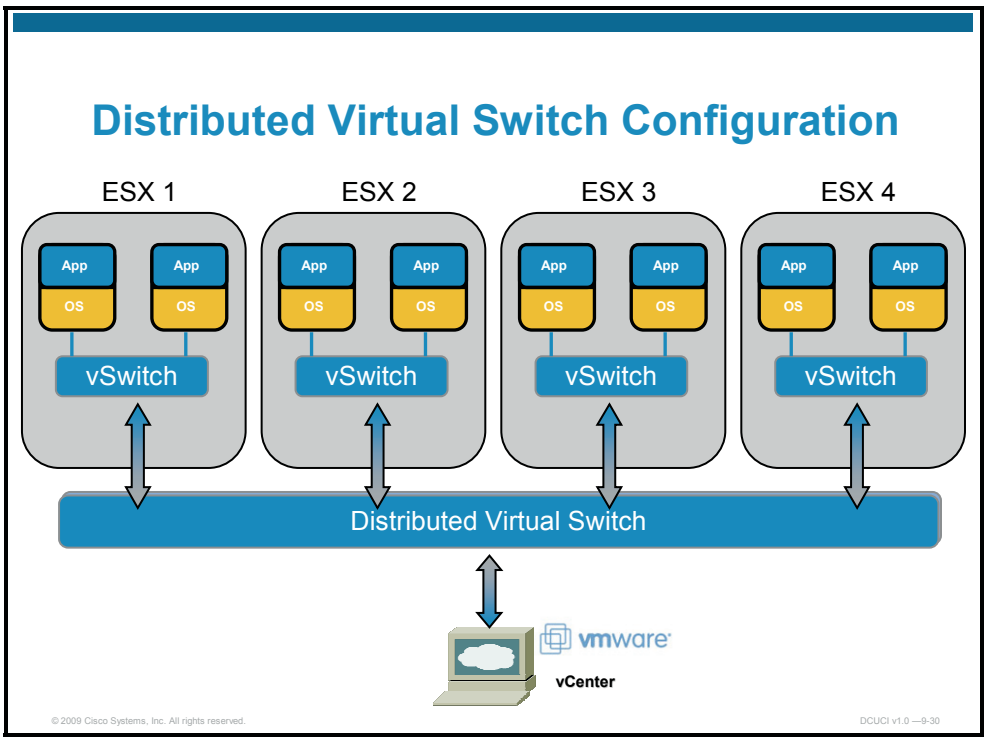

Utilizing VMware's DVS technology, multiple vSwitches within a cluster can be configured from a central point. The DVS automatically applies changes to the individual vSwitches on each VMware ESX Host. DVS is a licensed feature that is new with the vSphere release of VMware 4.0. DVS is reliant on vCenter Server and is not available for individually managed hosts.

https://t.me/learningnets

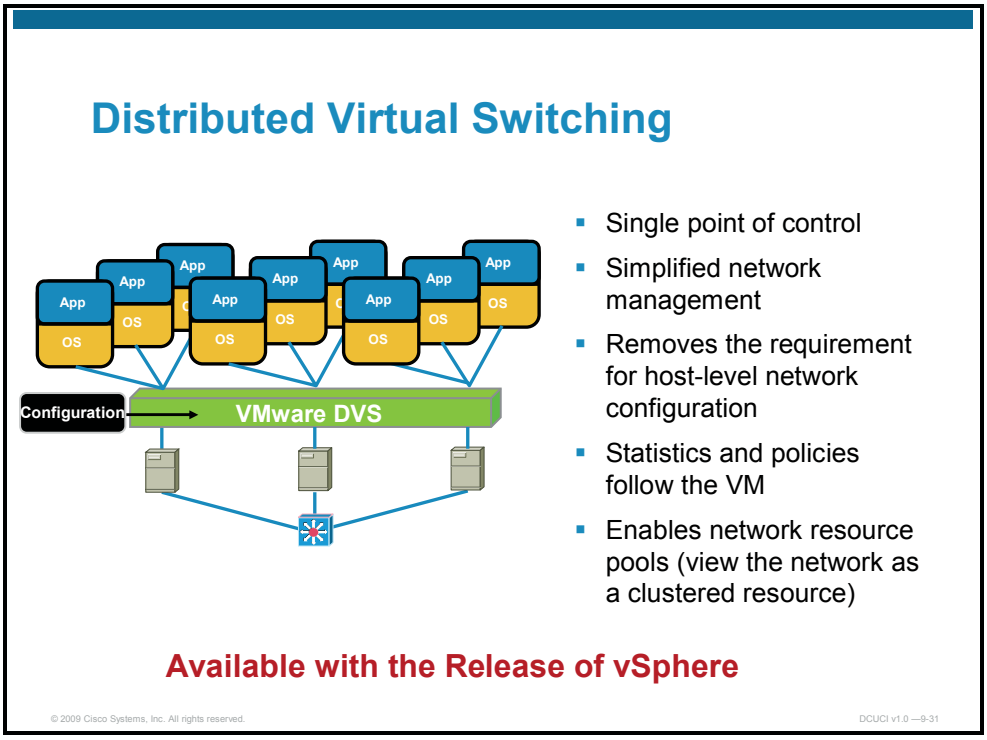

The figure describes the unified management advantages of VMware DVS.

https://t.me/learningnets

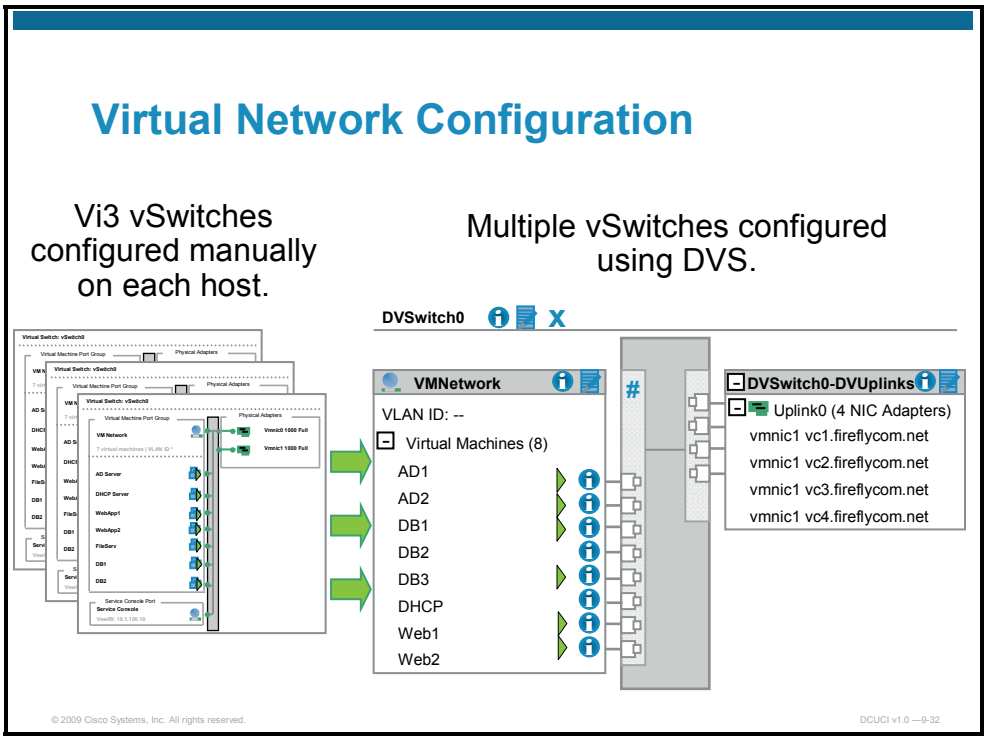

The use of VMware DVS allows for a single point of configuration for clustered networking. This approach allows each VMware ESX Host to be treated as a cluster resource rather than an individually managed entity.

https://t.me/learningnets

DVS Enhancements

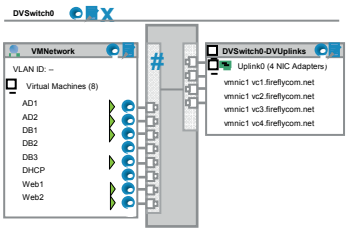

The VMware DVS offers several enhancements to VMware switching:

- Port state migration (statistics and port state follows VM)
- Rx Rate Limiting
- PVLANs

© 2009 Cisco Systems, Inc. All rights reserved.

VMware's DVS adds additional functionality and simplified management to the VMware network. DVS adds the ability to utilize PVLANs, perform inbound rate limiting, and tracks VM port state with migrations. Additionally, the DVS is a single point of network management for VMware networks. The VMware DVS is a requirement for the Cisco Nexus 1000V.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

vSwitch and DVS

VMware DVS and vSwitch are not mutually exclusive:

- Physical NIC ports are assigned to either a vSwitch or DVS
- Separate ports can be assigned to a vSwitch and DVS on the same VMware ESX Host

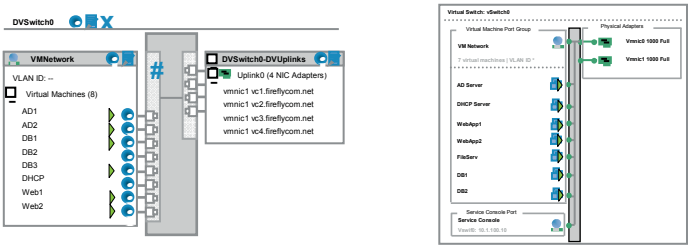

© 2009 Cisco Systems, Inc. All rights reserved.

The VMware DVS and vSwitch are not mutually exclusive. Both devices can be run in tandem on the same VMware ESX Host. An example of when this type of configuration is necessary would be when using the Cisco Nexus 1000V VSM running on a host that it is controlling. In this scenario, the VSM runs on a vSwitch configured for VSM connectivity while it controls a DVS running a Virtual Ethernet Module (VEM) on the same host.

https://t.me/learningnets

Summary

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

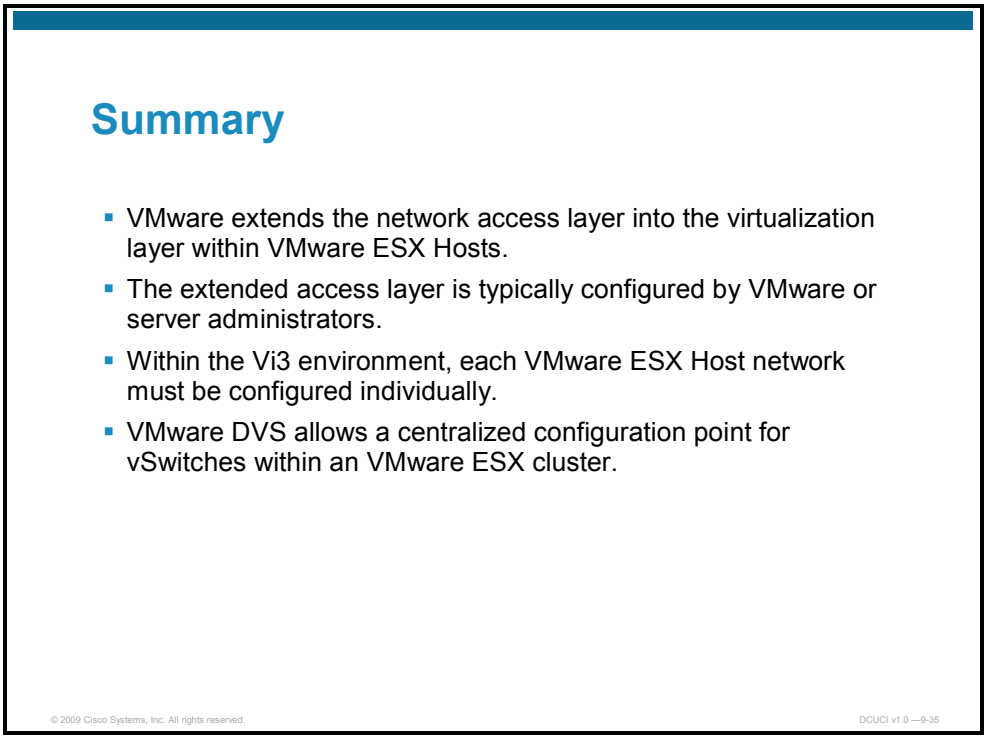

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

The Cisco Nexus 1000V Architecture

Overview

This lesson describes the Cisco Nexus 1000V network architecture.

Objectives

Upon completing this lesson, you will be able to describe the Cisco Nexus 1000V network architecture. This includes being able to meet these objectives:

- Describe the Cisco Nexus 1000V
- Describe the architecture of the Cisco Nexus 1000V
- Describe the Cisco Nexus 1000V installation and operation

https://t.me/learningnets

Cisco Nexus 1000V

This topic explains the Cisco Nexus 1000V.

Cisco Nexus 1000V Series Switch

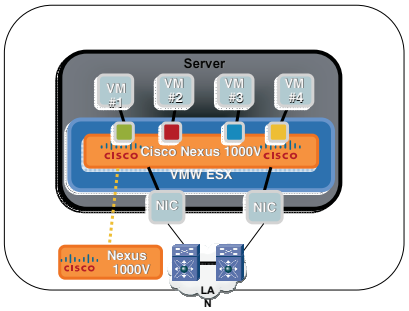

- Cisco Nexus 1000V provides the following per VM:
  - VLAN, PVLAN settings
  - ACL, port security, ACL Redirect
  - NetFlow collection
  - QoS marking (CoS/DSCP)
  - Switched Port Analyzer (SPAN)
  - Mobility of network policy and security to the virtual network
- Maintains policies as VMs migrate using vMotion

© 2009 Cisco Systems, Inc. All rights reserved.

The Cisco Nexus 1000V switch is a software solution that is used in place of VMware vSwitches to provide advanced functionality. The Cisco Nexus 1000V allows all of the same physical network best practices to be applied at the virtual machine (VM) level. This is not traditionally possible with VMware networks.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

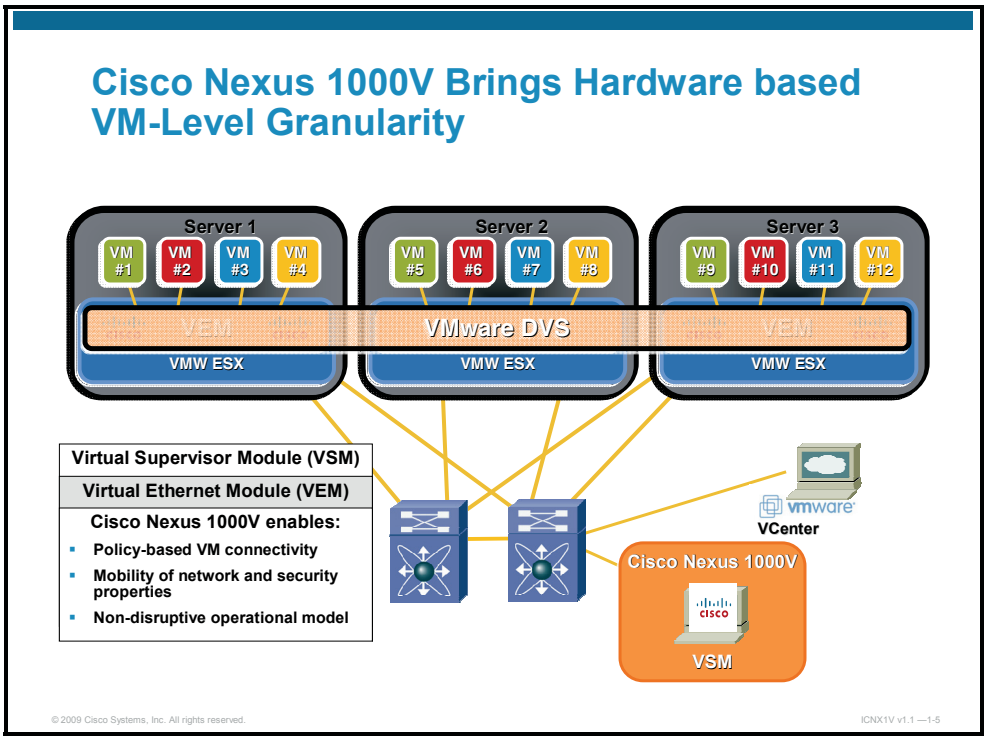

Virtual Supervisor Module (VSM):

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

Virtual Ethernet Module (VEM):

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

Note “DVS” stands for Distributed Virtual Switch.

https://t.me/learningnets

Centralized Control Plane Architecture

VMware ESX server cluster is modeled as a single Cisco Nexus switch. One NX-OS control plane manages multiple data planes in a cluster:

- One data plane per VMware ESX server
- Maximum 32 VMware ESX servers per cluster

NX-OS VSM can run as:

- A VM on one of the servers in the cluster (CPVA)
- A physical appliance connected to an upstream switch (CPPA)

© 2009 Cisco Systems, Inc. All rights reserved.

The distributed architecture of the Cisco Nexus 1000V deployment allows a single Cisco Nexus Operating System (Cisco NX-OS) based supervisor module to manage the switching capabilities of up to 32 VMware ESX servers. Each VMware ESX server Virtual Ethernet Module (VEM) acts as a “remote line card” of the Virtual Supervisor Module (VSM) and is configured as such. The VSM itself can run as a virtual machine known as a Control Plane Virtual Appliance (CPVA), or as a physical appliance known as a Control Plane Physical Appliance (CPPA).

https://t.me/learningnets

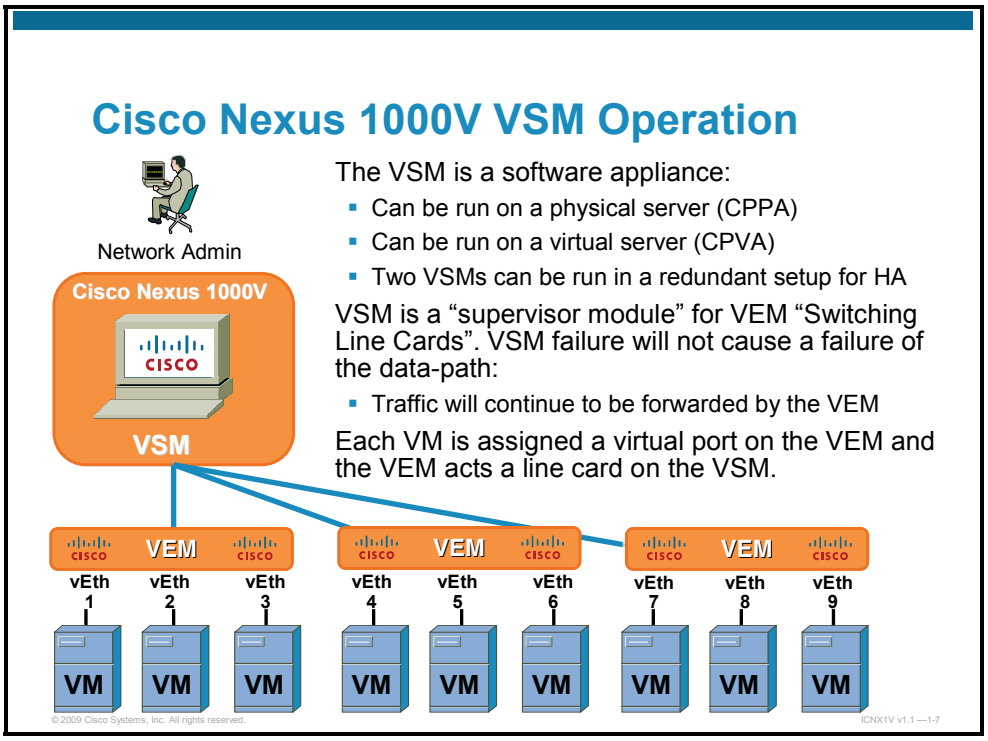

The Cisco Nexus 1000V VSM acts as a supervisor module for remote VEMs. The VEM is a software replacement for the VMware vSwitch on a VMware ESX Host. All traffic forwarding decisions are made by the VEM. All configuration and supervisor functions are handled by the VSM. A Cisco Nexus 1000V cluster can be thought of as a distributed software director switch. It operates in a similar fashion to a Catalyst 6500. Control plane traffic and configuration are handled by a supervisor module and switching is handled by line cards.

In Cisco Nexus 1000V environments, two software components make up a virtual switching chassis. The VSM is the central point of management for multiple VEMs.

Each VEM acts as a separate switching line card with no concept of a fabric between VEMs. This means there are no port channels or connections between VEMs and they do not rely on one another for operation. Switching decisions and frame forwarding all happen on the VEM and are not reliant on the VSM.

The VSM is responsible for configuration of the VEM modules. Using Console, Telnet, or SSH, an administrator makes all configuration changes on the VSM. When a change is made on the VSM, the configuration is passed to vCenter and the changes are made on the DVS. DVS changes are passed down to the corresponding VEM. Each VEM will act as a module on the VSM and the VSM/VEM(s) appear as a single switch to Cisco Discovery Protocol (CDP) neighbors. The VSM does not reside in the data path and therefore can not directly receive or respond to CDP messages. CDP and other network management packets are transferred between the VEM to the VSM on one of three required VLANs known as the Packet VLAN.

https://t.me/learningnets

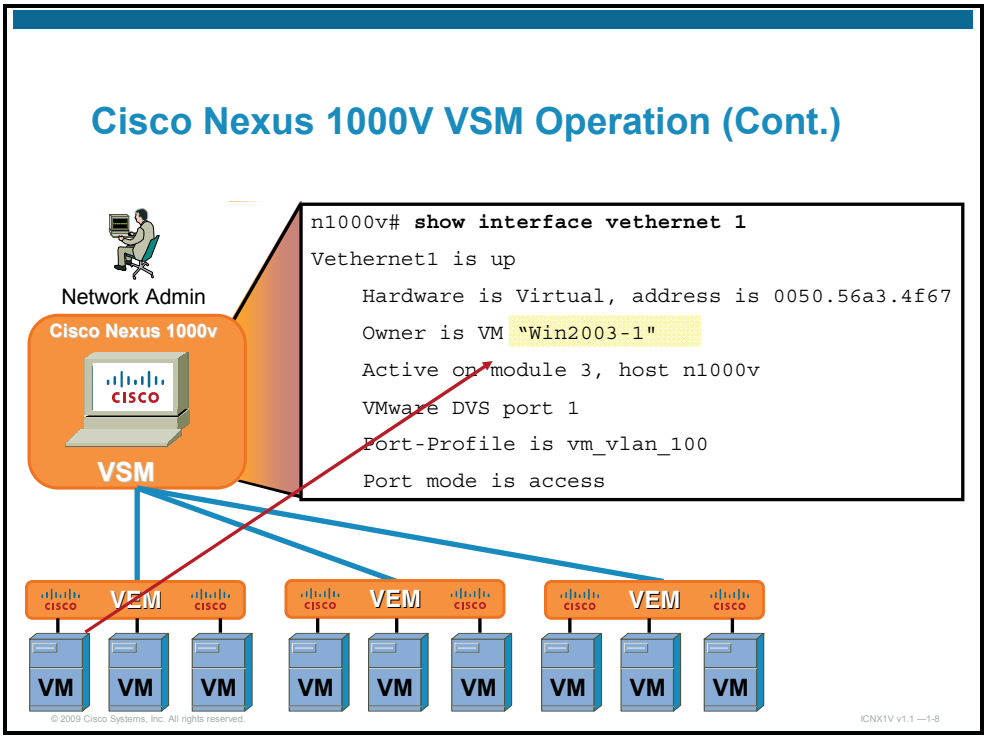

Each VM is assigned a virtual Ethernet port which is configured using the VSM. The figure shows sample output from the **show interface** command on the VSM. The "owner" output shown in the figure is acquired automatically from vCenter. This allows an administrator to quickly identify virtual Ethernet ports.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

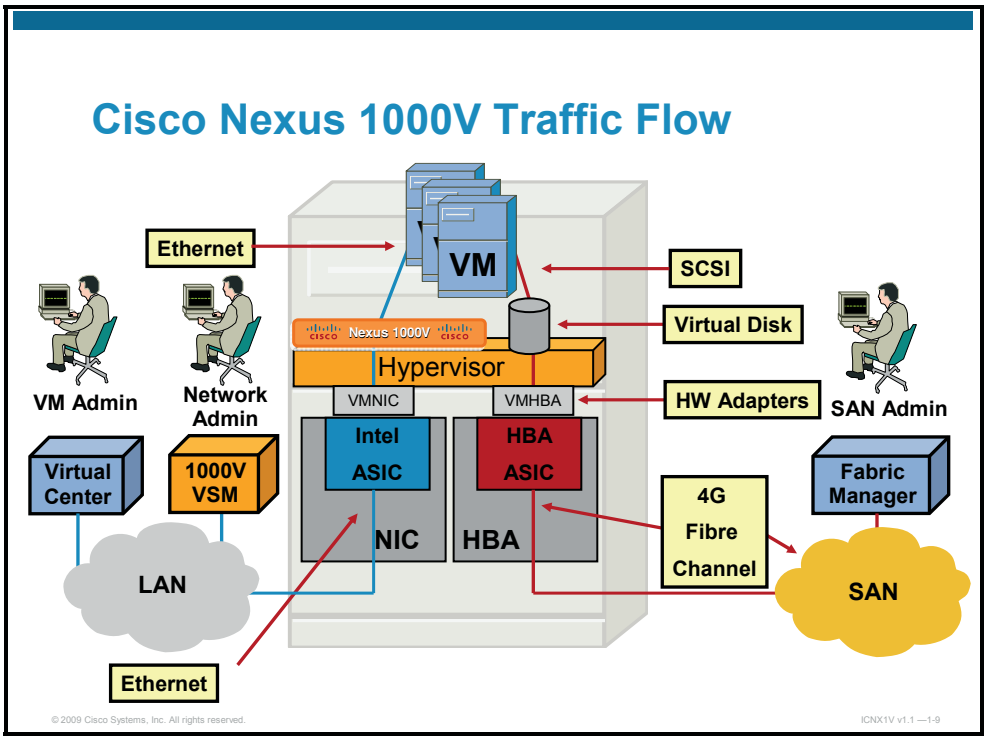

The figure shows the traffic flow from a VM to the LAN and SAN networks using the Cisco Nexus 1000V and Fibre Channel over Ethernet (FCoE).

Step 1

- Storage traffic originates from the VM destined for a virtual SCSI disk using a virtual Small Computer Systems Interface (SCSI) Host Bus Adapter (HBA).
- Ethernet traffic leaves the VM through a Virtual Ethernet network interface card (NIC) destined for the LAN.

Step 2

- The Ethernet frame is received by the Cisco Nexus 1000V VEM. Security checks are performed, VLAN and CoS tagging occurs and a forwarding decision is made.
- Finally, the frame is passed to the hypervisor.

Step 3

- The SCSI traffic is received at the Hypervisor Layer encapsulated in Fibre Channel frames and transmitted to the HBA ASIC onboard the Converged Network Adapter (CNA).
- Ethernet frames are received by the hypervisor and passed to the Ethernet ASIC on the CNA.

Step 4

- The HBA ASIC on the CNA passes the Fiber Channel frame to the Menlo ASIC where it is encapsulated in an FCoE type Ethernet frame and sent to the Cisco Nexus 5000.
- The Ethernet frame is passed through the Menlo ASIC and transmitted to the Cisco Nexus 5000.

https://t.me/learningnets

Step 5

- The Cisco Nexus 5000 sends the FCoE frame to an FCoE Logical End Point (LEP) for de-encapsulation. The LEP strips the FCoE type Ethernet frame and the Fibre Channel frame is forwarded to the SAN.
- The Ethernet frame is handled as any standard Ethernet frame and forwarded to the LAN.

Administration of these three device is handled using the current tools:

- SAN management is performed using the Cisco Fabric Manager (FM) suite.
- LAN management is performed using a NX-OS CLI and the Cisco Datacenter Network Manager (DCNM).
- VMware administration is performed using VMware's vCenter management server.

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

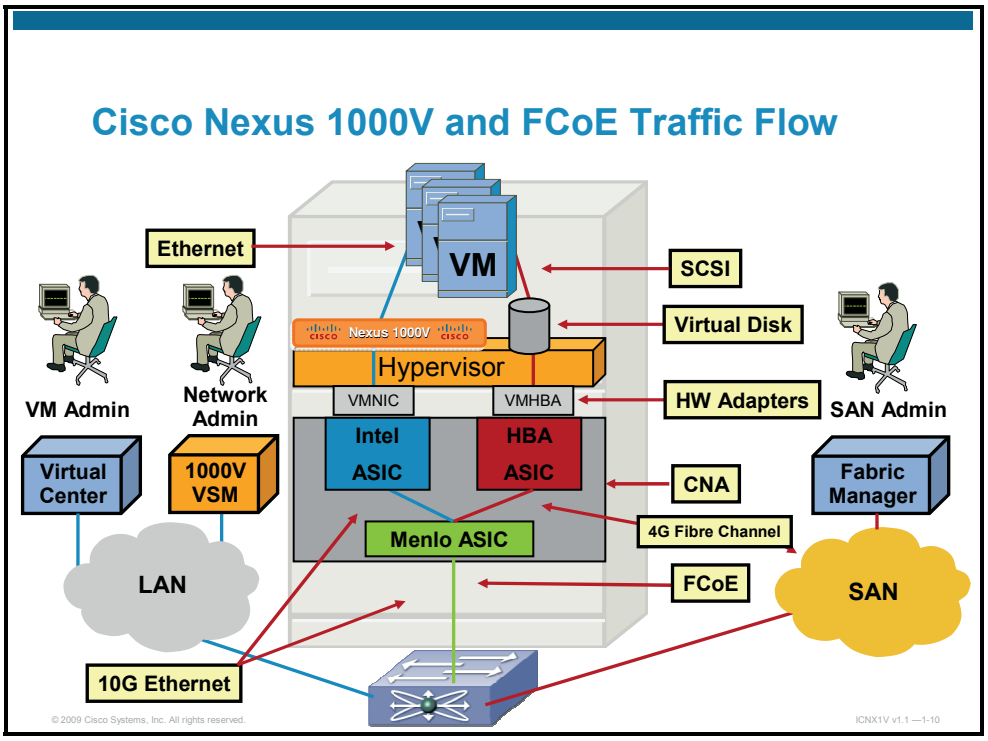

The figure shows the traffic flow from a VM to the LAN and SAN networks using the Cisco Nexus 1000V and FCoE.

Step 1

- Storage traffic originates from the VM destined for a virtual SCSI disk using a virtual SCSI HBA.
- Ethernet traffic leaves the VM through a Virtual Ethernet NIC destined for the LAN.

Step 2

- The Ethernet frame is received by the Cisco Nexus 1000V VEM. Security checks are performed, VLAN and CoS tagging occurs and a forwarding decision is made.
- Finally, the frame is passed to the hypervisor.

Step 3

- The SCSI traffic is received at the Hypervisor Layer encapsulated in Fibre Channel frames and transmitted to the HBA ASIC onboard the CNA.
- Ethernet frames are received by the hypervisor and passed to the Ethernet ASIC on the CNA.

Step 4

- The HBA ASIC on the CNA passes the Fibre Channel frame to the Menlo ASIC where it is encapsulated in an FCoE type Ethernet frame and sent to the Cisco Nexus 5000.
- The Ethernet frame is passed through the Menlo ASIC and transmitted to the Cisco Nexus 5000.

Step 5

- The Cisco Nexus 5000 sends the FCoE frame an FCoE LEP for de-encapsulation. The LEP strips the FCoE type Ethernet frame and the Fibre Channel frame is forwarded to the SAN.

https://t.me/learningnets

- The Ethernet frame is handled as any standard Ethernet frame and forwarded to the LAN.

Administration of these three device is handled using the current tools:

- SAN management is performed using the Cisco FM suite.
- LAN management is performed using a NX-OS CLI and Cisco DCNM.
- VMware administration is performed using VMware's vCenter management server.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Architectural Overview

This topic describes the architecture of the Cisco Nexus 1000V.

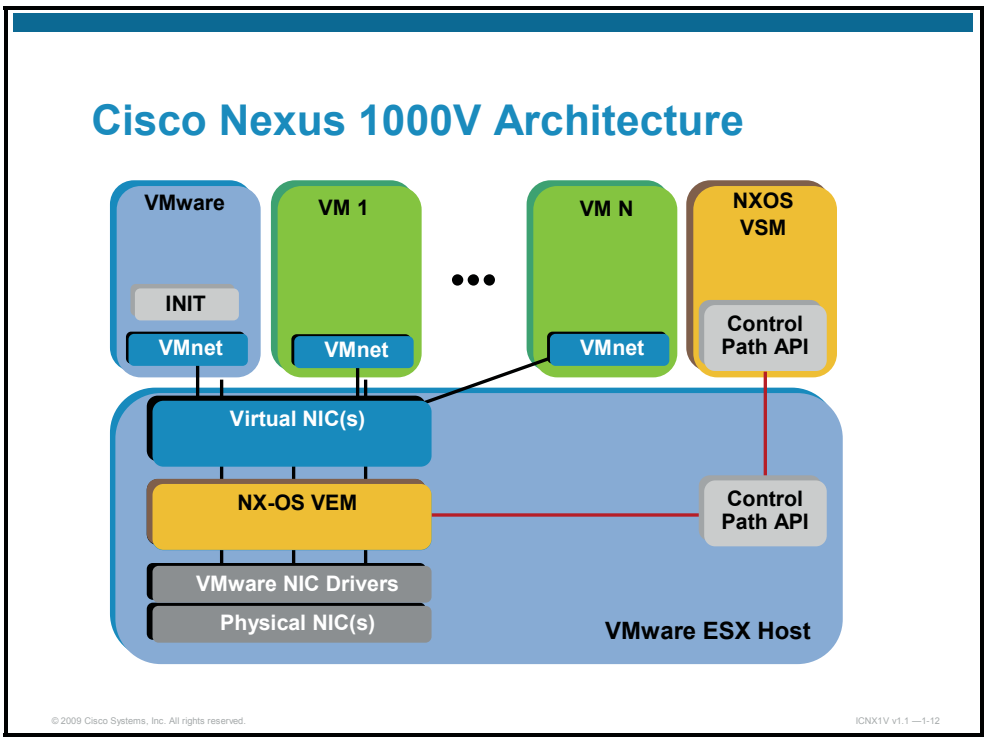

The figure depicts the basic architecture of the Cisco Nexus 1000V. The Cisco Nexus 1000V uses a control path application programming interface (API) in order to communicate with the data plane. This simulates out-of-band management and does not interfere with the data plane. In the figure, the NX-OS VEM replaces the vSwitch for assigned vmNICs and virtual network interface cards (vNICs). The VEM is controlled by the VSM on an out-of-band channel (control VLAN.) VMware's vCenter provides the API used by the NX-OS VSM for control of the VEM. The VSM resides logically outside of the host, although it can be a virtual machine residing on the host.

In 1000V deployments VMware provides the vNIC and drivers while the Cisco Nexus 1000V provides the switching and management of switching.

Note	A NIC in VMware is represented by an interface called a vmNIC. The vmNIC number is allocated during VMware installation.

https://t.me/learningnets

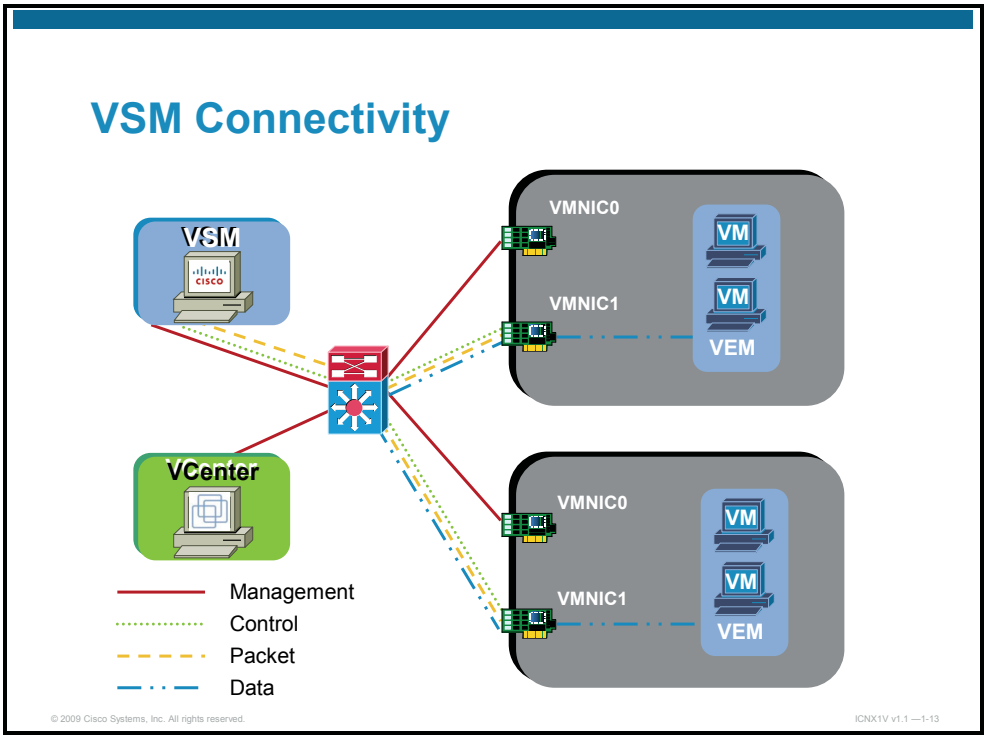

The Cisco Nexus 1000V VLANs consist of a management, control, packet, and one or more data networks. The data networks carry VM packet traffic (server data). Each VMware ESX Host, the VSM, and the vCenter server must all reside in the same management network and be part of the same Layer 2 domain. Control and Packet VLANs from the VSM must be accessible by uplink profiles on each VEM. Data traffic from the VM is not sent to the VSM and the VSM does not require access to the data VLANs. All VSM management is out of band and switching decisions do not rely on the VSM.

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

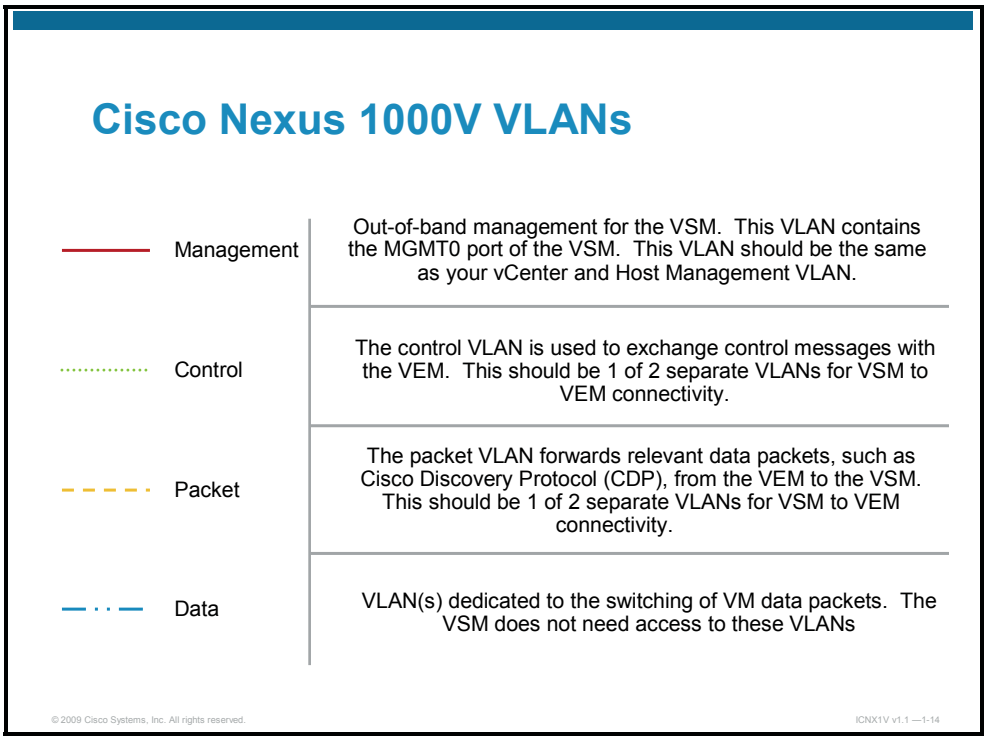

This figure shows the use of Cisco Nexus 1000V networks.

https://t.me/learningnets

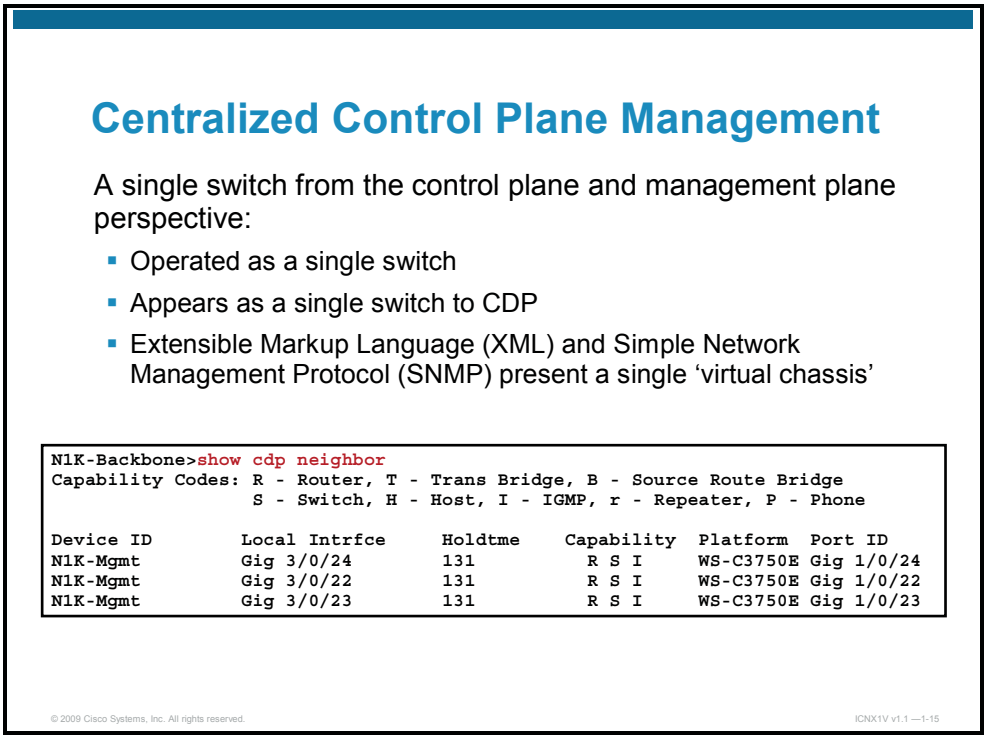

All upstream devices see the Cisco Nexus 1000V as a single switch, regardless of the number of Cisco Nexus 1000V VEMs it is supporting. Management of the VEMs is handled through the VSM, as a single switch.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Distributed Data Plane

Each data plane (DP) behaves like an independent switch:

- No address learning/synchronization across DPs
- No concept of fabric between the DPs
- No concept of forwarding from an ingress line card to an egress line card
- No EtherChannel across DPs
- Internal switching (no need to access the control plane)

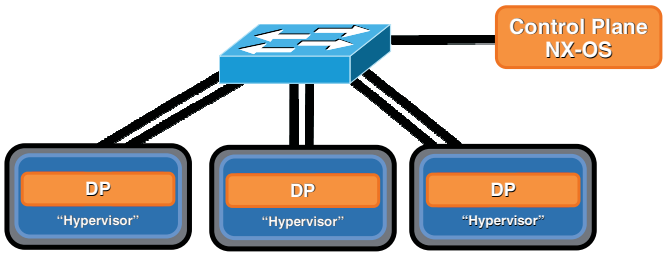

© 2009 Cisco Systems, Inc. All rights reserved.

Each control plane (CP) maintains it's own forwarding table and is able to switch independently of the CP. There is no concept of a centralized fabric that frames must transverse, as each DP is able to switch internally or forward frames to the physical network.

https://t.me/learningnets

Physical Ports

Represented as eth ports on line cards:

- 1 line card per VMware ESX server (MODULE_TYPE_NIC_LC)
- 1 Eth port instance on line card per vmNIC on server

```
Nexus1000v# show interface brief

Ethernet VLAN    Type Mode   Status  Reason   Speed   Port   Interface   Ch#
--------------------------------------------------------------------------
Eth3/1    1       eth  trunk  up      none    auto             (D)      --
Eth3/2    1       eth  trunk  up      none    auto             (D)      --
```

NIC numbering within VMware ESX begins with 0, while port numbering within NX-OS begins with 1. This means that Eth1/1 will map to NIC0 on the host.

© 2009 Cisco Systems, Inc. All rights reserved.

Each VMware ESX server is represented as a line card on the VSM. Physical NIC (vmNIC) ports are shown as Ethernet ports on that line card. In the example, you see the vmNIC ports on the VMware ESX server assigned to slot 3. Slots 1 and 2 are reserved for supervisor modules, therefore the first switching line card will always be slot 3. Because VMware begins its NIC numbering at 0 and Cisco begins port numbering at 1, Eth3/1 will represent the 1st NIC (NIC0) on the VMware ESX Host; Eth3/2 would represent the 2 NIC (NIC1).

https://t.me/learningnets

Virtual Ethernet Ports

vEth ports are a new interface type that map to a vNIC.

Interface status:

- Can be administratively enabled/disabled
- Operationally down until VM is activated on a server

Interface info:

- Stats preserved through migration from one server to another
- **show interface** indicates which server the interface is currently active on

```
NX-OS# show interface veth10

Vethernet10 is up
    Hardware is Virtual, address is 0050.5681.2bc4
    Owner is VM "WebServerVM"
    Host is ESX1
```

© 2009 Cisco Systems, Inc. All rights reserved.

A new type of port has been created to identify the individual virtual machines. This port type is a virtual Ethernet (vEth) port. This port type represents a vNIC regardless of what physical server it is on. vEth ports are not tied to a specific server or vmNIC, but instead represent the vNIC of a virtual server. The vEth port for a VM will remain the same even if the VM migrates. This means that if a VM obtains “vEthernet 1” on host 1 and is moved to host 2, it will retain vEthernet 1 as it’s port. This allows for network configuration and policy mobility.

https://t.me/learningnets

Logical Virtual Ethernet Ports

- Represents the local port on the Cisco Nexus1000V data plane (VEM) within a VMware ESX Host
- A vEth port is automatically bound to an LvEth port when the VM is powered on
- LvEth ports are not user configurable and are simply an internal system mapping
- vEth ports will be remapped to an LvEth port by the system when a VM is migrated between servers
- All configuration is performed on the vEth port and is persistent with the VM through migrations

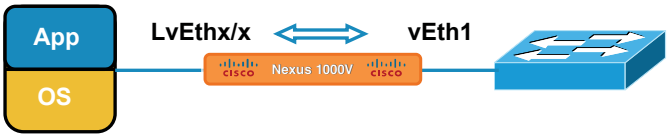

© 2009 Cisco Systems, Inc. All rights reserved.

Logical vEth (LvEth) ports are used internally to tie a vEth port to a specific slot. Because each slot represents a data plane of an individual VMware ESX Host and VMs can migrate between hosts, there must be a dynamic mapping between the vEth and the host. Administrators will configure and manage vEth ports (which represent individual vNICS). The system will map these vEth ports to a LvEth port, to represent the host that the VM is running on. LvEth mapping occurs when a machine is powered on. When a VM is migrated to another server, the system will automatically modify the mapping to reflect the change. The migration of VMs is monitored by the VSM through vCenter APIs.

https://t.me/learningnets

Cisco Nexus 1000V Installation and Operation

This topic describes the Cisco Nexus 1000V installation and operation.

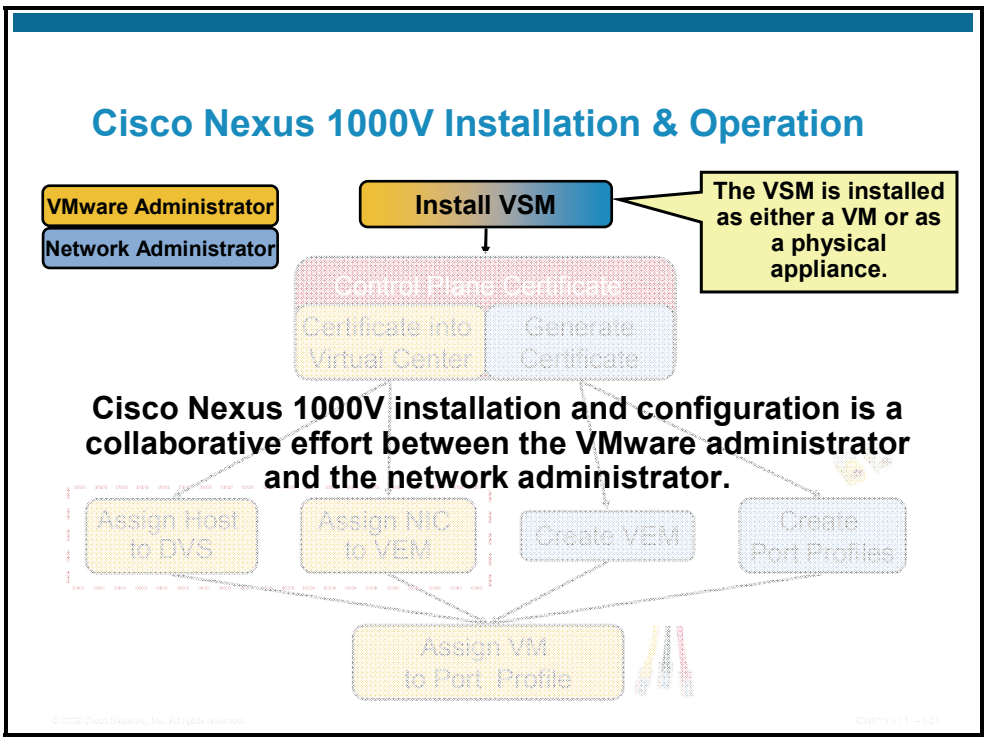

Installing a new Cisco Nexus 1000V or adding a VMware ESX Host to an existing Cisco Nexus 1000V environment is a collaborative process between the VMware administrator and the network administrator. The process follows:

Step 1 VMware administrator installs the VSM as a VM (CPVA), or network administrator installs it as a physical appliance (CPPA).

Step 2 For new VSM only: Network administrator creates a certificate for the VSM; VMware administrator installs the certificate in vCenter.

Step 3 Network administrator creates uplink port profiles.

Step 4 VMware administrator assigns the hosts vmNIC to the uplink profiles.

Step 5 Network administrator creates VM port profiles which are pushed to vCenter for use by the VMware administrator.

Step 6 VMware administrator assigns the VM to an available port profile.

https://t.me/learningnets

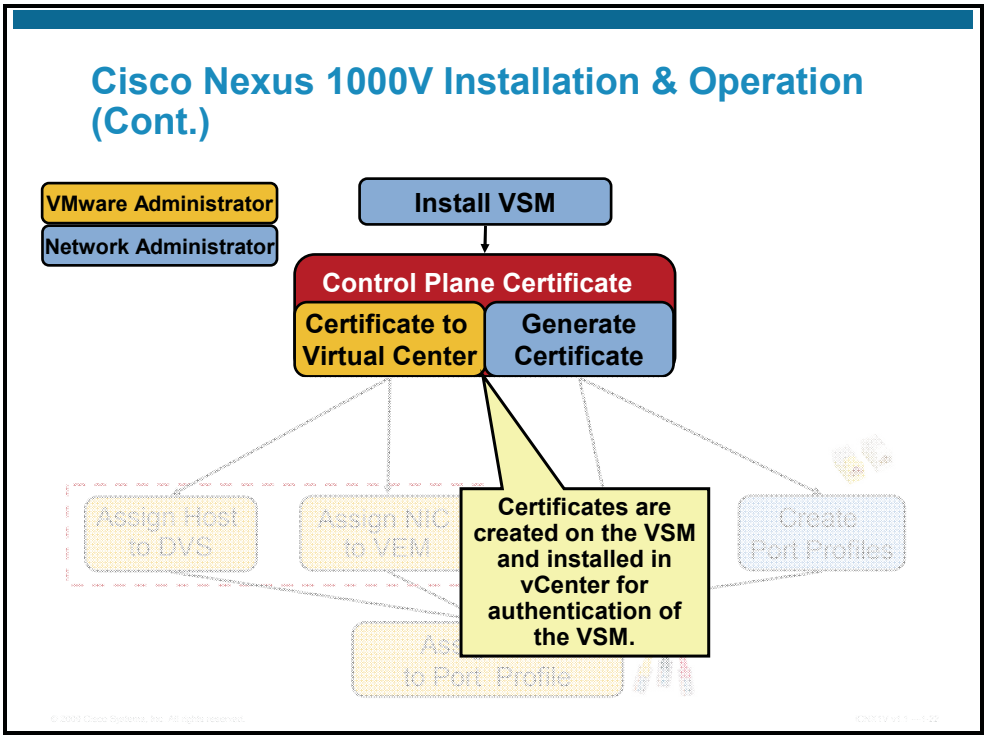

Installing a new Cisco Nexus 1000V or adding a VMware ESX Host to an existing Cisco Nexus 1000V environment is a collaborative process between the VMware administrator and the network administrator. The process follows:

Step 1 VMware administrator installs the VSM as a VM (CPVA), or network administrator installs it as a physical appliance (CPPA).

Step 2 For new VSM only: Network administrator creates a certificate for the VSM; VMware administrator installs the certificate in vCenter.

Step 3 Network administrator creates uplink port profiles.

Step 4 VMware administrator assigns the hosts vmNIC to the uplink profiles.

Step 5 Network administrator creates VM port profiles which are pushed to vCenter for use by the VMware administrator.

Step 6 VMware administrator assigns the VM to an available port profile.

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

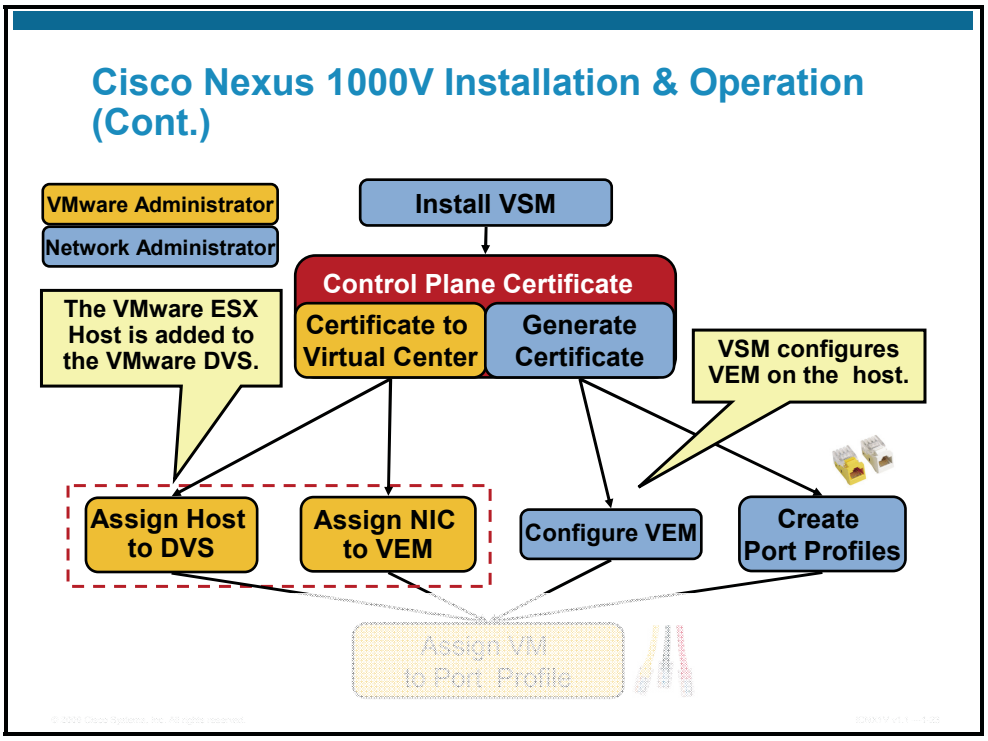

Installing a new Cisco Nexus 1000V or adding a VMware ESX Host to an existing Cisco Nexus 1000V environment is a collaborative process between the VMware administrator and the network administrator. The process follows:

Step 1 VMware administrator installs the VSM as a VM (CPVA), or network administrator installs it as a physical appliance (CPPA).

Step 2 For new VSM only: Network administrator creates a certificate for the VSM; VMware administrator installs the certificate in vCenter.

Step 3 Network administrator creates uplink port profiles.

Step 4 VMware administrator assigns the hosts vmNIC to the uplink profiles.

Step 5 Network administrator creates VM port profiles which are pushed to vCenter for use by the VMware administrator.

Step 6 VMware administrator assigns the VM to an available port profile.

https://t.me/learningnets

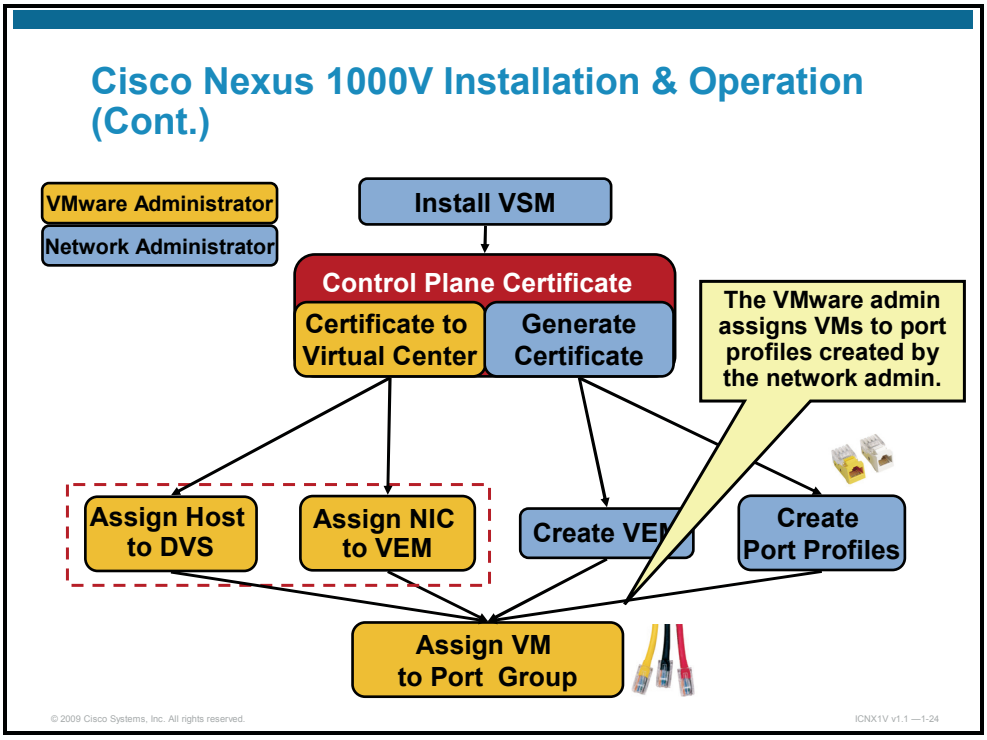

Installing a new Cisco Nexus 1000V or adding a VMware ESX Host to an existing Cisco Nexus 1000V environment is a collaborative process between the VMware administrator and the network administrator. The process follows:

Step 1 VMware administrator installs the VSM as a VM (CPVA), or network administrator installs it as a physical appliance (CPPA).

Step 2 For new VSM only: Network administrator creates a certificate for the VSM; VMware administrator installs the certificate in vCenter.

Step 3 Network administrator creates uplink port profiles.

Step 4 VMware administrator assigns the hosts vmNIC to the uplink profiles.

Step 5 Network administrator creates VM port profiles which are pushed to vCenter for use by the VMware administrator.

Step 6 VMware administrator assigns the VM to an available port group.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Deploying the Cisco Nexus 1000V

1. vCenter and Cisco Nexus 1000V relationship established
2. Network administrator configures Cisco Nexus 1000V uplink profile(s) for VMware ESX Hosts
3. VMware administrator assigns new host vmNICs to uplink port profiles on the Cisco Nexus 1000V DVS

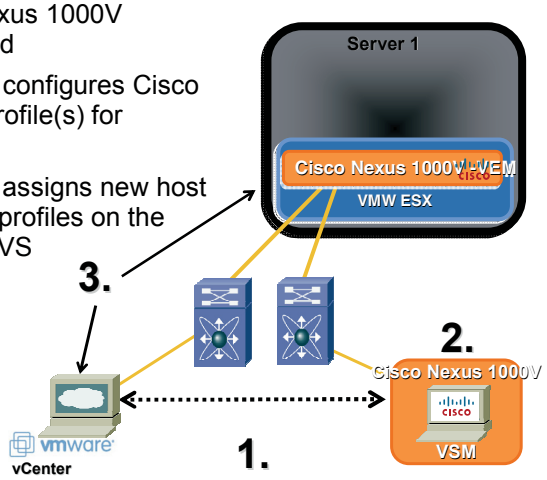

© 2009 Cisco Systems, Inc. All rights reserved.

When deploying the Cisco Nexus 1000V, port profiles are created by the network administrator on the VSM. These port profiles push port groups to vCenter which the VMware administrator uses to assign vNICs and vmNICs to the Cisco Nexus 1000V controlled DVS.

https://t.me/learningnets

Deploying the Cisco Nexus 1000V (Cont.)

3. VMware administrator assigns new host vmNICs to uplink port profiles on the Cisco Nexus 1000V DVS

Repeat step 3 to add another host.

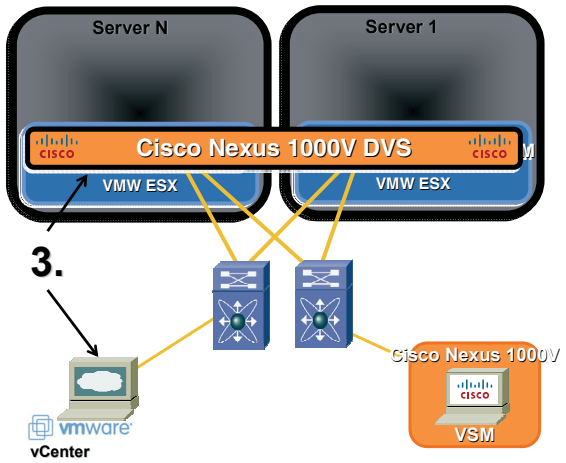

© 2009 Cisco Systems, Inc. All rights reserved.

The VMware administrator assigns each host, or additional host by adding available vmNICs to available uplink port groups.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

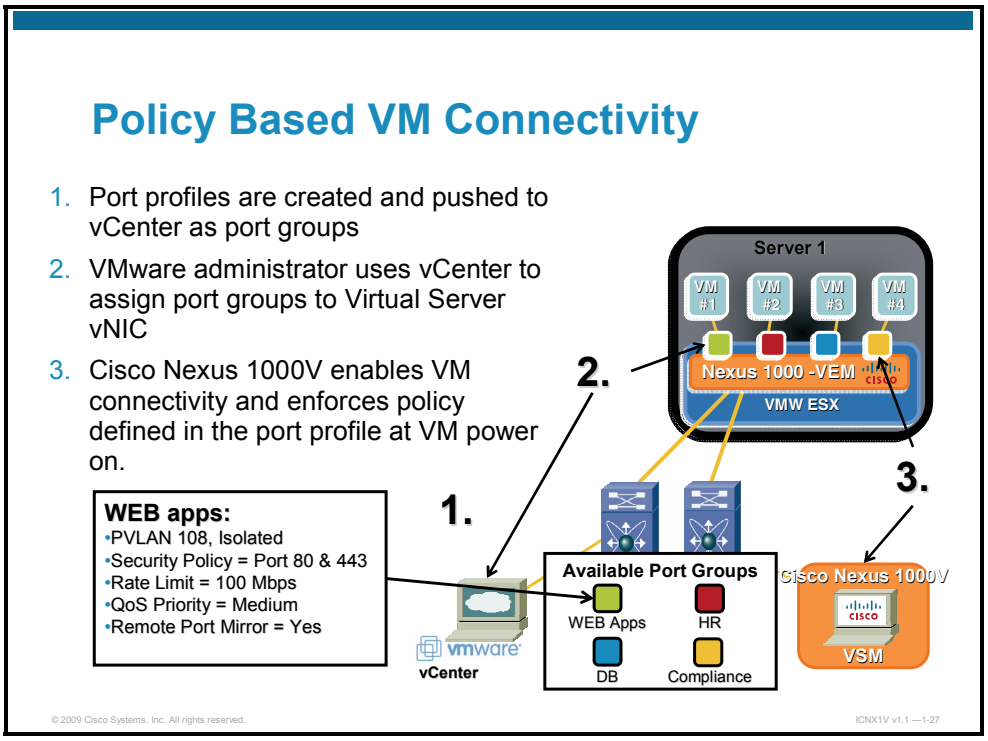

Network administrators configure VM port profiles that contain the policies for specific server types. These port profiles are passed to vCenter as port groups and assigned to individual vNICs by the VMware administrator.

https://t.me/learningnets

Policy Based VM Connectivity (Cont.)

Port profile options:

- VLAN, PVLAN settings
- ACL, Port Security, ACL Redirect
- Cisco Trust Sec (SGT)
- NetFlow Collection
- Rate Limiting
- QoS Marking (COS/DSCP)
- Remote Port Mirror (ERSPAN)

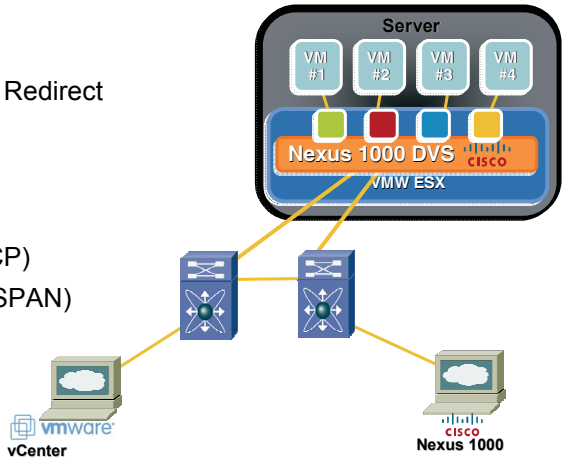

© 2009 Cisco Systems, Inc. All rights reserved.

Various security and traffic flow options can be applied to port profiles:

- VLAN, PVLAN settings
- ACL, Port Security, ACL Redirect
- Cisco Trust Sec (SGT)
- NetFlow Collection
- Rate Limiting
- QoS Marking (COS/DSCP)
- Remote Port Mirror (ERSPAN)

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

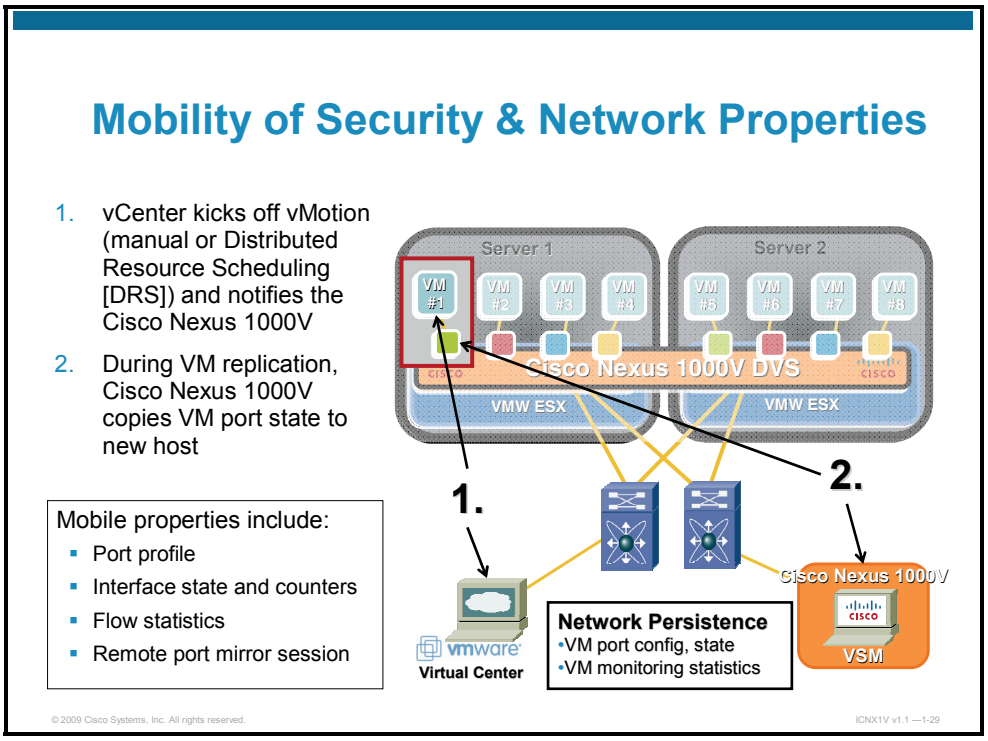

As VMs migrate between hosts through manual or automatic processes, vCenter and the Cisco Nexus 1000V VSM work together to maintain the VMs state and port policies. This ensures network security and consistency in rapidly adapting VMware environment.

https://t.me/learningnets

Mobility of Security & Network Properties (Cont.)

3. Once vMotion completes, the port on the new VMware ESX Host is brought up and the VM's MAC address is announced to the network

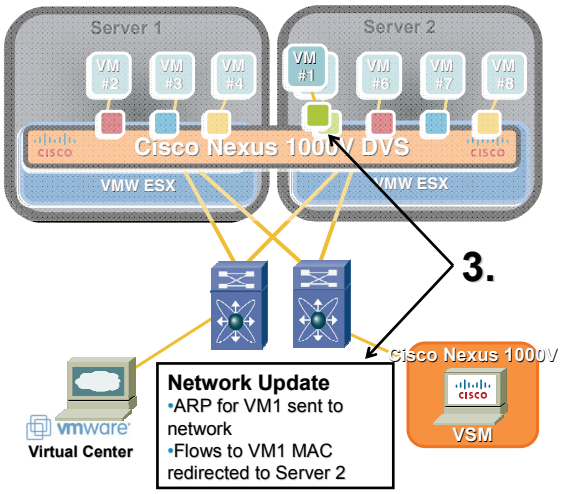

© 2009 Cisco Systems, Inc. All rights reserved.

1. Prevents disparate vSwitch configurations from impacting DRS or live VM migration
2. Enables security setting, network policy and connectivity state to move with vMotion
3. Offers continuous traffic mirroring as VMs move between servers

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

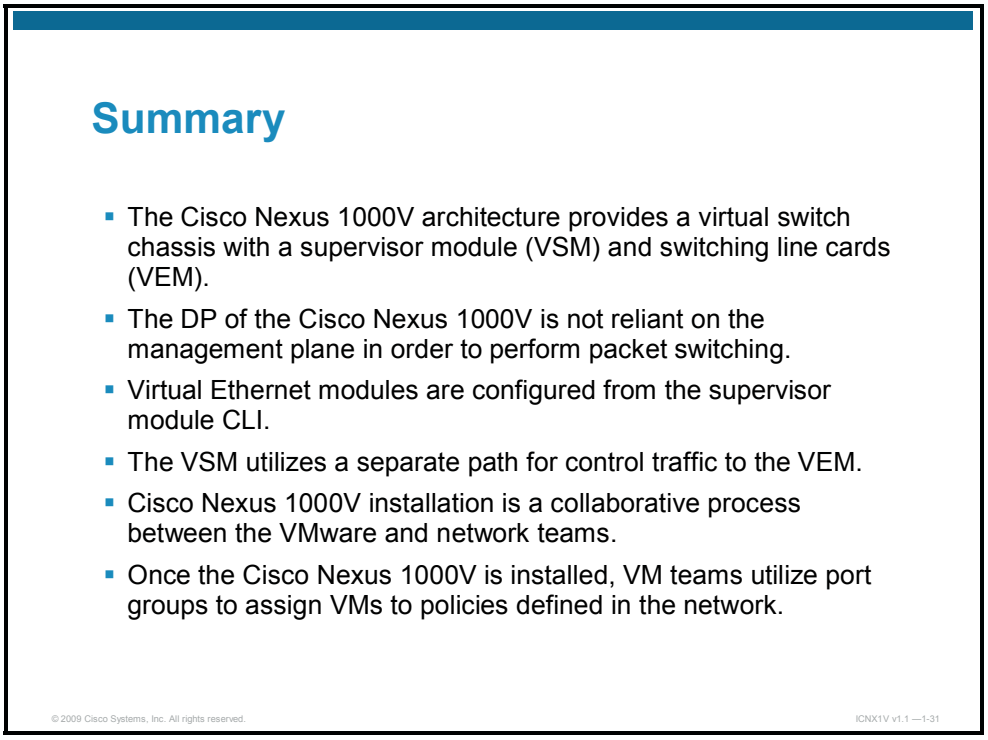

https://t.me/learningnets

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

Lesson 5

Configure VMware Switching

Overview

This lesson explains how to configure VMware switching.

Objectives

Upon completing this lesson, you will be able to explain how to configure VMware switching. This includes being able to meet these objectives:

- Describe vSwitch configuration
- Explain DVS configuration

https://t.me/learningnets

Configure vSwitches

This topic explains vSwitch configuration.

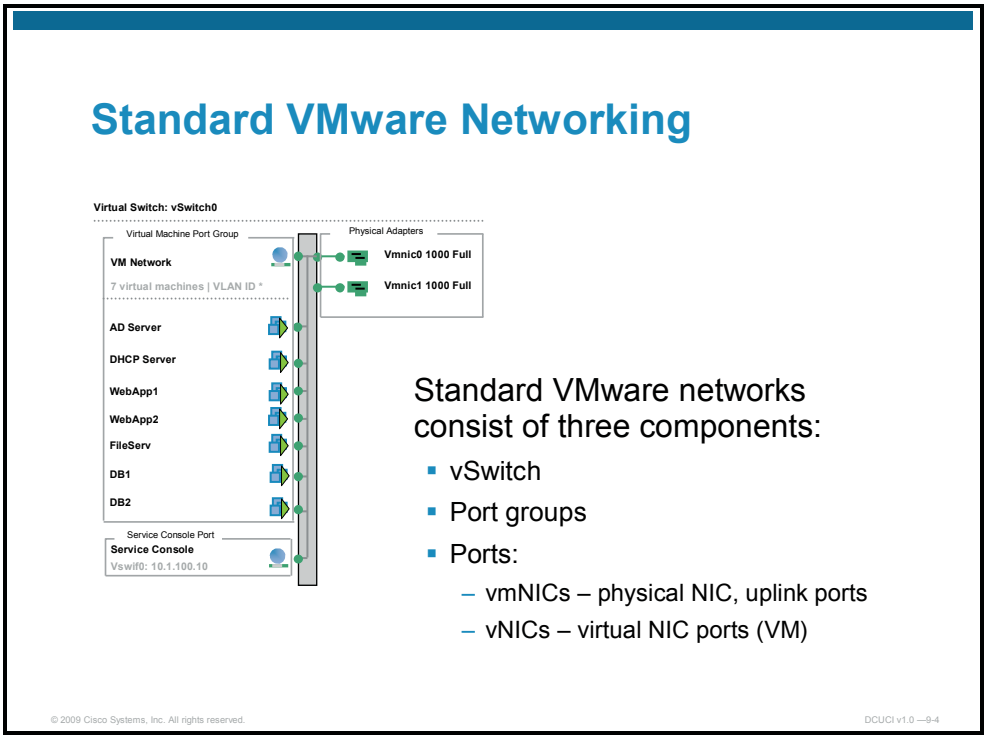

Networking within the VMware host environment is performed using three major components:

- **vSwitch:** Virtual switches
- **Port groups:** Collection of ports sharing similar configuration (such as VLAN)
- **Ports:** Physical network interface cards (NICs) used as uplinks and virtual NICs (vNICs) used for VMware connectivity

Note	A NIC in VMware is represented by an interface called a vmNIC. The vmNIC number is allocated during VMware installation.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Standard VMware Networking (Cont.)

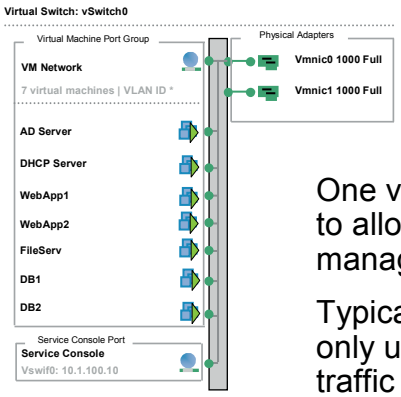

One vSwitch will always be present to allow outbound connectivity for management.

Typically, it is recommended to only use this vSwitch and break up traffic types using port groups and VLANs.

© 2009 Cisco Systems, Inc. All rights reserved.

Typically, it is recommended to only use one vSwitch and segregate traffic types using port groups and VLANs. While some practices may differ, this method simplifies management while providing the same controls as separate vSwitches.

While both port groups and vSwitches segregate traffic in software, typical recommendations are to use multiple port groups on a vSwitch rather than many vSwitches. When a vSwitch is created, a small amount of system resources must be reserved for each port. This is not the same for port groups which are logical break downs of a vSwitch. As the resources have already been reserved by the vSwitch, there is no need to make additional reservations for each port group within a vSwitch.

Multiple port groups can be assigned to a vSwitch and each can belong to a separate VLAN. The port group can provide the VLAN tagging functionality or leave that to the VMs or physical switches. Port groups can either utilize all available NICs assigned to the vSwitch, or they can be assigned specific NICs. Port groups can also be denied outbound connections to create internal networks within the VMware environment.

https://t.me/learningnets

Configuring Existing vSwitches

Existing vSwitches can be modified in order to:

- Add, remove or change port groups:
  - Port group name, VLAN, and so on
- Add physical NIC adapters
- Enable or disable features such as vMotion

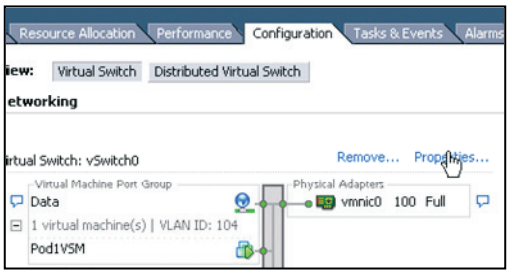

© 2009 Cisco Systems, Inc. All rights reserved.

To open the vSwitch configuration dialogue box, start by selecting the host you wish to configure, then navigate to **Configuration > Networking > Properties** for the switch you want to configure.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Configuring Existing vSwitches (Cont.)

To modify a vSwitch:

- Click **Configuration > Networking**
- Click **Properties** above the vSwitch you want to modify

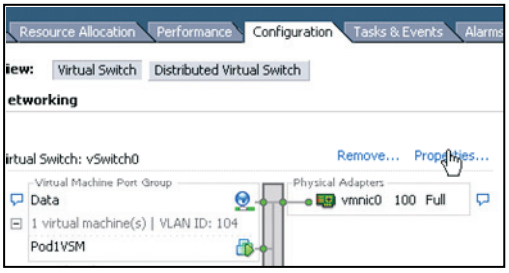

To create a new port group, click **Add.**

To edit an existing port group, highlight that group and click **Edit.**

© 2009 Cisco Systems, Inc. All rights reserved.

When modifying vSwitches, new port groups can be created, existing port groups can be modified, and physical network adapters can be configured, added, or removed.

https://t.me/learningnets

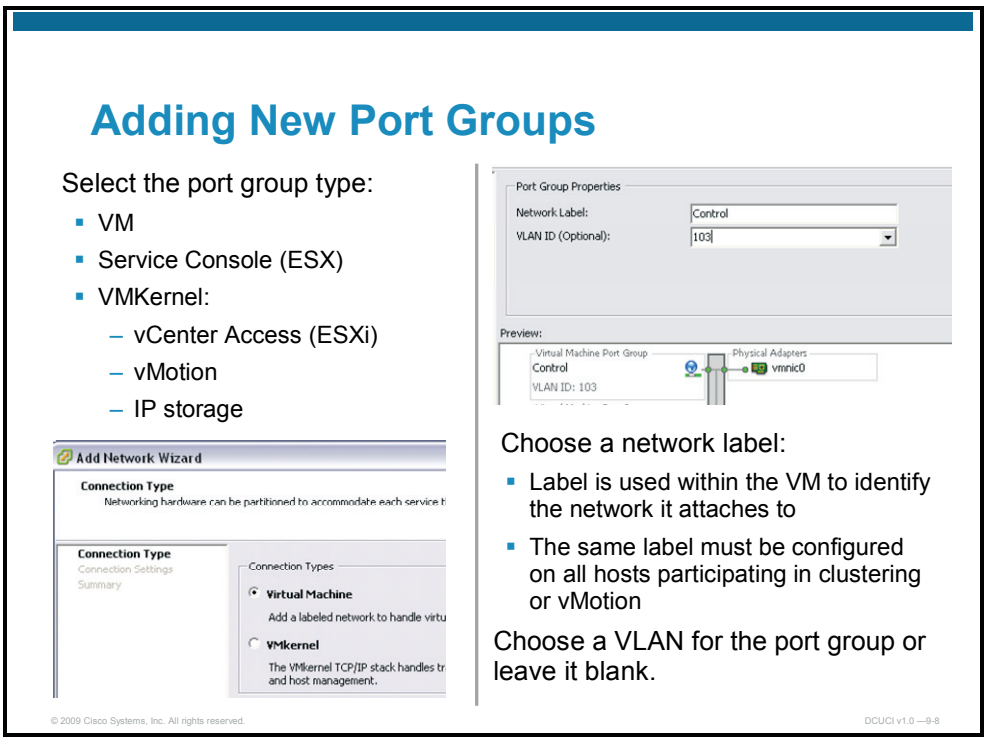

There are two types of additional port groups within VMware ESXi, and three types within standard VMware networks.

ESXi:

- **VMKernel:** Used for the VMware kernel (base VMware operating system) to communicate with the network and network resources such as IP storage.
- **Virtual Machine:** Used for virtual machine network traffic.

ESX:

- **VMkernel:** Used for the VMware kernel (base VMware operating system) to communicate with the network and network resources such as IP storage.
- **Virtual Machine:** Used for VM network traffic.
- **Service Console:** Used for VMware ESX Host management.

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

Attaching a Host to a Port Group

VMs are attached to port groups through the Add hardware form:

- Right click the VM
- Select **Edit Settings**
- Select the vNIC you want to change
- Change the network label to match the port group you wish the VM to attach to. Each vNIC can be attached to only one port group at a time.

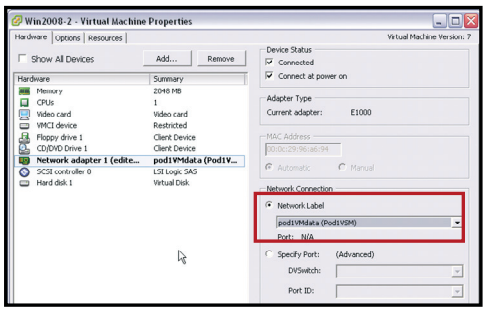

© 2009 Cisco Systems, Inc. All rights reserved.

Once a port group has been created, it becomes a selectable option within VM configuration. VMs are assigned to a port group based on network label.

https://t.me/learningnets

Configuring VMware DVS

This topic explains how to configure the vSphere Distributed Virtual Switch (DVS).

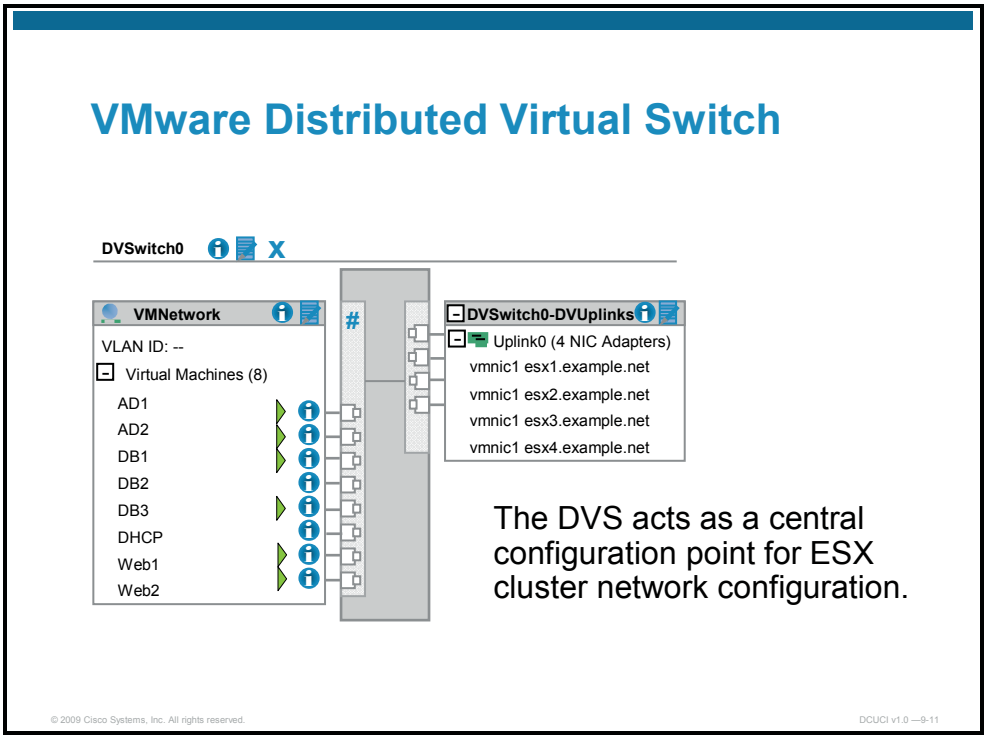

The VMware DVS is created within vCenter. One or more hosts can be added to a DVS. The DVS acts as a central management point for the network resources on attached hosts.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

DVS Creation

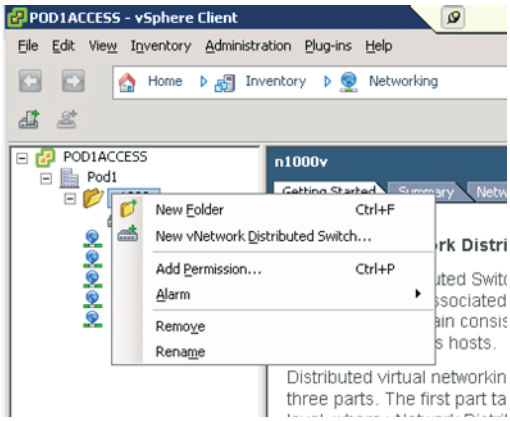

From the Networking view, right click the data center or folder that will receive the new DVS.

© 2009 Cisco Systems, Inc. All rights reserved.

DVS creation is performed through the Networking view within vSphere. DVS switches can be added or removed from this view.

https://t.me/learningnets

DVS Creation (Cont.)

Name the DVS and choose the number of uplink ports for the new DVS.

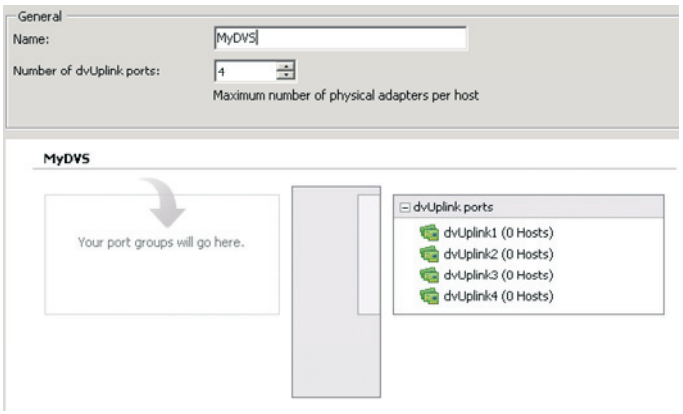

© 2009 Cisco Systems, Inc. All rights reserved.

DVS switches must be given a unique name and assigned a number of uplink ports. The figure shows the first screen of DVS creation.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

DVS Creation (Cont.)

Hosts can be added to the DVS during creation or later.

If adding hosts during creation, select the host and adapters to add.

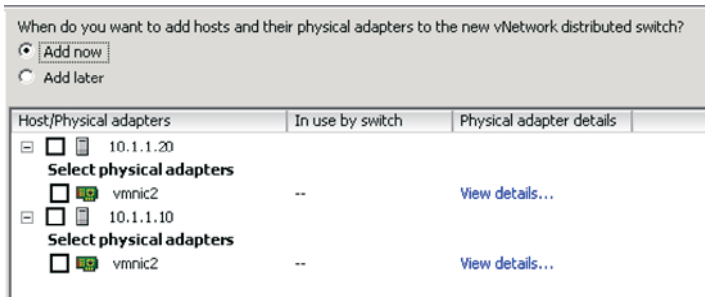

© 2009 Cisco Systems, Inc. All rights reserved.

Hosts are added as members of a DVS. Hosts can be added during or after creation.

https://t.me/learningnets

DVS Creation (Cont.)

Verify your settings and choose whether or not to configure a default port group.

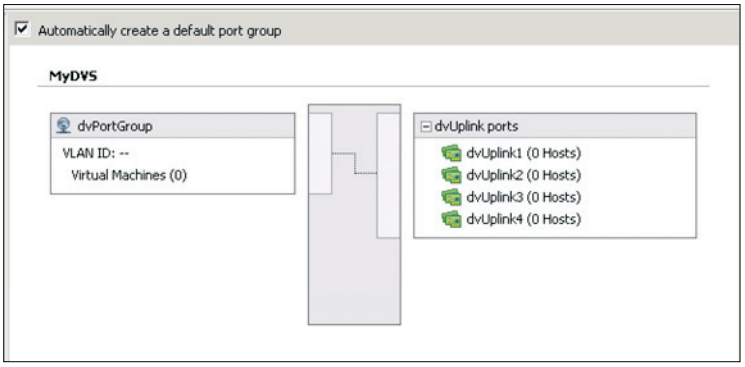

© 2009 Cisco Systems, Inc. All rights reserved.

The final step in DVS configuration is verification. During this final stage, administrators are given a chance to verify the configuration and to choose if a default port group should be assigned.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

DVS Configuration (Cont.)

The DVS contains two major components.

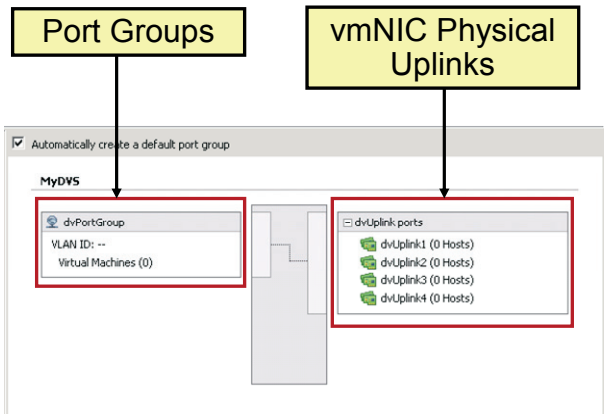

© 2009 Cisco Systems, Inc. All rights reserved.

The DVS consists of port groups for uplinks, and port groups for VMs. Both types can be configured independently.

https://t.me/learningnets

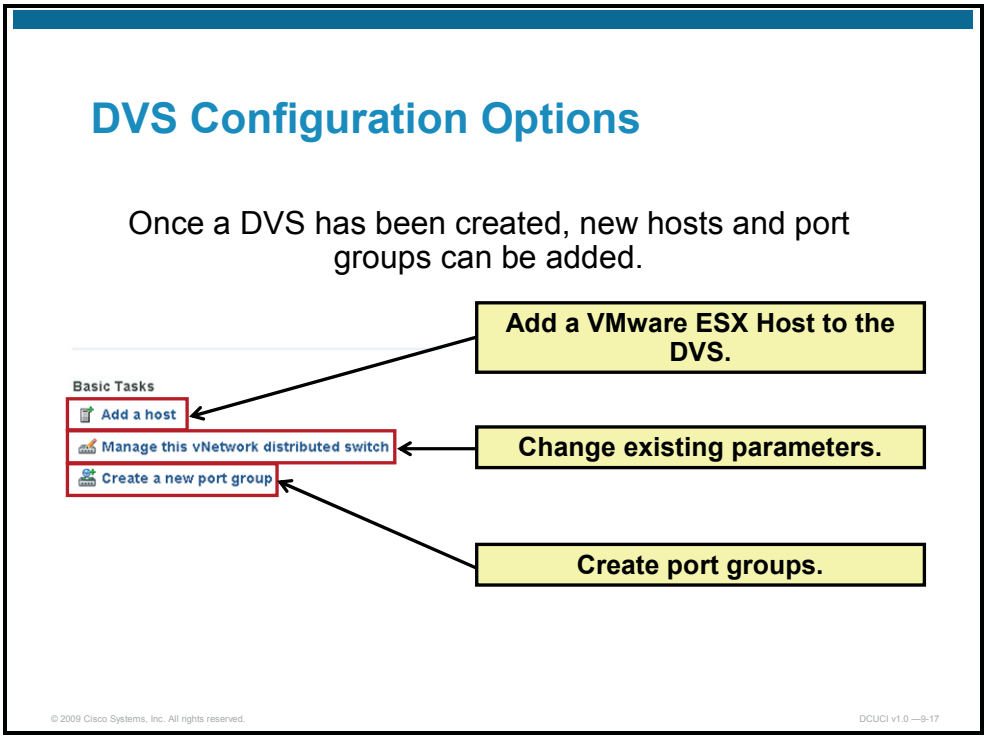

The figure shows options for configuring a DVS.

https://t.me/learningnets

Managing the DVS

From the Properties tab, maximum transmission unit (MTU), Cisco Discovery Protocol (CDP), contact information and general properties such as network label and port count can be changed:

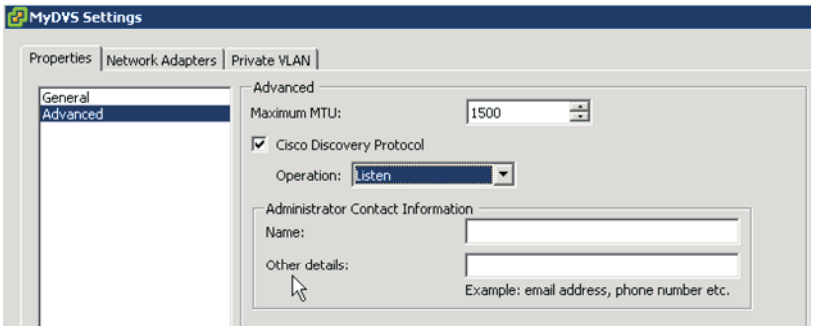

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the advanced general options for the DVS.

https://t.me/learningnets

Network Adapters

The Network Adapters tab shows available hosts and adapters.

These settings can be configured from the individual hosts networking tab.

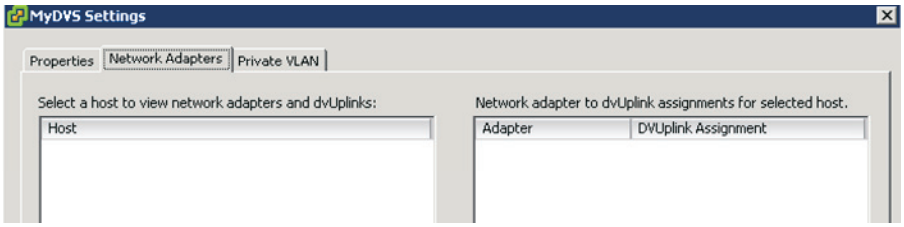

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the Network Adapters view of DVS settings.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Private VLAN settings

Private VLANs can be configured from the Private VLAN tab.

The left pane allows creation of primary VLANs.

The right pane allows configuration of secondary VLANs.

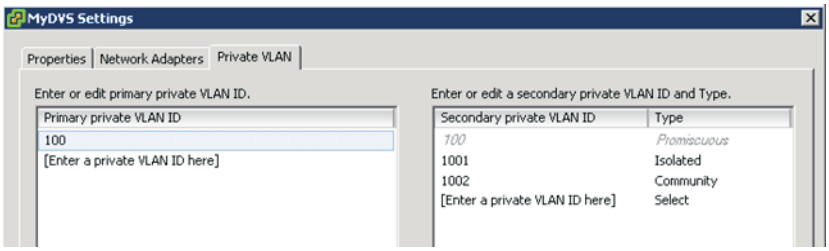

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the Private VLAN configuration options.

https://t.me/learningnets

DVS Port Groups

New port groups can be added; one of three VLAN types can be assigned:

- VLAN:
  - Assign a specific VLAN ID
- VLAN trunking:
  - Trunked VLANS
- Private VLAN:
  - Assign a previously created PVLAN

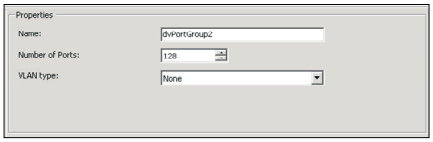

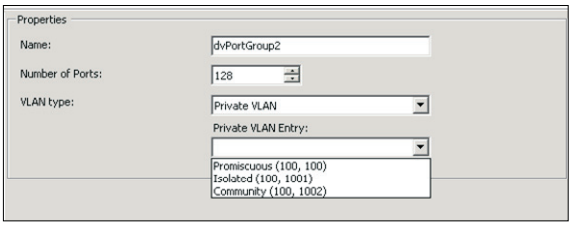

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the VLAN configuration of DVS port groups.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Ports Groups

Uplink port groups have several settings which configure the way outbound traffic is handled.

VM port groups control how VM traffic is handled.

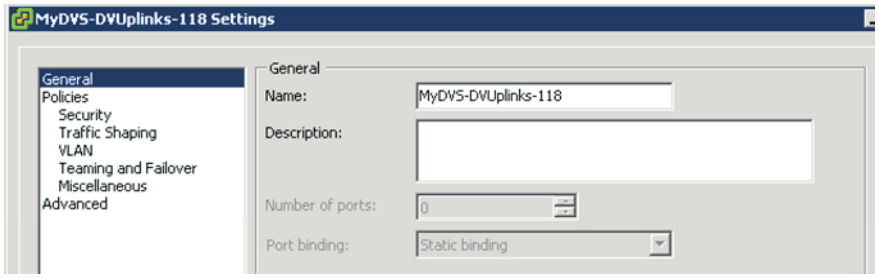

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the uplink port configuration.

https://t.me/learningnets

Security Policies

Security policies for the VMware DVS are similar to vSwitch security policies and are configured for VMware port groups:

- Promiscuous mode
- MAC address changes
- Forged transmits

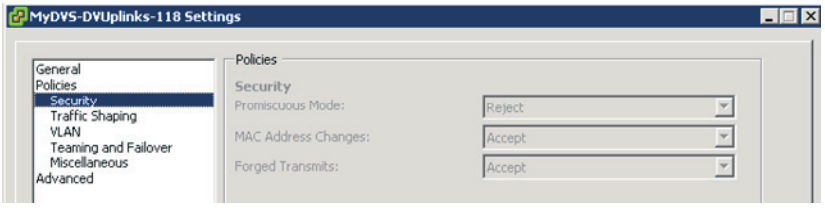

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows security settings for VMware port groups.

https://t.me/learningnets

Traffic Shaping

Traffic shaping features for the DVS consist of egress traffic shaping and now ingress traffic shaping, which was not previously available.

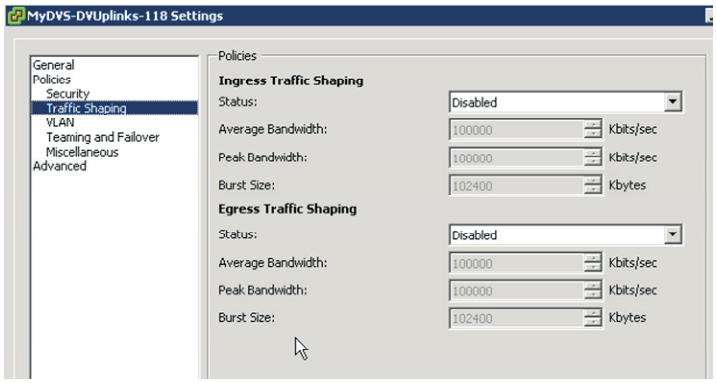

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the traffic shaping options for port groups.

https://t.me/learningnets

VLAN

VLAN trunking defaults to trunking all VLANs. This assignment can be configured within the VLAN uplink settings.

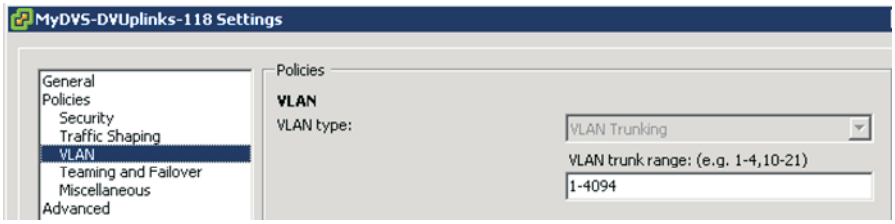

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the VLAN options within a port group.

https://t.me/learningnets

Teaming and Failover

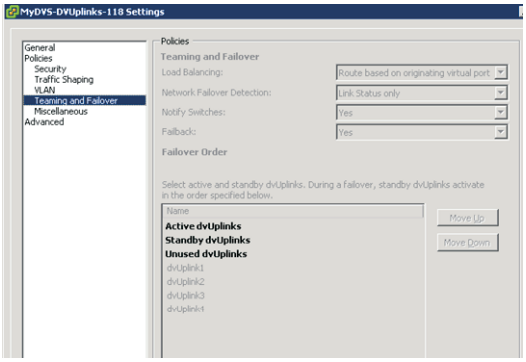

Load balancing, failover detection, failback and NIC active/standby configuration is available from the Teaming policy menu.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the teaming and failover options for port groups.

https://t.me/learningnets

Summary

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

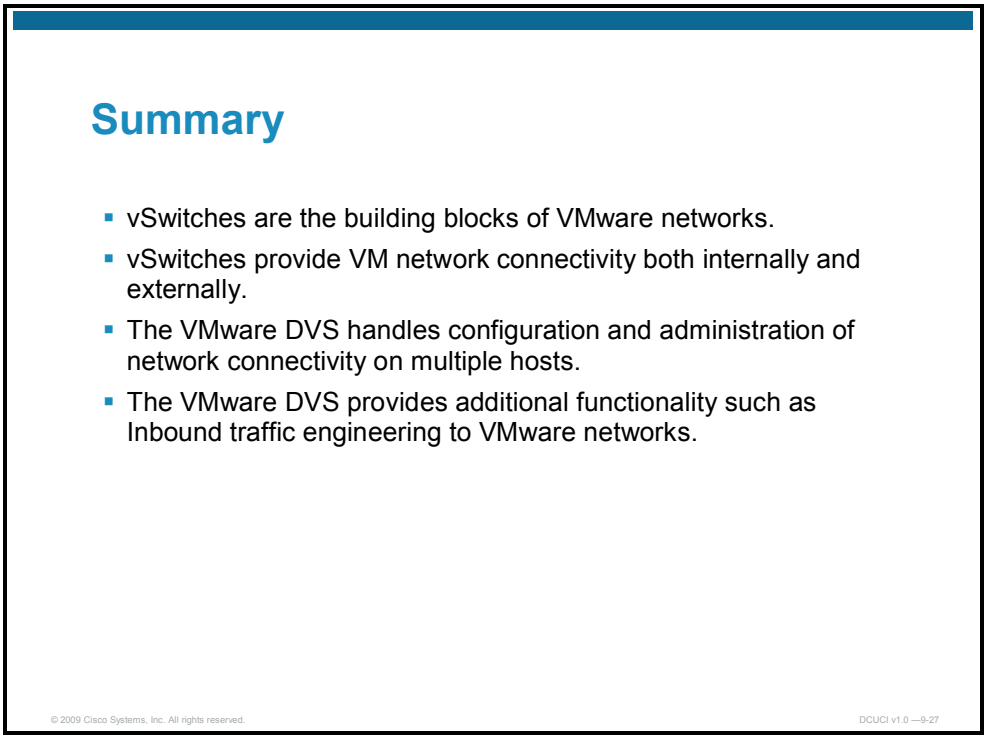

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 6

Installing and Configuring the Cisco Nexus 1000V

Overview

This lesson describes how to install and configure a Cisco Nexus 1000V Switch.

Objectives

Upon completing this lesson, you will be able to install and configure a Cisco Nexus 1000V Switch. This includes being able to meet these objectives:

- Describe how to configure the VSM vSwitch network
- Describe how to create a VSM VM
- Describe the initial configuration of the VSM
- Describe the steps for connecting the VSM to vCenter
- Describe Cisco Nexus 1000V high-availability configuration

https://t.me/learningnets

Configure VSM vSwitch Networks

This topic describes how to configure the Virtual Supervisor Module (VSM) vSwitch network.

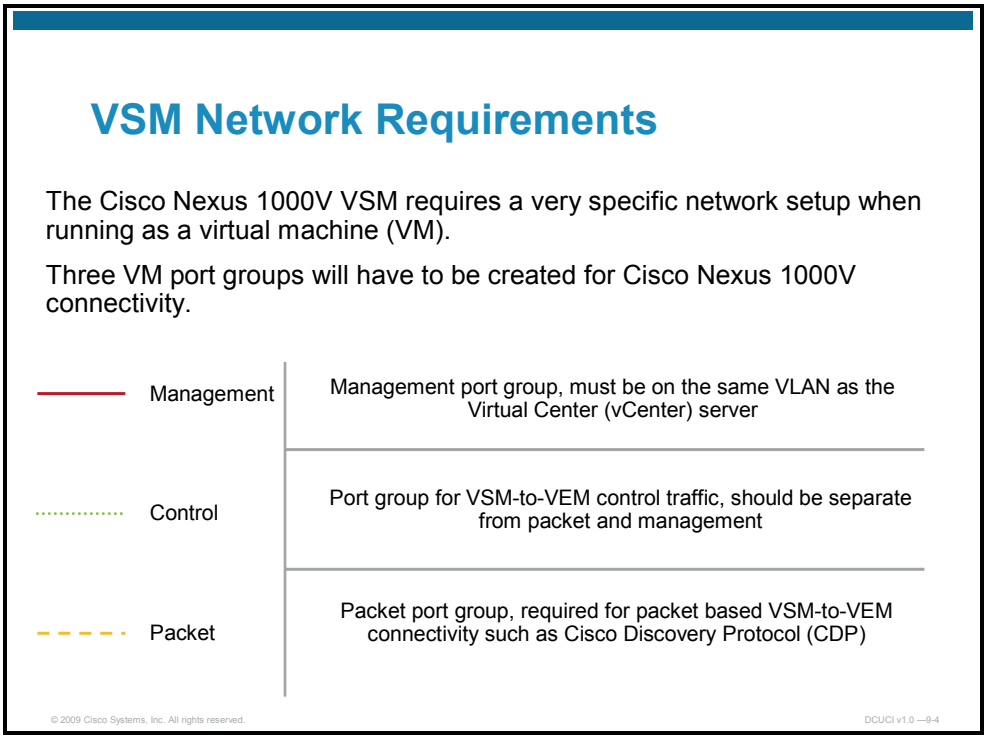

In order to install the Cisco Nexus 1000V VSM as a virtual appliance, three port groups must be created on an existing or new vSwitch. The three port groups correspond to the VSM communication VLANs described in the Cisco Nexus 1000V architecture section. Port groups for management, control, and packet traffic must be available for VSM communication.

Note	VEM stands for Virtual Ethernet Module.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Create VSM Port Groups

- Select the host where the VSM resides
- Click the **Configuration** tab
- Select **Networking**
- Click **Properties**

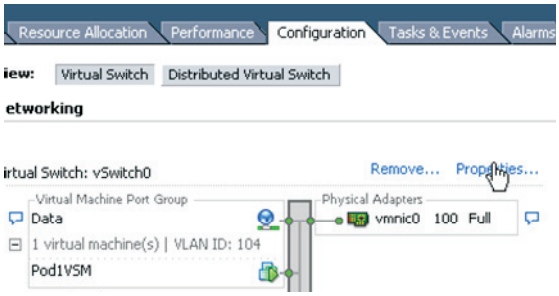

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes how to add VMware port groups for VSM use.

https://t.me/learningnets

Create VSM Port Groups (Cont.)

- In the properties pane, click **Add**

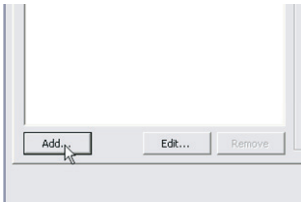

- Select **Virtual Machine**
- Click **Next**

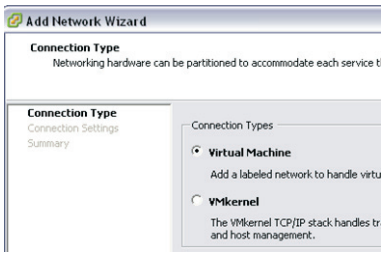

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes how to add VMware port groups for VSM use (continued).

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Create VSM Port Groups (Cont.)

- For Network Label, type **Control**
- For VLAN ID, use the VLAN you've chosen for VSM control traffic
- Click **Next**

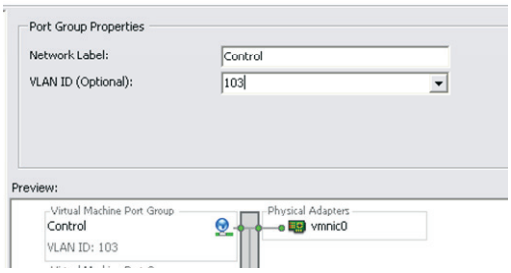

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes how to add VMware port groups for VSM use (continued).

https://t.me/learningnets

Create VSM Port Groups (Cont.)

- Verify your configuration
- Click **Finish**
- Repeat these steps to configure a **management** and **packet** port group

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes how to add VMware port groups for VSM use (continued).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

VSM Virtual Machine Creation

This topic explains how to create a VSM VM.

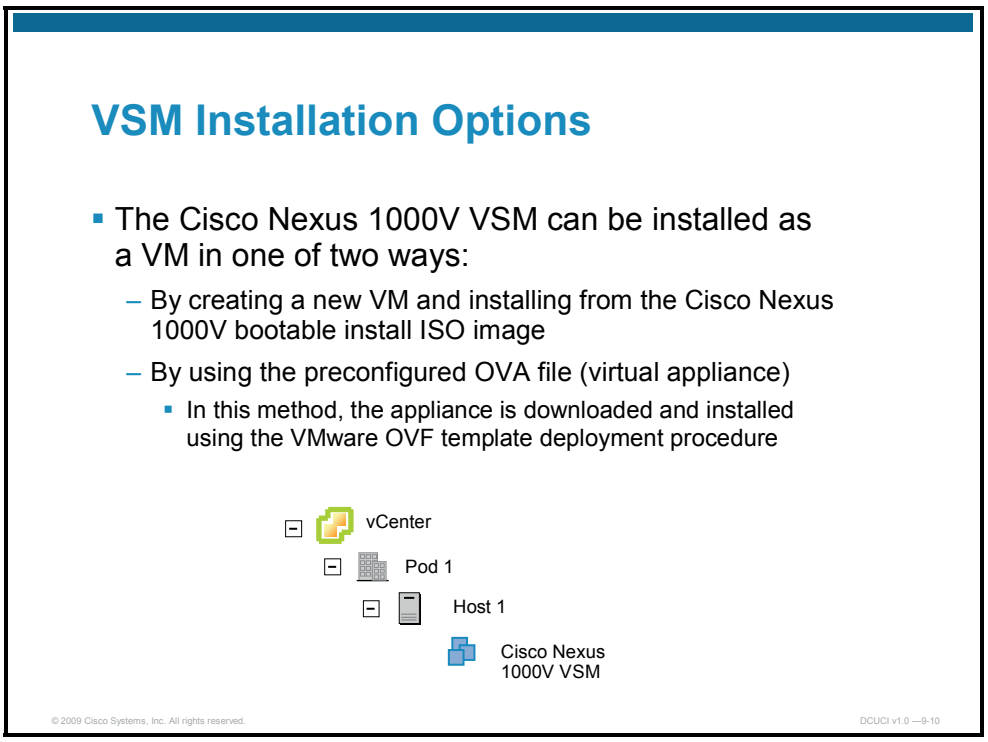

The figure describes VSM installation options.

https://t.me/learningnets

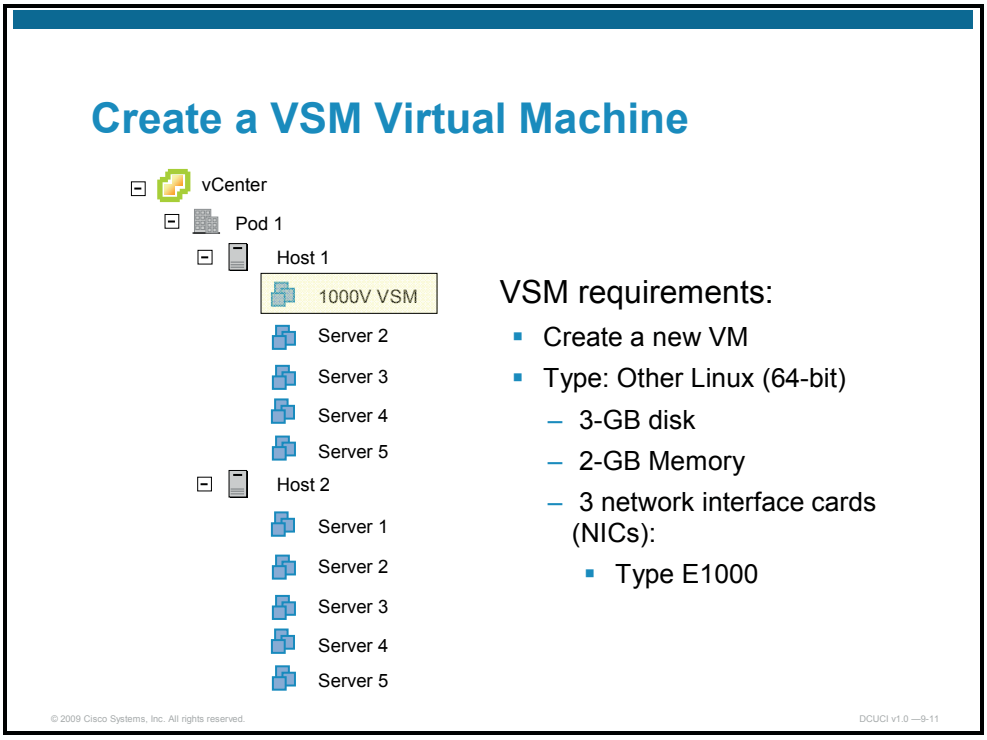

The figure describes the steps required to create a VSM VM.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating a Virtual Machine

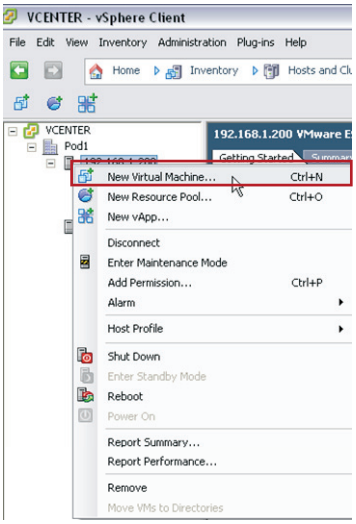

- Right click the host where the VSM will be installed
- Select **New Virtual Machine**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM.

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

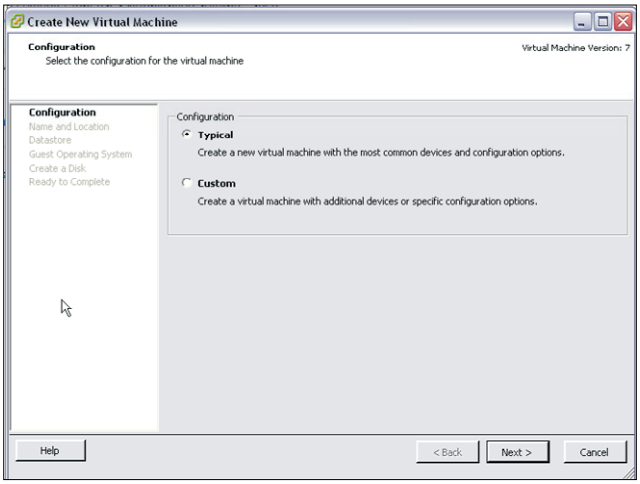

- Select **Typical**
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

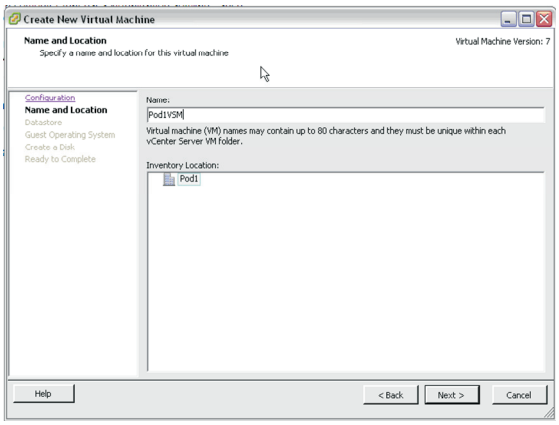

- Name the VM
- Select a data center
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

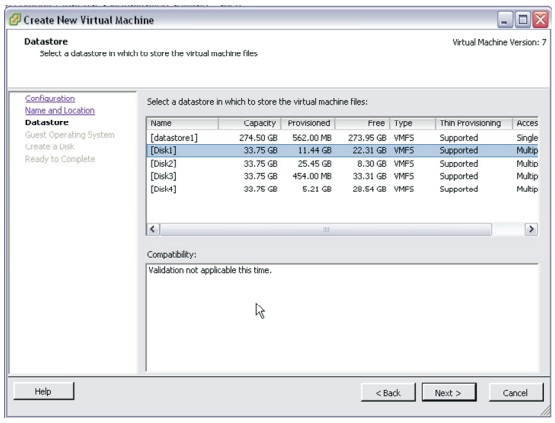

- Choose a Datastore
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

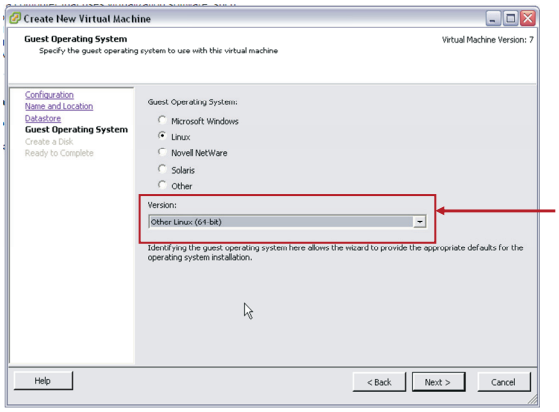

- Select **Linux**
- Choose **Other Linux (64-bit)**
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

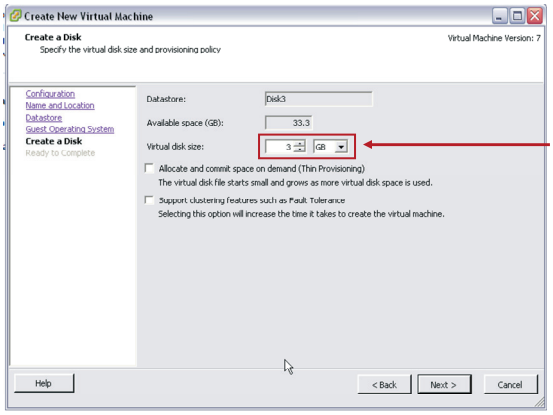

- Choose **3 GB** size
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

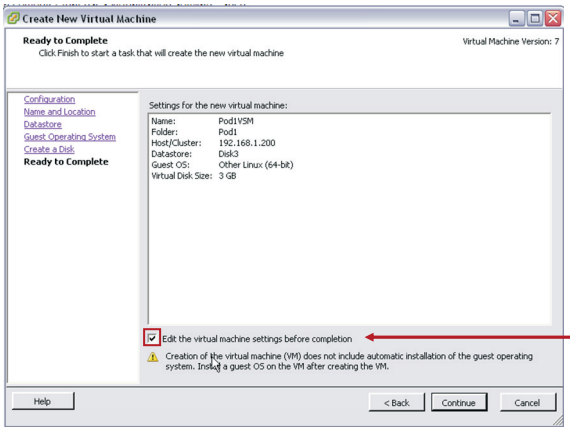

- Select **Edit the virtual machine settings before completion**
- Verify settings
- Click **Continue**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of creating a new VM (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

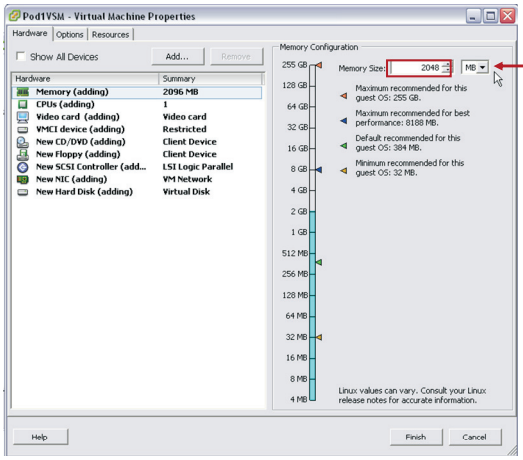

- Select **Memory**
- Change size to **2048**

The VSM requires a minimum of 2-GB of Memory.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings.

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

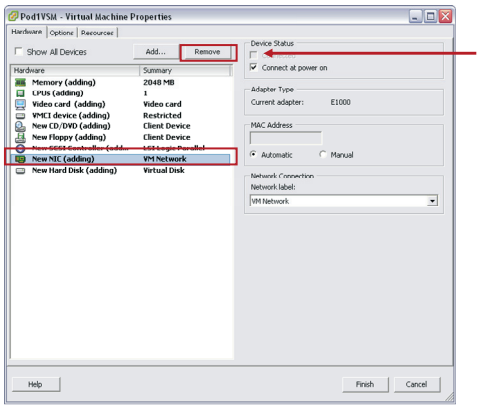

- Select **New NIC**
- Click **Remove**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

- Click **Add**

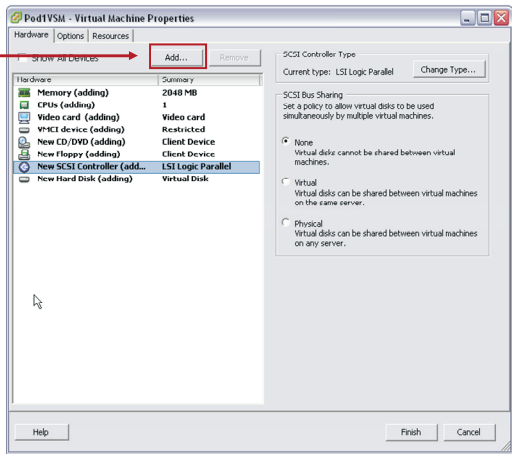

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

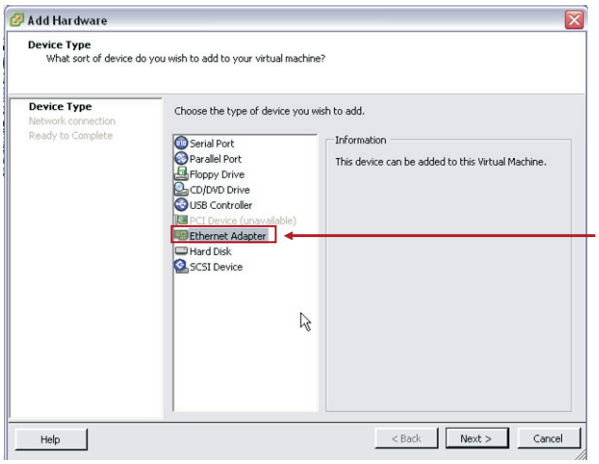

- Choose **Ethernet Adapter**
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings (continued).

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

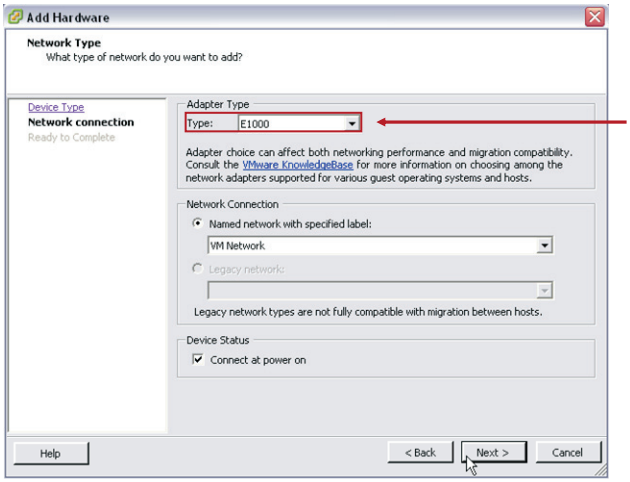

- Choose type **E1000**
- Leave Network Label
- Click **Next**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings (continued).

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

- Verify the configuration
- Click **Finish**
- Repeat this process to add two more NIC adapters

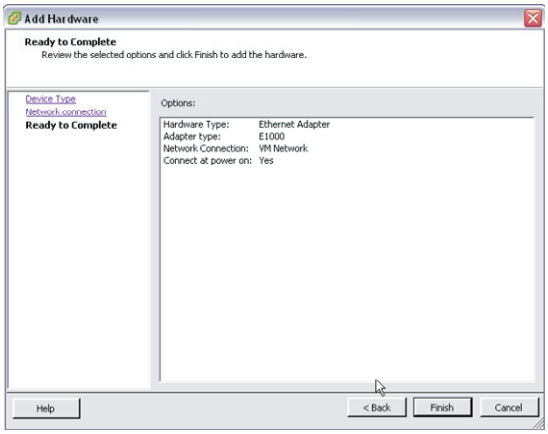

The VSM requires three NICs for network connectivity.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings (continued).

https://t.me/learningnets

Configure the VSM Virtual Machine Networking

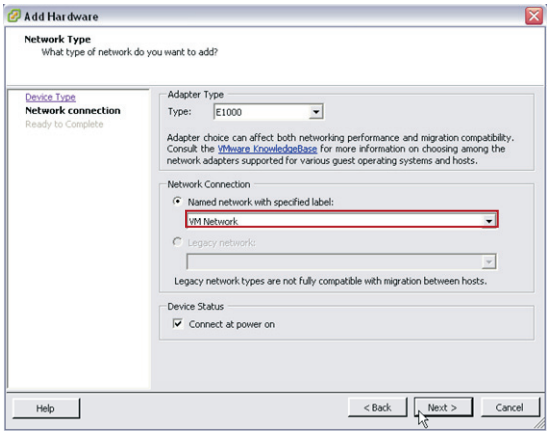

- Configure NIC0 for the **Control** Network
- Configure NIC1 for the **Management** Network
- Configure NIC2 for the **Packet** network

NIC numbering may not exactly match; ensure NICs are assigned in ascending order as described above.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure describes how to add VMware port groups for VSM use.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Verify Correct Network Label Order

Ensure that the network labels are assigned to NICs in the correct order.

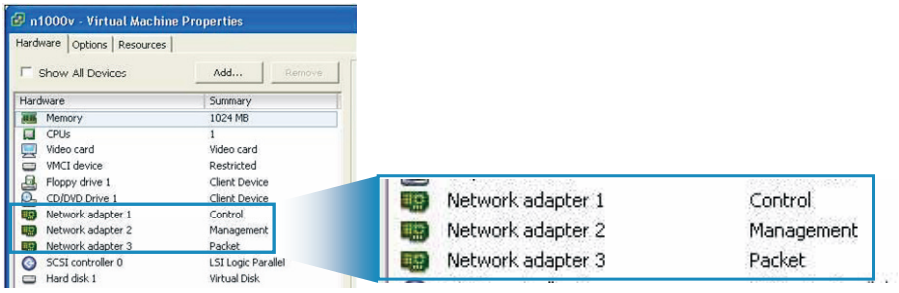

© 2009 Cisco Systems, Inc. All rights reserved.

Virtual NIC adapters on the VSM must be assigned in ascending numeric order for Control, Management, then Packet. Depending on how NICs were added to the VSM, the numbering may not always be NIC0, NIC1, NIC2. The numbering does not matter as long as control is assigned to the lowest numbered NIC, followed by Management and finally Packet. The VSM operating system will assign each to the appropriate traffic based on lowest to highest numbering.

https://t.me/learningnets

Creating a Virtual Machine (Cont.)

Attach the VSM to a bootable Cisco Nexus 1000V install ISO or disk.

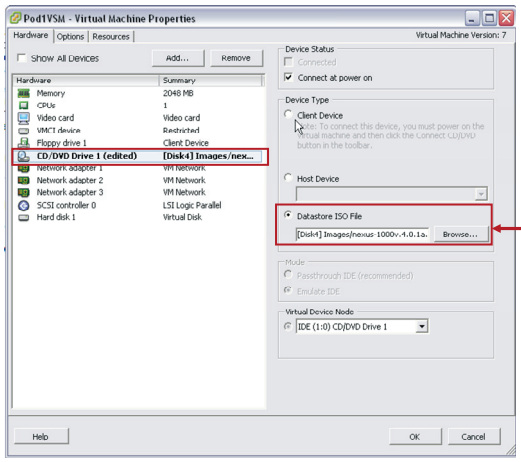

- Choose **CD/DVD Drive**
- Select **Datastore ISO**
- Choose the location of your Cisco Nexus 1000V VSM ISO file
- Click **OK**

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of editing VM settings.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Initial VSM Configuration

This topic describes the initial configuration of the VSM.

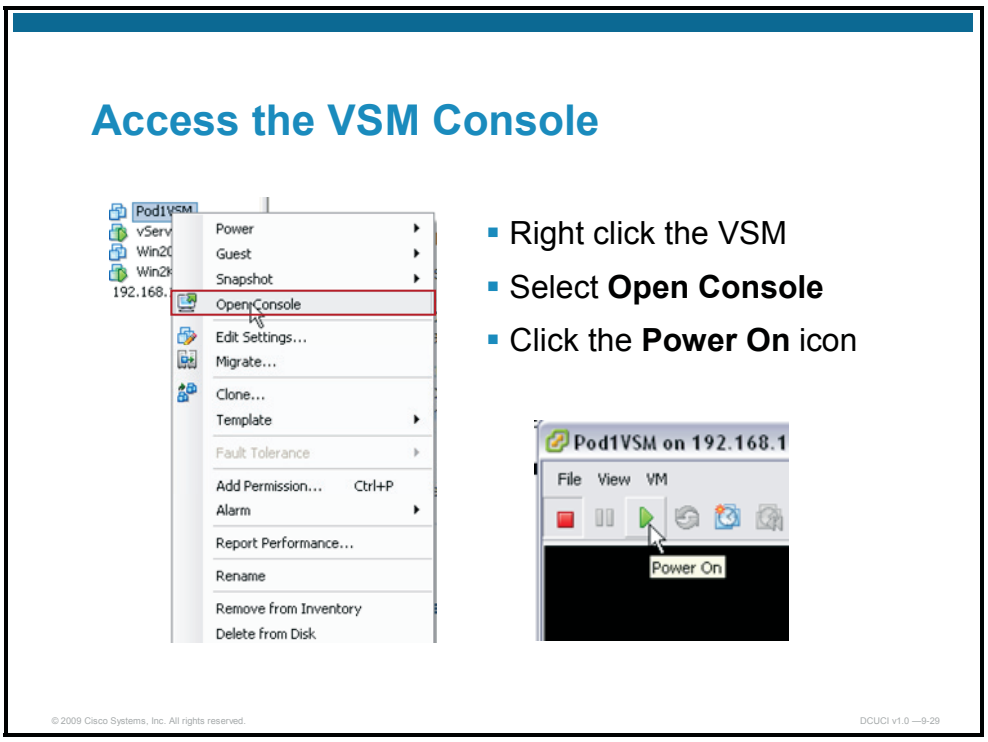

The figure shows the process of configuring the VSM.

https://t.me/learningnets

Access the VSM Console (Cont.)

When the boot menu appears, select **Install Cisco Nexus 1000V and bring up the new image.**

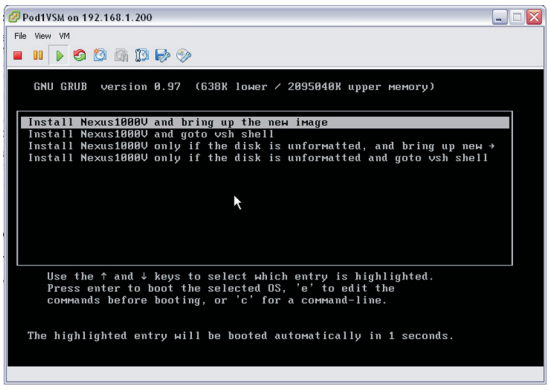

This process may take up to five minutes.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure shows the process of configuring the VSM.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Initial Setup

When prompted, enter and confirm the administrator password. Type **yes** to enter basic configuration.

```
---- System Admin Account Setup ----
Enter the password for "admin": 1234Qwer
Confirm the password for "admin": 1234Qwer
[########################################] 100%


---- Basic System Configuration Dialog VDC: 1 ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
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
```

© 2009 Cisco Systems, Inc. All rights reserved.

Most Cisco devices provide a basic setup script. You will be prompted for this script in one of three events:

- When booting a new switch
- After a write erase, reload
- By typing **setup** from the command-line interface (CLI)

This script will step the user through configuration of basic connectivity options and switch identification.

https://t.me/learningnets

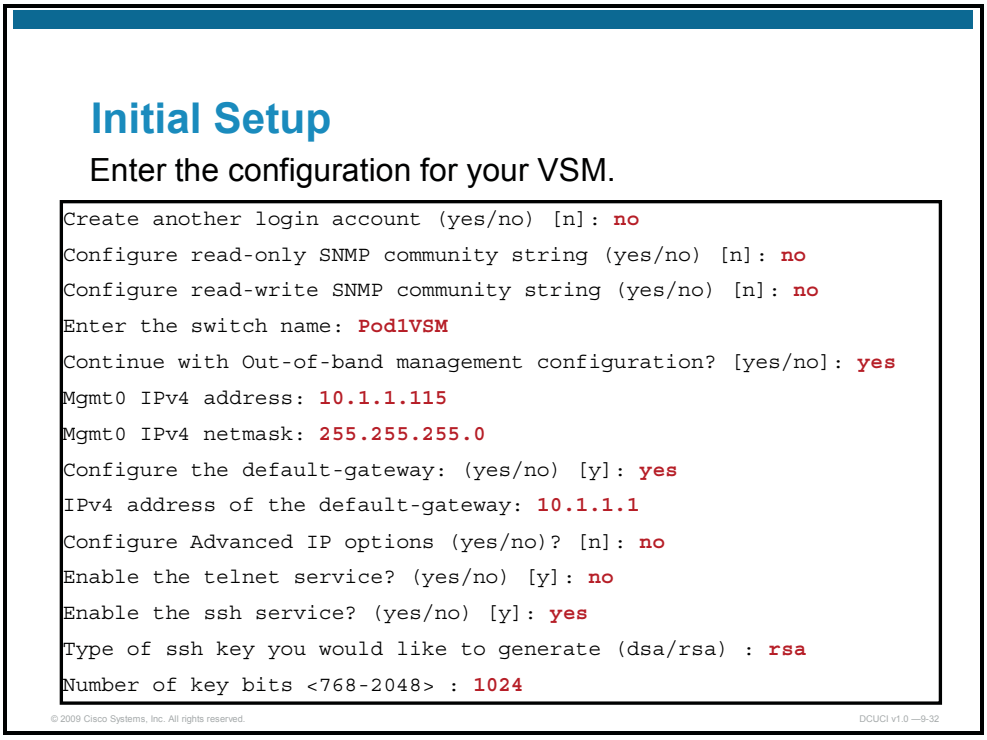

The figure shows the process of configuring the VSM.

https://t.me/learningnets

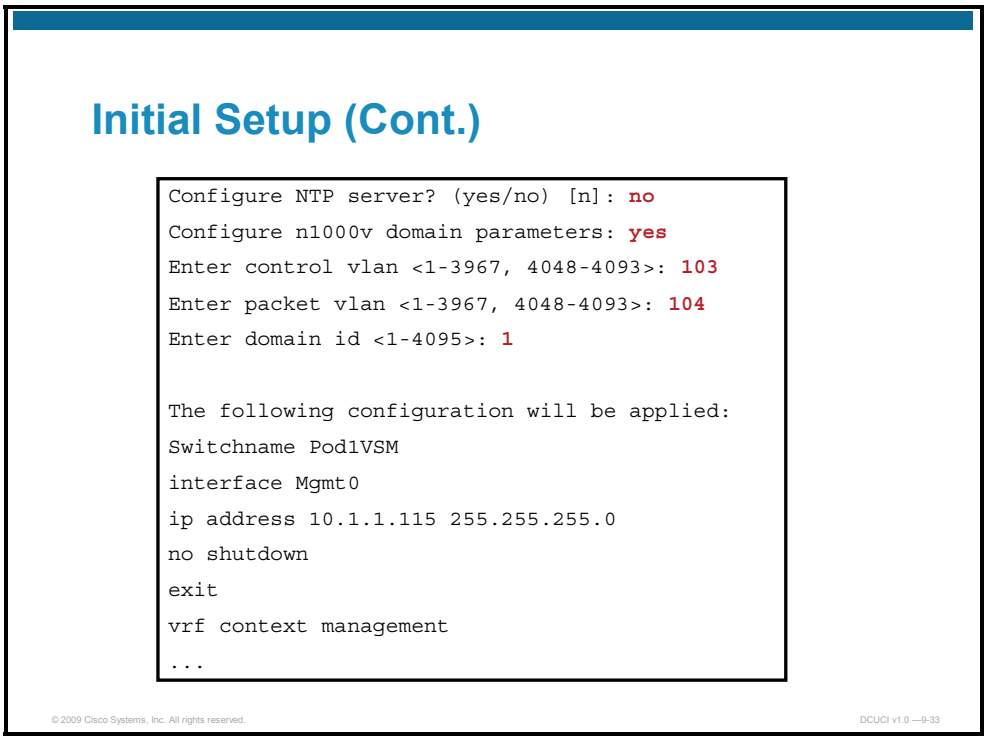

The figure shows the process of configuring the VSM (continued).

https://t.me/learningnets

Connect the VSM to vCenter

This topic describes the steps for connecting the VSM to vCenter.

Install the Plug-In for the New VSM

- Open a web browser on your management network and direct it to the IP address configured for management of the VSM
- Download the **extension.xml** file

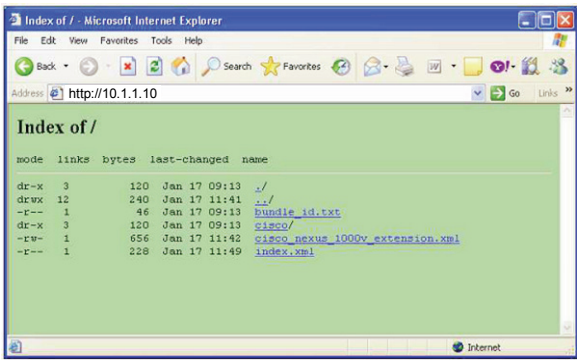

This plug-in is specific to each VSM and must be downloaded and installed individually for each VSM.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure explains the process of connecting a VSM to vCenter.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

VSM Plug-In

The VSM plug-in is an authentication method for the VSM. Without the plug-in installed, the VSM is not authorized to communicate with vCenter.

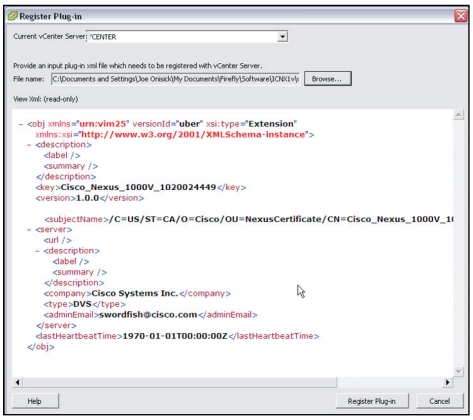

© 2009 Cisco Systems, Inc. All rights reserved.

The VSM plug-in is unique for each VSM and must be manually installed. This acts as an authentication method for each VSM and prevents unauthorized DVS configuration from un-registered VSMs.

https://t.me/learningnets

Install the Plug-In for the New VSM (Cont.)

Step 1: Open the vCenter **Plug-Ins** menu

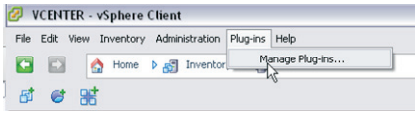

Step 2: Right click the white space and click **New Plug-In**

Step 3:

- Select the plug-in downloaded from the VSM
- Click **Open**

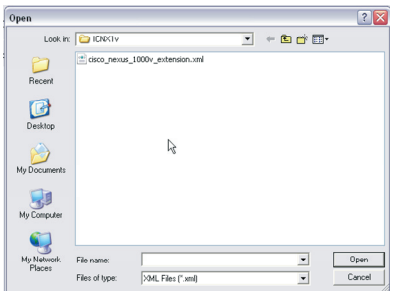

You may need to expand the window to obtain clear white space.

© 2009 Cisco Systems, Inc. All rights reserved.

The figure explains the process of connecting a VSM to vCenter.

https://t.me/learningnets

Install the Plug-In for the New VSM (Cont.)

Verify your plug-in and click **Register Plug-In.**

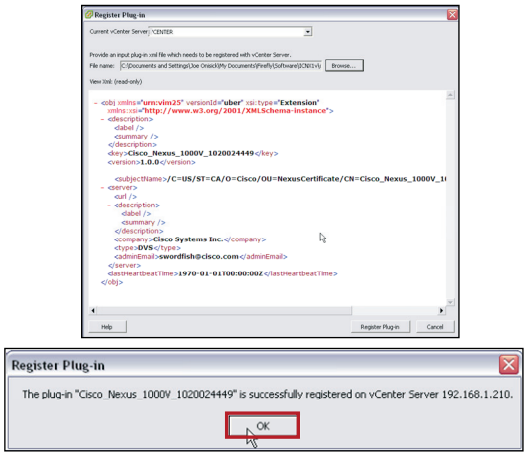

© 2009 Cisco Systems, Inc. All rights reserved.

The figure explains the process of connecting a VSM to vCenter (continued).

https://t.me/learningnets

Verify Connectivity

Ping the default gateway to verify connectivity.

```
PING 10.1.1.1 (10.1.1.1): 56 data bytes
Request 0 timed out
64 bytes from 10.1.1.1: icmp_seq=1 ttl=63 time=3.848 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=63 time=1.996 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=63 time=1.724 ms
64 bytes from 10.1.1.1: icmp_seq=4 ttl=63 time=1.98 ms
--- 10.1.1.1 ping statistics ---
5 packets transmitted, 4 packets received, 20.00% packet loss
round-trip min/avg/max = 1.724/2.387/3.848 ms
```

The first ping is typically lost as the switch uses Address Resolution Protocol (ARP) to gain the gateway MAC address.

© 2009 Cisco Systems, Inc. All rights reserved.

Before beginning the connectivity configuration on the VSM, ensure the VSM has connectivity to the vCenter server from the management interface. This connectivity will be required for the VSM to provide configuration commands to vCenter for Cisco Nexus 1000V VEMs.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

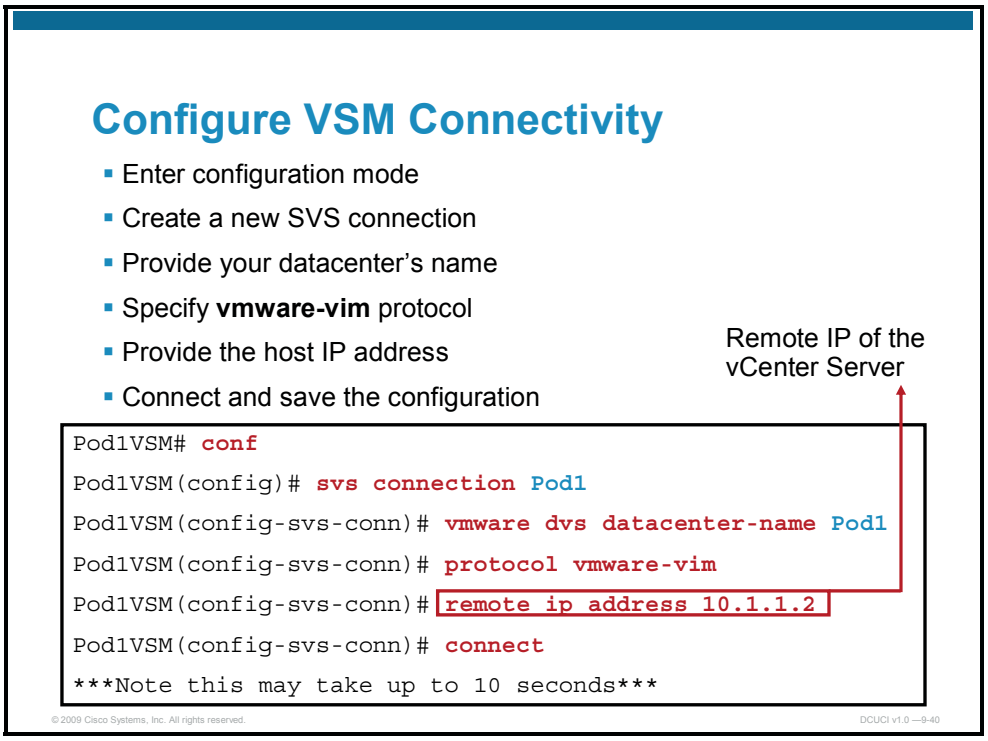

A Server Virtualization Switch (SVS) connection is the configuration point for VSM to vCenter connectivity. The SVS connection is given a name and then provided with parameters that pertain to the connection address and type. In order to connect to vCenter and create a DVS, the SVS connection must contain:

- The VMware datacenter name configured in vCenter (case sensitive)
- The protocol type (currently only **vmware-vim** type is supported)
- IP address of the vCenter server

https://t.me/learningnets

Verify VSM Connectivity

After the VSM properly connects, you should see output showing the creation of a DVS.

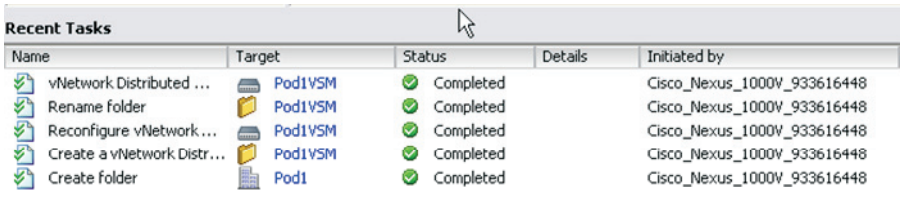

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the VMware event output shown after successful VSM connection.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Cisco Nexus 1000V High Availability Configuration

This topic describes Cisco Nexus 1000V high-availability configuration.

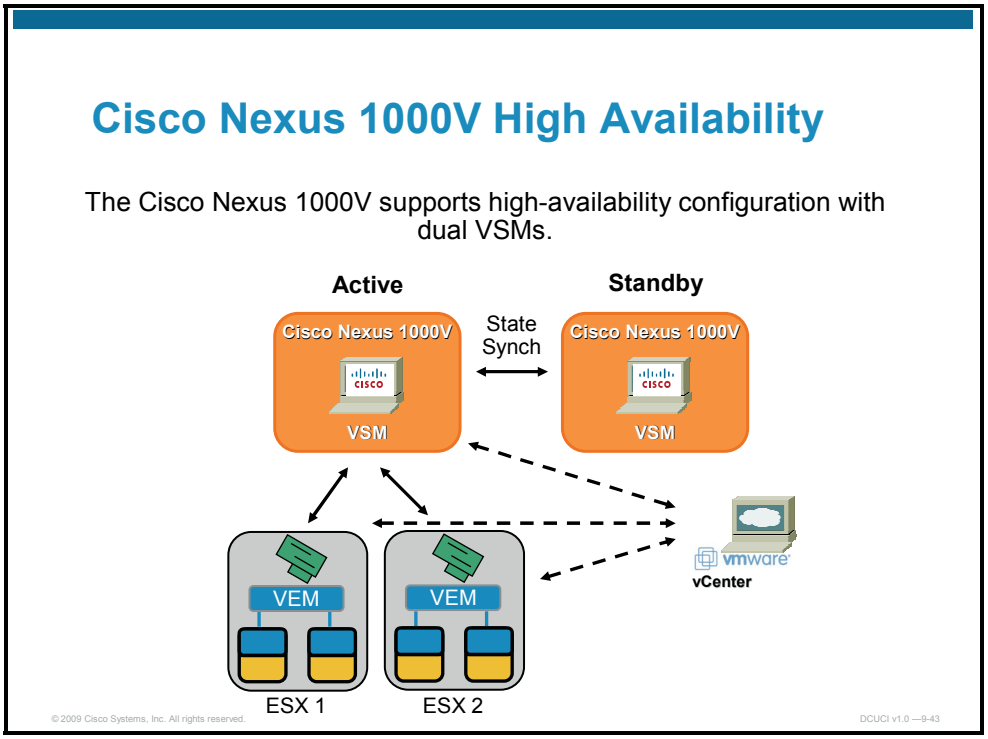

The Cisco Nexus 1000V supports dual supervisors configured for high availability. In this mode, one supervisor is active and the other is passive. The two supervisors maintain state synch using Mobile TeleSystems (MTS) or application inter-process communication (AIPC) protocols. All services are run on the active supervisor and are kept in a ready mode for failover on the standby supervisor.

https://t.me/learningnets

Supervisor Modes

- In Dual Supervisor mode, the active (primary) supervisor will be the supervisor in slot 1; the standby (secondary) supervisor will reside in slot 2.
- All services are run on the active supervisor and exist in a standby mode on the second supervisor. These services are maintained in synch for failover.
- All management and configuration is performed on the active supervisor.
- Configuration of primary and secondary can be performed during initial switch setup.

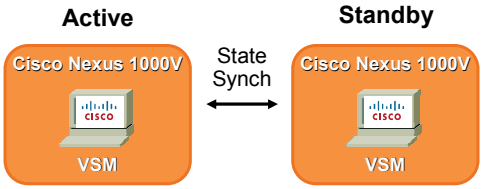

© 2009 Cisco Systems, Inc. All rights reserved.

Dual supervisor modes operate in a similar fashion to dual supervisor physical chassis. The active supervisor runs processes and is available for configuration. The standby supervisor remains in stateful synch for failover purposes. There are three possible roles for a VSM: Primary, secondary and standalone. In single supervisor configurations, the VSM is in standalone mode. In dual supervisor configurations, one supervisor will be primary (active) and one supervisor will be secondary (standby). The **show system redundancy status** command shows the currently running configuration.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

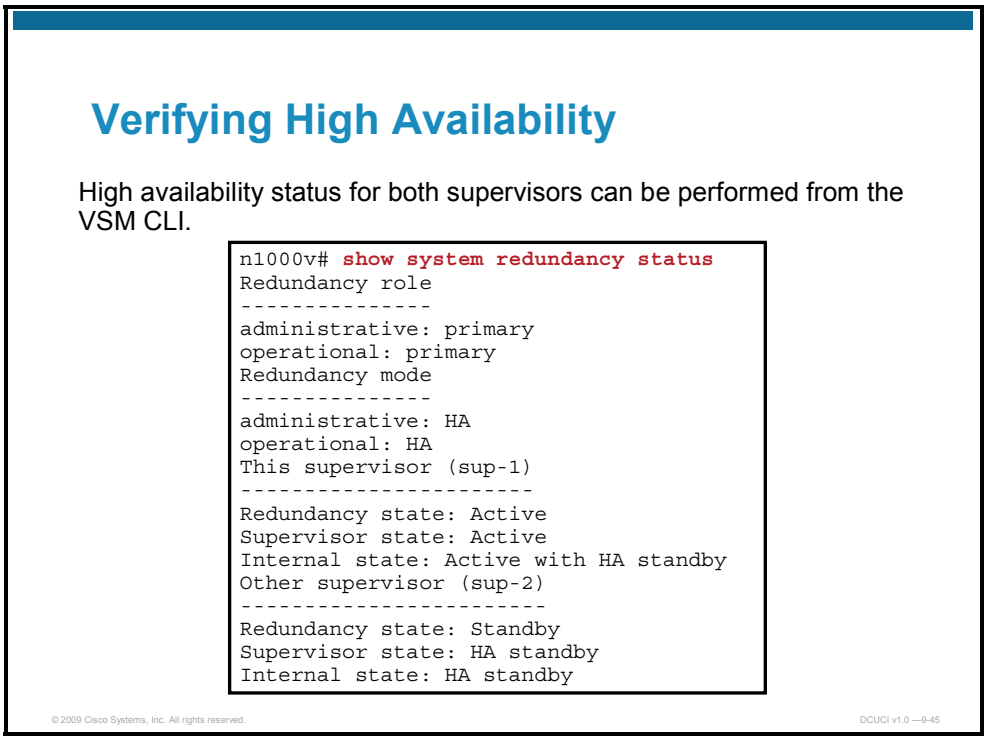

This figure shows the steps needed to verify high availability status on the VSMs.

https://t.me/learningnets

Verifying High Availability (Cont.)

The **show module** command shows all modules controlled by the VSM, and the high availability status of available supervisors.

```
n1000v# show module
Mod  Ports   Module-Type                 Model       Status
---  -----  --------------------------  ----------  ---------
1     0     Virtual Supervisor Module  Nexus1000V  active *
2     0     Virtual Supervisor Module  Nexus1000V  ha-standby
3     248   Virtual Ethernet Module    NA          ok
...
```

© 2009 Cisco Systems, Inc. All rights reserved.

This figure shows the steps needed to verify high availability status on the VSMs (continued).

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

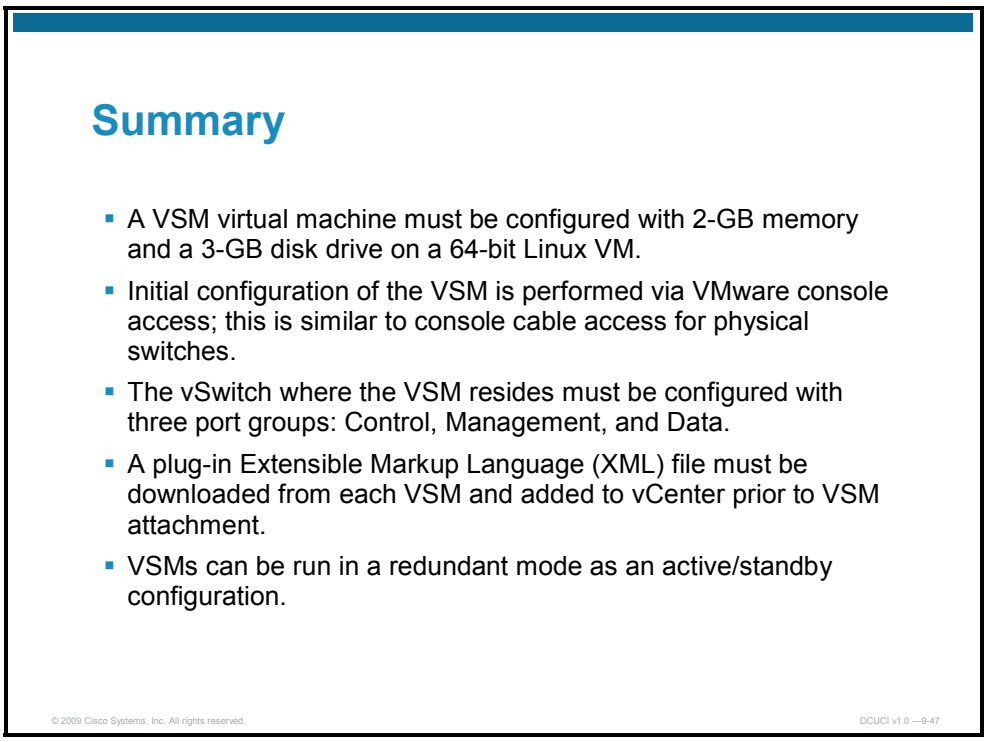

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 7

Configure Basic Cisco Nexus 1000V Networking

Overview

This topic describes the basic Cisco Nexus 1000V network configuration.

Objectives

Upon completing this lesson, you will be able to configure basic Cisco Nexus 1000V networking. This includes being able to meet these objectives:

- Describe port profiles
- Describe VLAN configuration
- Describe the configuration of uplink profiles
- Describe the configuration of VM data port profiles
- Describe Cisco Nexus 1000V port channel configuration
- Describe the process of adding hosts to a VSM
- Describe how to add VMs to a VSM

https://t.me/learningnets

Port Profiles

This topic describes port profiles.

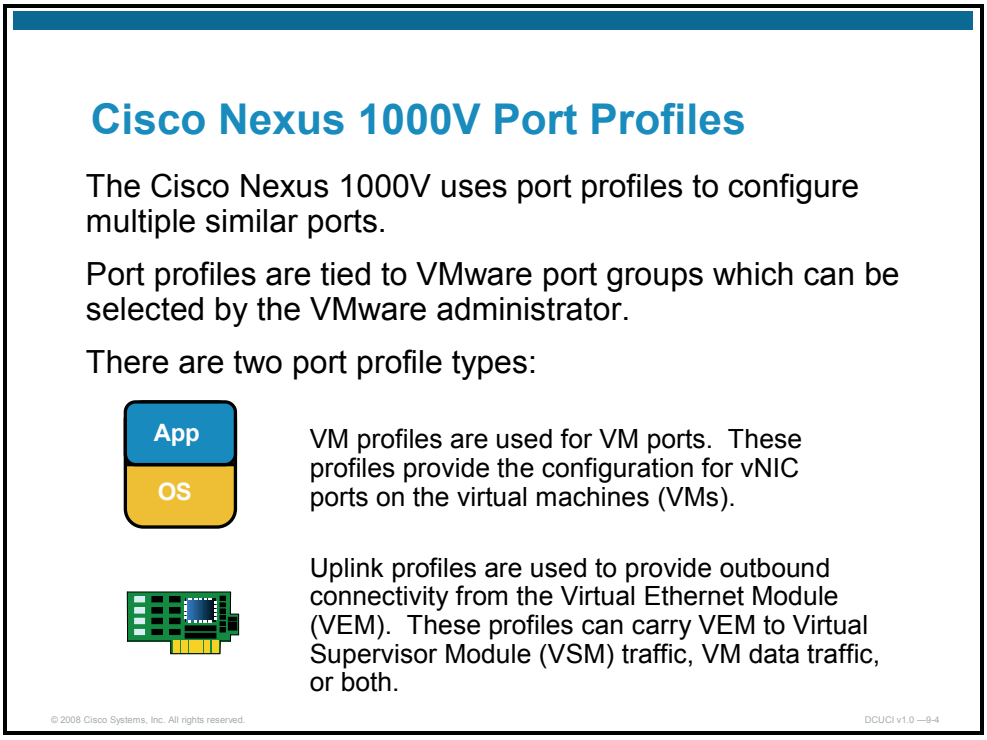

The Cisco Nexus 1000V uses a new concept called port profiles. Port profiles are used to provide port configuration for the ports assigned to the profile. There are two types of port profile; VM and uplink. Uplink profiles provide all outbound communication for VMs, as well as all VSM to VEM communication on the control and packet Virtual Local Area Networks (VLANs). VM port profiles provide the configuration for VM virtual network interface card (vNIC) ports.

Port profiles are used as a central configuration point for multiple ports. By utilizing port profiles, groups of ports requiring the same network configuration can be rapidly configured. Once a port profile has been created and configured, ports can be assigned to the profile, and will receive the configuration contained in that profile.

Port profiles tie directly to VMware port groups on the Distributed Virtual Switch (DVS) controlled by the Cisco Nexus 1000V. Once a profile has been created, a VMware port group name can be assigned to the profile. By default, the port group name will be the same as the port profile name, but this is configurable. When a port profile is enabled, the corresponding port group is created on the DVS within Virtual Center (vCenter).

VMware administrators assign ports to a port group from the vSphere client. When a port is made a member of a port group, the corresponding port profile configuration is applied to the port on the Virtual Ethernet Module (VEM). Any configuration changes to the port profile are immediately assigned to member ports and stored in the Cisco Nexus 1000V running configuration file.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

By using port profiles, hundreds or thousands of VMs can be provisioned rapidly with detailed network configurations such as port state, quality of service (QoS) tagging, and access control list (ACL) controls. Additionally, port profiles reduce the risk of misconfiguration on groups of similar ports by maintaining the configuration in a central location. While individual port configuration is still possible using the Cisco Nexus 1000V, it is recommended to use port profiles rather than configuring ports directly. This method reduces administration time and simplifies network troubleshooting.

https://t.me/learningnets

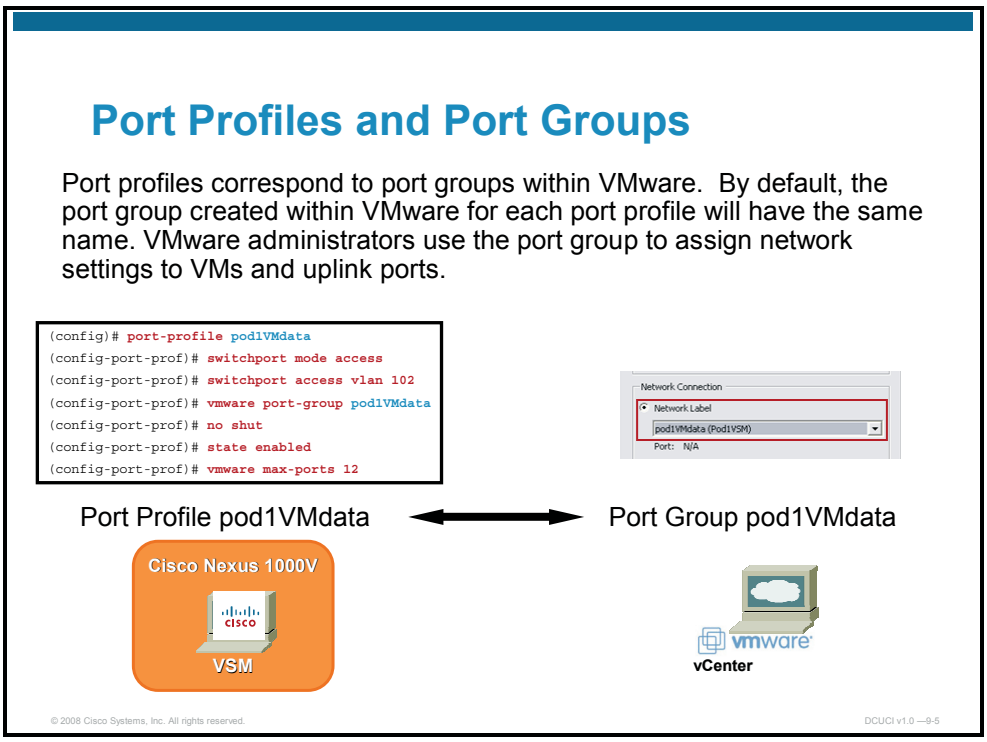

When a port profile is created and enabled, a corresponding port group is created in VMware. By default, this port group will have the same name as the profile, but this name is configurable. VMware administrators use the port profile to assign network settings to VMs and uplink ports. When a VMware ESX Host port (vmNIC) is added to a Cisco Nexus 1000V controlled DVS, an available uplink port group is assigned and those settings are applied. When a network interface card (NIC) is added to a VM, an available VM port group is assigned and the network settings associated with that profile are inherited.

Note	A NIC in VMware is represented by an interface called a vmNIC. The vmNIC number is allocated during VMware installation.

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

Uplink Profiles

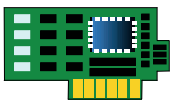

- Uplink profiles can be assigned to physical NICs (vmNIC) on VMware ESX Hosts.
- These profiles can be one of two types:
  - System uplinks which carry packet and control VLANs between the VEM and the VSM
  - VM uplinks which carry data traffic from VMs destined for the network
- A single uplink profile can also be configured to carry both types of traffic.

© 2008 Cisco Systems, Inc. All rights reserved.

Uplink profiles are used to provide outbound connections from the VEM to the VSM, as well as carry VM data traffic to the network. Uplink profiles are used to configure vmNICs or physical host ports. These profiles typically contain information such as trunking and port channel behavior.

https://t.me/learningnets

VM Profiles

- VM profiles provide the configuration for VM ports.
- VM vNIC ports are assigned to a VM port profile.
- VM port profiles will require an uplink port profile in order to access the physical network.
- VM profiles can be configured in a VLAN with no corresponding uplink profile to create internal VM networks.

© 2008 Cisco Systems, Inc. All rights reserved.

VM profiles are used to provide configuration for VMs. These profiles require an uplink profile to access the physical network.

VM profiles can be configured without a corresponding uplink profile to create internal VM networks. If a VM profile is created accessing a VLAN that is not trunked to the physical network, then the assigned VMs will only be able to communicate with other VMs assigned to the profile in the same host. This configuration is similar to creating internal only vSwitches or port groups within a standard VMware networking environment.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

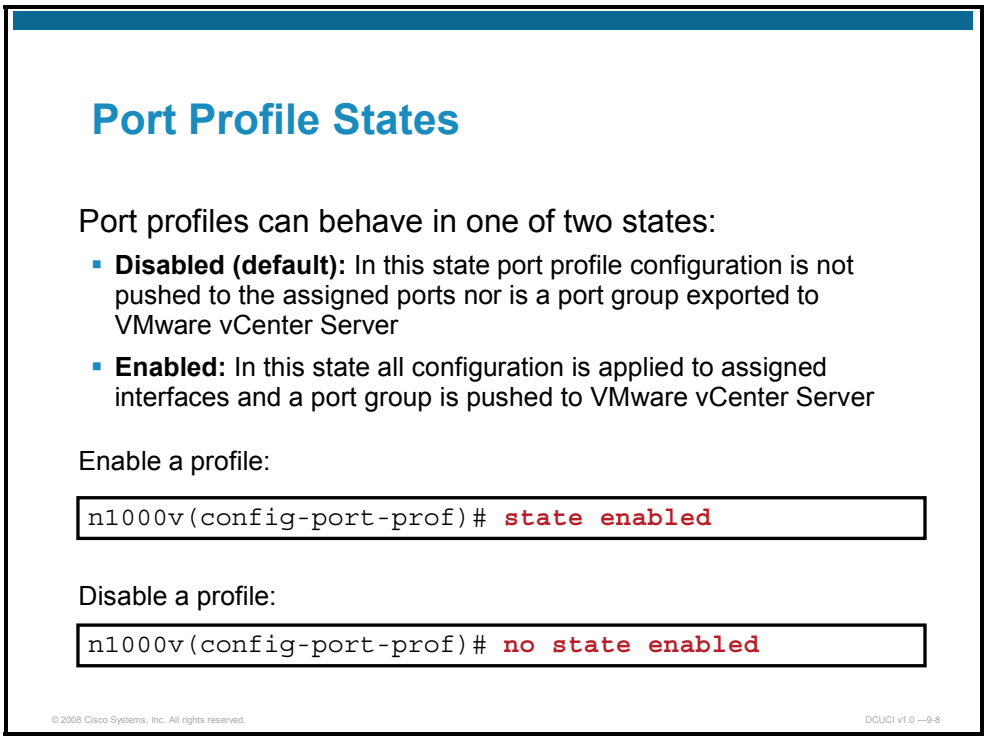

Port profiles can be either enabled or disabled as a whole. The configuration of a port profile is only applied to assigned ports if it is in an enabled state. Additionally, no VMware port group is created for disabled port groups.

The enabled and disabled state of a port group is separate from the *shut* and *no shut* parameter within the profile. The enabled and disabled state applies to the profile itself, where *shut* and *no shut* applies to the port configuration of member ports.

https://t.me/learningnets

Port Profile Usage

- Port profiles should be used for all interface configurations. This ensures consistency for configuration of like devices.
- Any configuration performed directly to a port overrides port profile settings and should only be used for testing and troubleshooting.
- All port profile configuration and changes are automatically pushed to assigned ports.

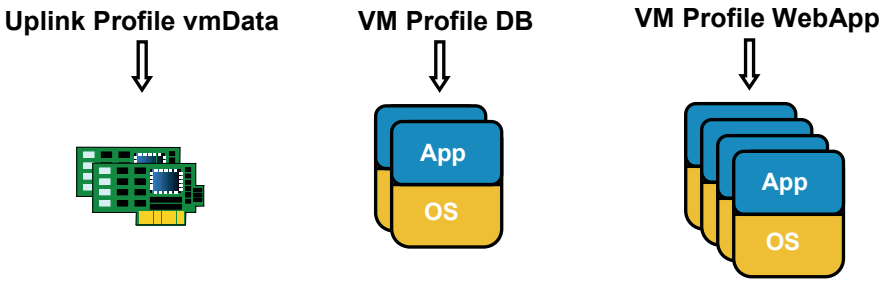

© 2008 Cisco Systems, Inc. All rights reserved.

It is recommended to use port profiles for all port configurations on the Cisco Nexus 1000V. This helps to ensure consistency among ports with similar characteristics, and speeds the deployment process. Any configuration performed on an individual port overrides the port profile configuration. Configuring individual ports is typically only necessary for testing and troubleshooting purposes.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

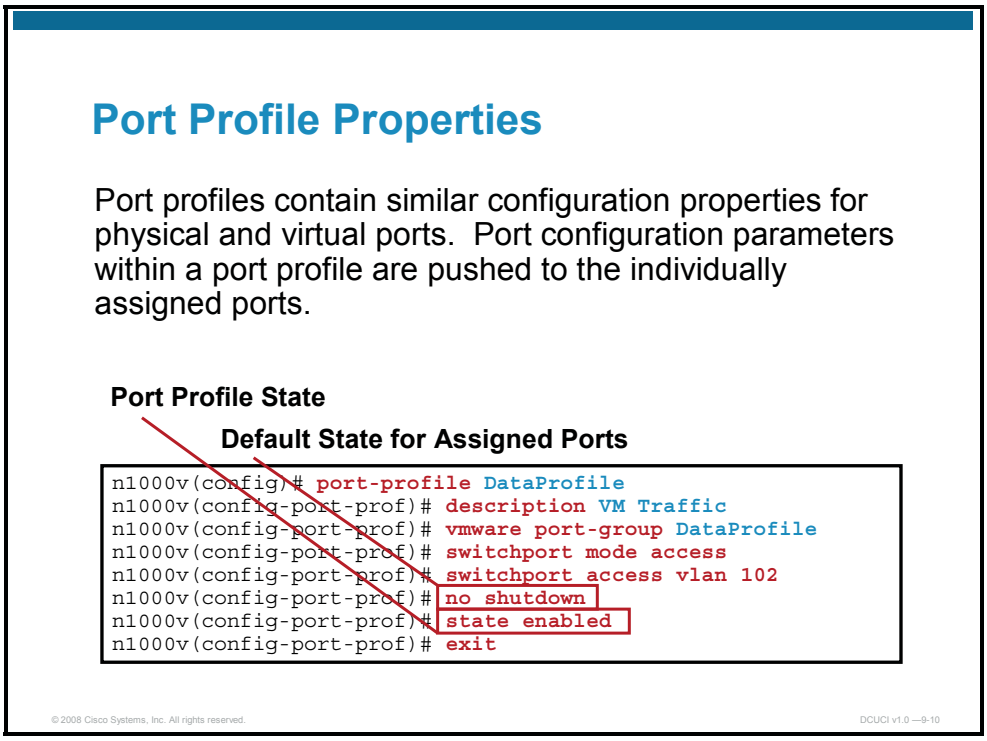

Port profiles are a key element of port configuration on the Cisco Nexus 1000V. Ports within VMware are assigned to port profiles. Best practices dictate that all port configuration should be performed using port profiles.

`acl`

- ACLs to be applied to a port profile

`name`

- Name for the profile

`description`

- Text description for the profile

`capability (uplink)`

- Designates the profile for uplink ports (vmNICs)

`channel-group auto mode on`

- Enables auto port channel configuration for available vmNICs

`interface state (shut/no shut)`

- Enable or disable the ports

`netflow`

- Configure the port for NetFlow

`port security`

- Configure port security settings such as maximum number of connections and MAC binding

`private vlan configuration`

- Configure PVLAN settings

`state (enabled or disabled)`

- Enable or disable the port profile

`switchport mode (access or trunk)`

https://t.me/learningnets

- Single VLAN or trunking multiple VLANs

```
system vlan <vlan list>
```

- Configure the systems VLANs; a system VLAN is used to configure and bring up physical or veth ports before the control plane (CP) has established communication with the data plane (DP)

```
vlan configuration
```

- Configure VLANs for the port

```
vmware port-group name default
```

- Resets a characteristic to the default value

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

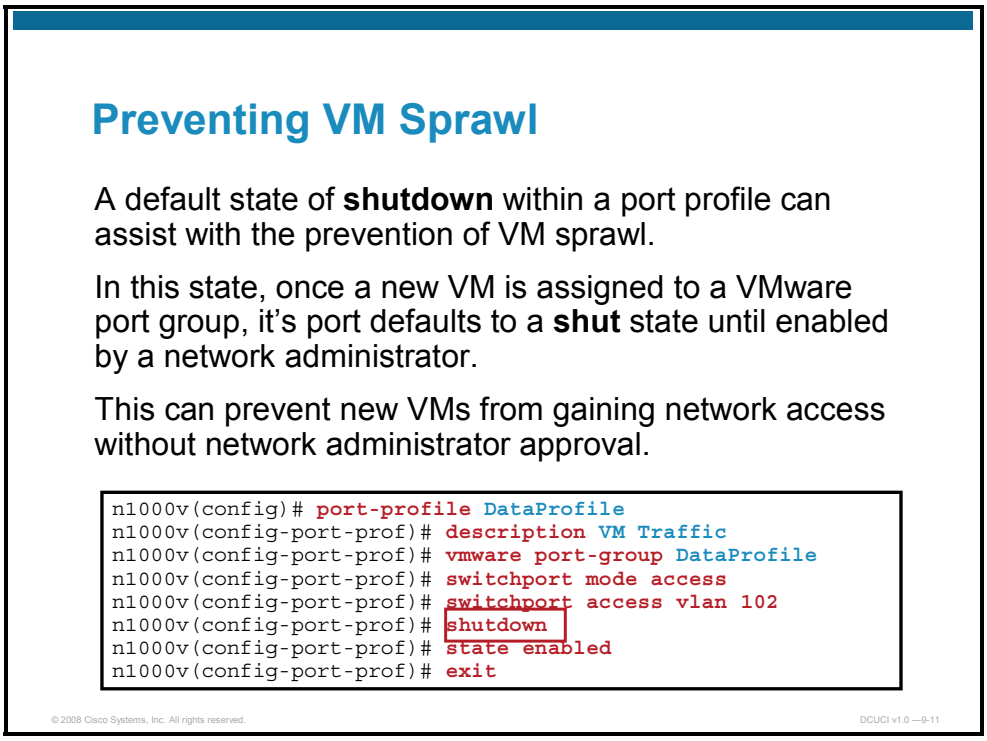

VM sprawl is a new data center challenge that has come about with virtualization. Virtual servers typically make new server deployment very easy, and as such, servers may to be deployed more than would be typical in a standard physical architecture. By utilizing the port profiles on the Cisco Nexus 1000V, VM sprawl can be minimized by not allowing VMs to access physical network resources without network administrator intervention.

By dictating a port profile default port state of *down*, a network administrator can ensure that no new VM has access to the network without approval from network teams. After a new VM is brought online, the VMware administrator needs to contact the network team to enable the port. Until the port is enabled, the virtual server will exist, but it will have no outbound network connectivity.

https://t.me/learningnets

VLAN Configuration

This topic explains VLAN configuration.

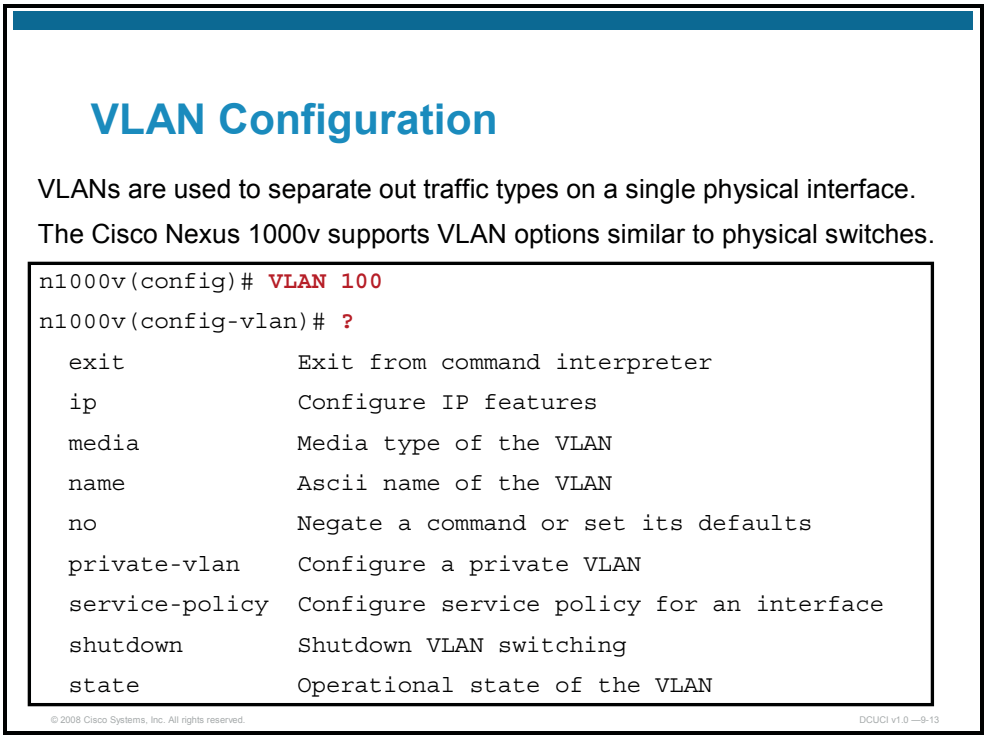

VLANs are used to separate network traffic into multiple virtual networks. This traffic separation is useful for reducing overhead, failure domains, and for increasing security. VLANs segregate traffic at Layer 2.

`exit`

- Exit to the configuration mode

`ip`

- Internet Group Management Protocol (IGMP) configuration

`media`

- Set the media type to Ethernet

`name`

- A name for the VLAN (for reference purposes)

`no`

- Removes a command (followed by the command to be removed)

`private-vlan`

- Configure private VLAN settings; creates primary and secondary private VLANs

`service-policy`

- Configure VLAN ACL

`shutdown`

- Disable switching in the VLAN

`state`

- Disable or activate the VLAN

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

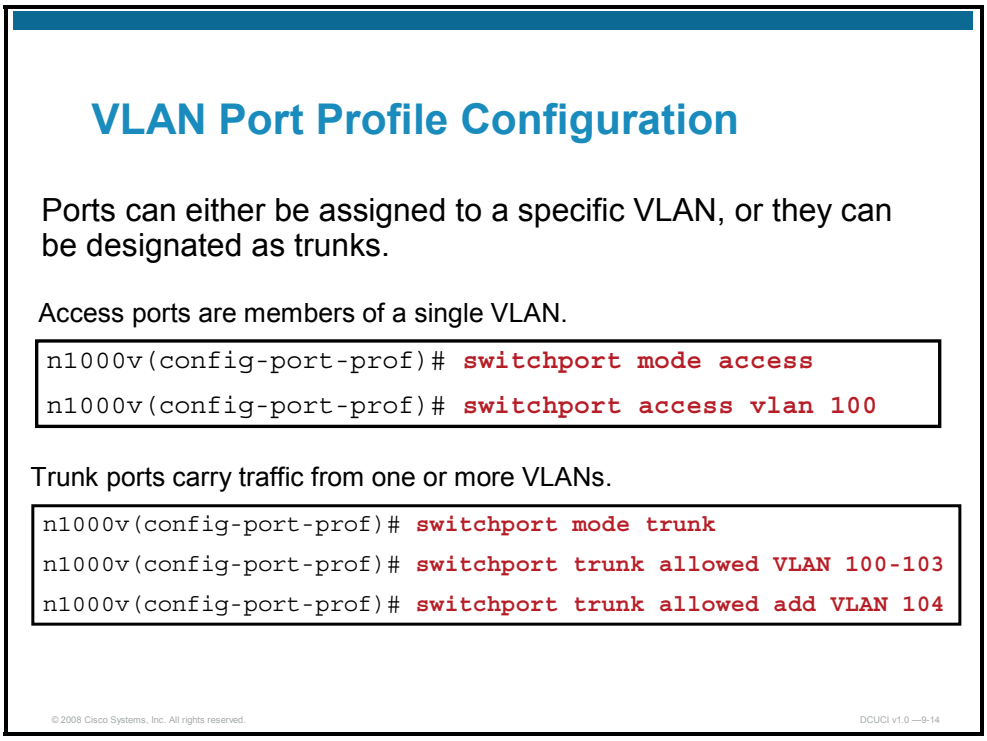

After VLANs have been configured on the switch, port profiles can be assigned to use them. Ports inherit VLAN settings from the port profile. Two port modes are typically used for VLANs:

- **Trunk ports:** These ports carry traffic for multiple VLANs. By default, trunk ports carry data for all VLANs but this can be pruned back to minimize traffic or increase security.
- **Access ports:** Access ports are members of a single VLAN. Access ports are typically configured for the VLAN that they are a member of. In most cases, single server instances (1 operating system, and 1 application) running on bare metal or virtual hardware will be assigned to access ports.

https://t.me/learningnets

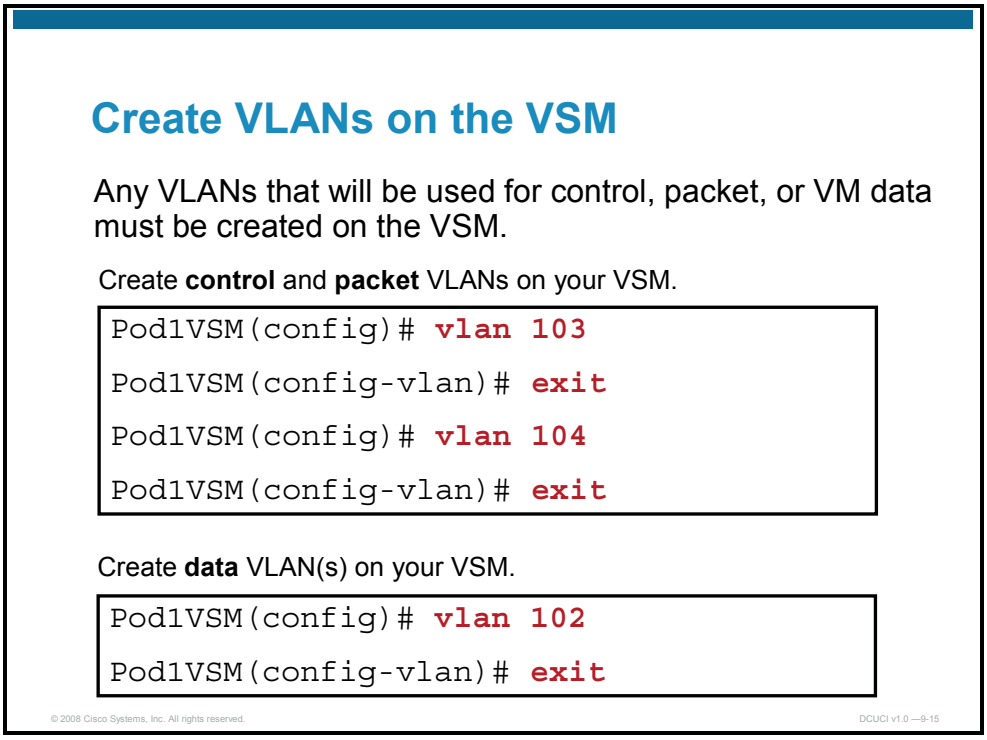

Because the Cisco Nexus 1000V acts as a fully functional switch within the virtual environment, any VLANs being used by attached ports need to be created on the Cisco Nexus 1000V to facilitate switching for those ports. VLAN creation is done from the configuration level of the Cisco Nexus Operating System (Cisco NX-OS) command-line interface (CLI).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating Uplink Profiles

This topic describes the configuration of uplink profiles.

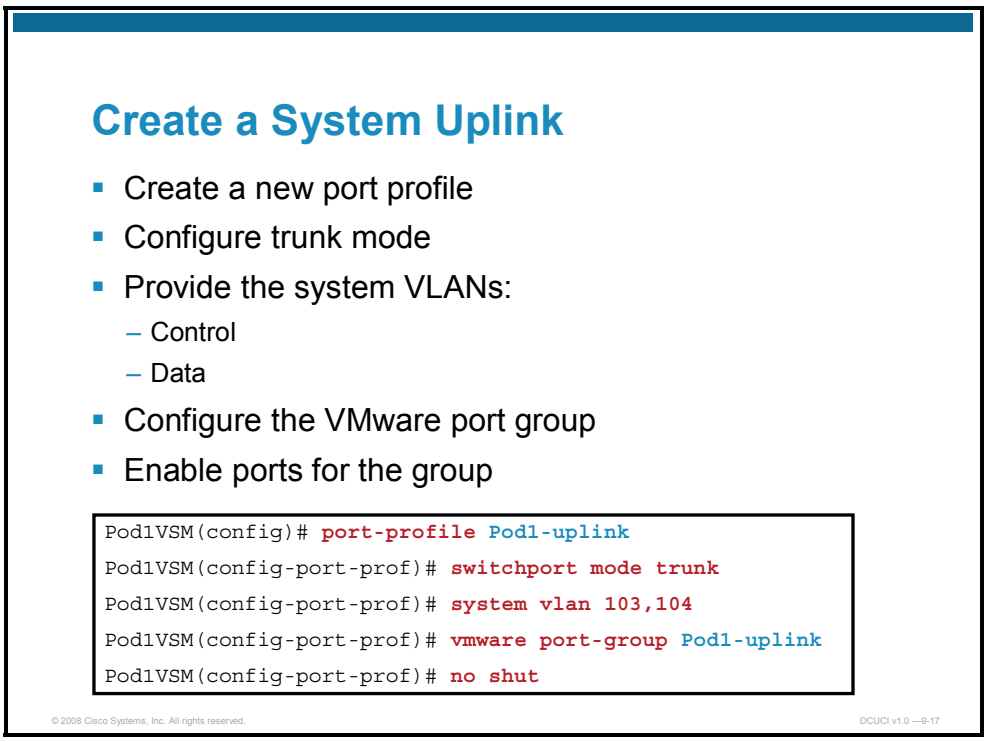

This figure shows how to create a port profile for uplink ports.

https://t.me/learningnets

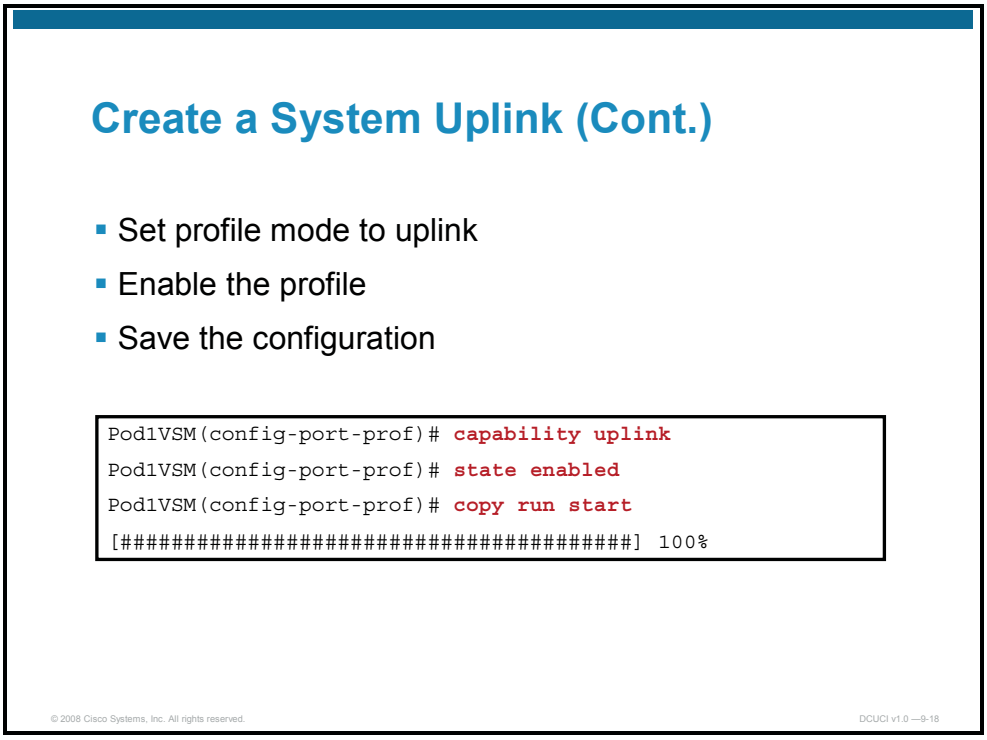

This figure shows how to create a port profile for uplink ports (continued).

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Verify the System Uplink

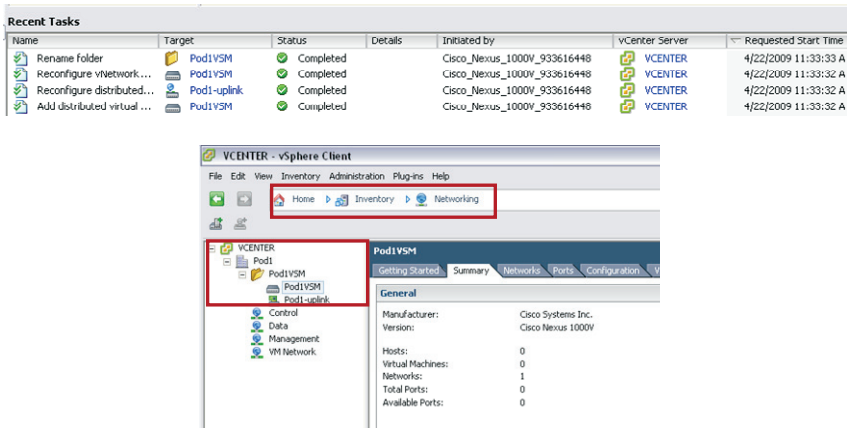

You should see similar output when the configuration is successfully completed.

© 2008 Cisco Systems, Inc. All rights reserved.

This figure shows how to create a port profile for uplink ports (continued).

https://t.me/learningnets

Create Uplink Profiles for VM Data

- Create a new port profile
- Configure the VMware port group
- Configure the profile to allow VM data VLANs across the trunk
- Enable ports for the group
- Designate the profile as an uplink
- Enable the profile

```
Pod1VSM(config)# port-profile vmData-uplink
Pod1VSM(config-port-prof)# switchport mode trunk
Pod1VSM(config-port-prof)# vmware port-group vmData-uplink
Pod1VSM(config-port-prof)# switchport trunk allowed VLAN 105-110
Pod1VSM(config-port-prof)# no shut
Pod1VSM(config-port-prof)# capability uplink
Pod1VSM(config-port-prof)# state enabled
```

© 2008 Cisco Systems, Inc. All rights reserved.

This figure describes creating a VM data uplink profiles.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Creating VM Data Port Profiles

This topic describes the configuration of VM data port profiles.

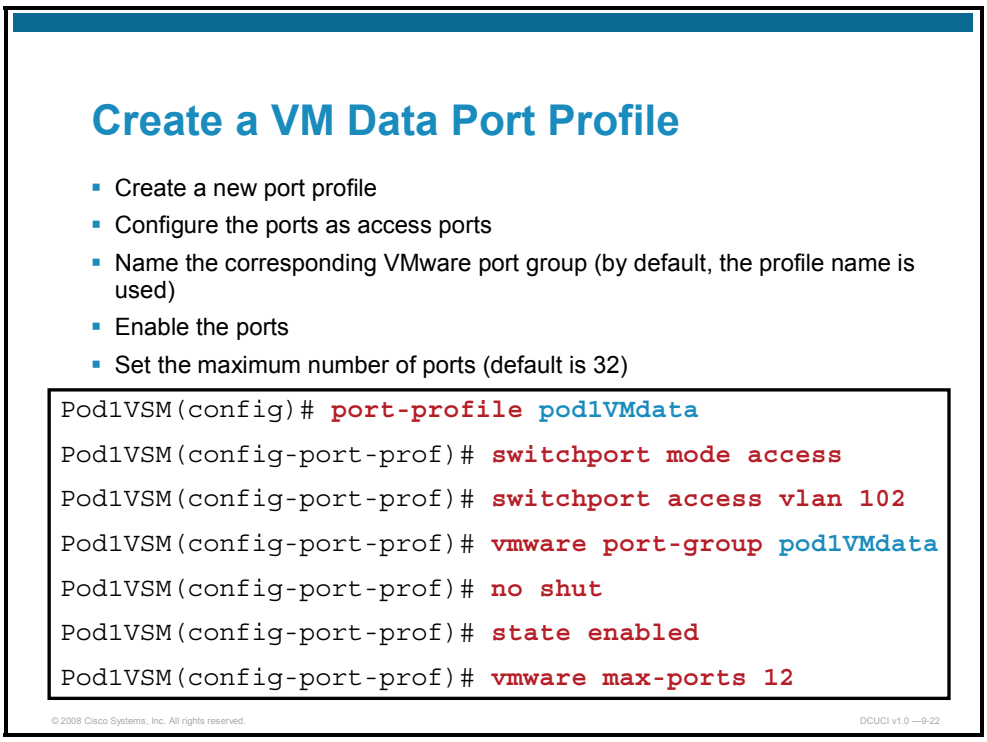

This figure describes the process of creating a VMware data port profile.

https://t.me/learningnets

Verify the New Port Group

You should observe output similar to the following:

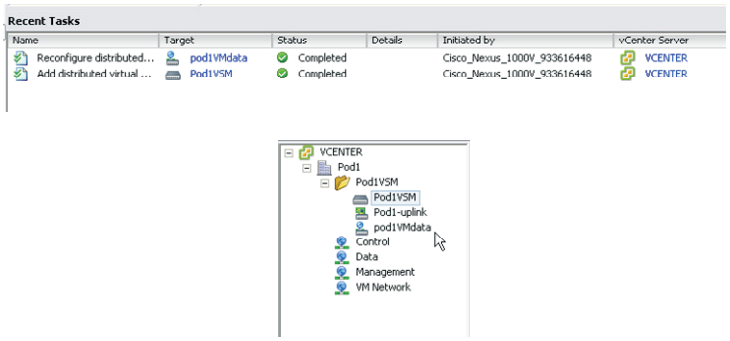

© 2008 Cisco Systems, Inc. All rights reserved.

This figure describes the process of verifying VMware data port profile configuration.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Configuring Cisco Nexus 1000V Port Channels

This topic describes Cisco Nexus 1000V port channel configuration.

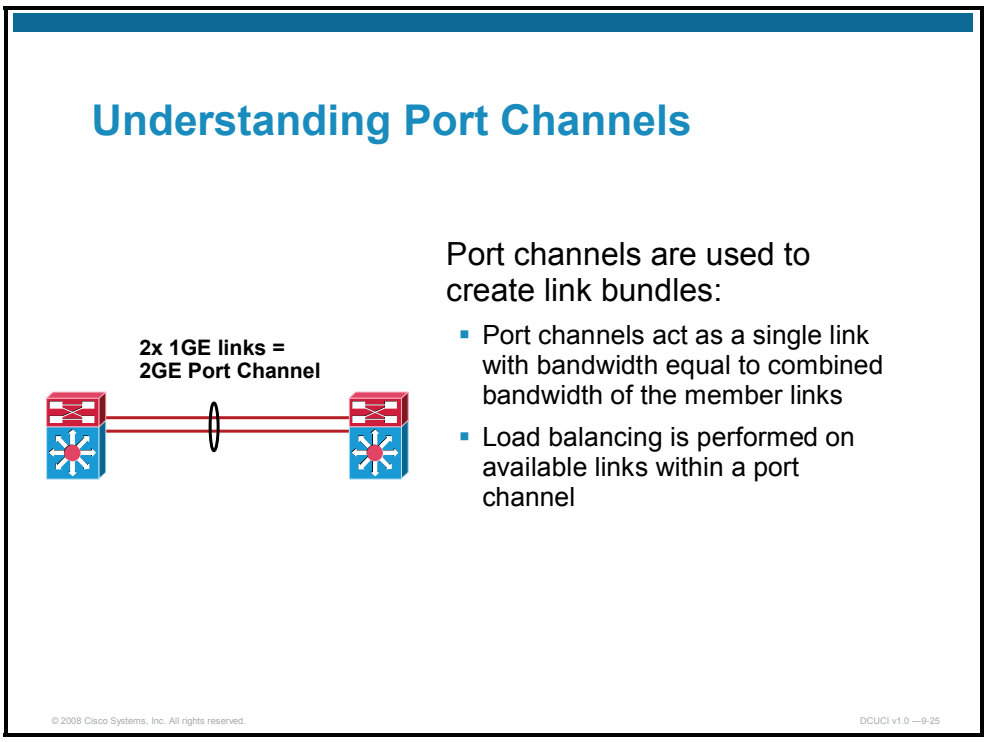

Port channels are used to bundle links to increase availability. The Cisco Nexus 1000V uses vmNICs to create port channels to upstream switches. These port channels allow VM data destined for the physical network to be load balanced across available links in an active fashion. Without port channels, NIC teaming from the Cisco Nexus 1000V is performed in an active-passive fashion.

https://t.me/learningnets

Port Channel Load Balancing

Load balancing is the method of selecting a physical link to transmit data within a logical port channel link.

Port channels can be load balanced using the following methods:

- Destination MAC
- Source MAC address
- Source and destination MAC
- Destination IP
- Source IP
- Source and destination IP
- Source TCP/UDP port number
- Destination TCP/UDP port number
- Source and destination TCP/UDP port number

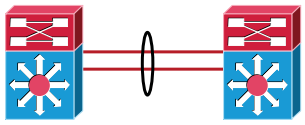

© 2008 Cisco Systems, Inc. All rights reserved.

Port channels create a single logical link comprised of multiple physical links. Once a port channel has been selected as an egress interface, the switch must choose a physical link within the port channel to transmit data on. The process of selecting a physical link is based on load balancing algorithms that are intended to provide fairness among the links. The Cisco Nexus 1000V supports several load balance algorithms, the most granular of which is *source* and *destination* TCP/UDP port number. This means that using source and destination port number will provided the greatest fairness when distributing traffic load across physical links.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Port Channel Guidelines

- Port channels across VEMs are not supported
- Port channels can be formed with multiple upstream links only when they satisfy the compatibility requirements, and under the following conditions:
  - The uplinks from the host are going to same upstream switch
  - The uplinks from the host are going to multiple upstream switches but the load balancing algorithm is source MAC-based
- Port channels can be configured using a port profile
- You can configure multiple port channels on a device

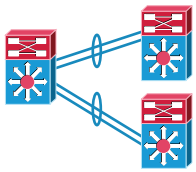

© 2008 Cisco Systems, Inc. All rights reserved.

The Cisco Nexus 1000V does not support port channels across multiple hosts. Each VEM on each host requires its own port channel. These port channels can exist between a single upstream switch or two upstream switches. Multiple port channels are supported on a single VEM or VSM.

https://t.me/learningnets

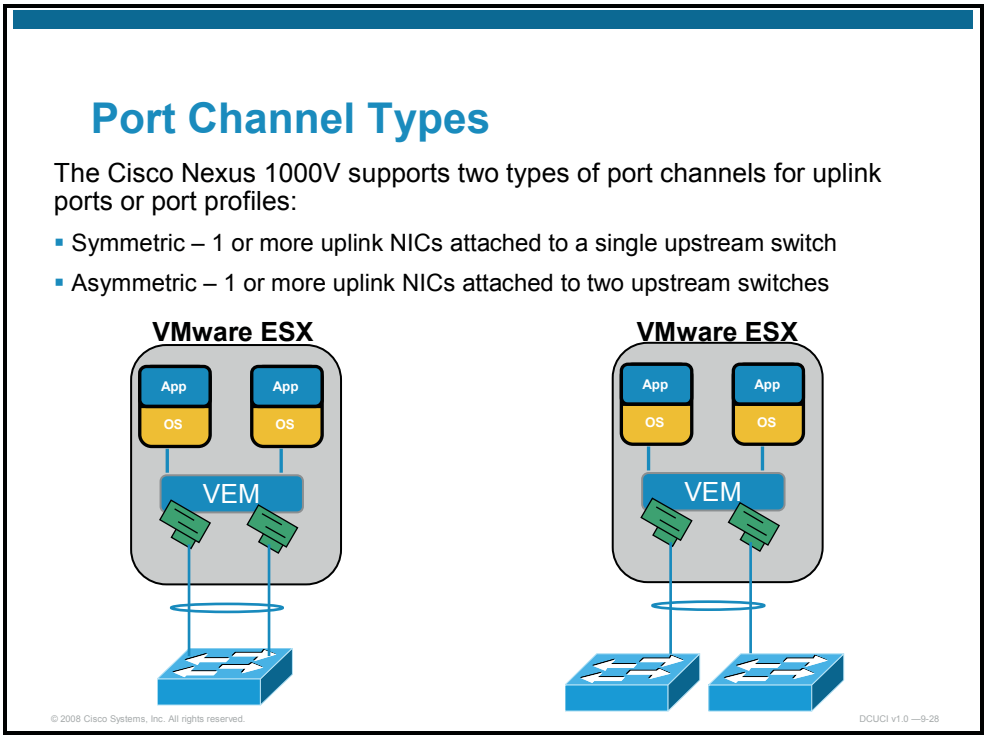

The Cisco Nexus 1000V supports two types of port channels.

Symmetric

This is the most common type of port channel and requires all members of the port channel to belong to the same device on each side of the port channel. In this configuration, all vmNICs that are members of the port channel must be cabled to the same upstream switch.

Asymmetric

In asymmetric mode, vmNIC port channel members can be cabled to two upstream switches for redundancy. This mode has the following requirements:

- The uplinks from the host are going to multiple upstream switches but the load balancing algorithm is source MAC-based.
- The upstream switch is capable of, and has CDP enabled
- The Cisco Nexus 1000V port channel is configured as CDP enabled

CDP information allows the Cisco Nexus 1000V to address issues that would typically arise in asymmetric configurations:

- MAC address flapping:
  - Prevented by utilizing MAC based load balancing.
- Duplicate packets received during floods, broadcasts, and multicasts:
  - The Cisco Nexus 1000V checks for a match between the subgroup ID where a packet was received and the destination subgroup ID, and only accepts those that match.
- Packets sent back to the host during floods and broadcasts:
  - The Cisco Nexus 1000V verifies that ingress packets are not the same packets that it sent out.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Symmetric Port Channel Configuration

1. Create the port profile.
2. Configure port profile parameters.
3. Turn on **channel-group auto.**

```
n1000v# conf
n1000v(config)# port-profile UplinkChannel
n1000v(config-port-prof)# channel-group auto mode on
n1000v(config-port-prof)# state enabled
n1000v(config-port-prof)# capability uplink
n1000v(config-port-prof)# switchport mode trunk
n1000v(config-port-prof)# description Upstream Port-Channel
```

After the first physical port has been added to the port profile, a channel group will be automatically created.

Separate port channels will be created using assigned vmNICs on each VEM.

© 2008 Cisco Systems, Inc. All rights reserved.

It is recommended that port channels be configured using port profiles rather than at the individual port level. This will ensure consistent configuration across hosts and ensure compatibility for advanced functionality such as: vMotion, high availability, and DRS.

For symmetric port channels, enter the **channel-group auto mode on** command under port profile configuration. When vmNICs (physical NICs on VMware ESX Hosts) are added to this profile, port channels will automatically be created on each host and additional vmNICs will be added to the channel. This is due to the fact that a port channel can not span multiple hosts, so an individual port channel will be created from the VEM of each VMware ESX Host.

https://t.me/learningnets

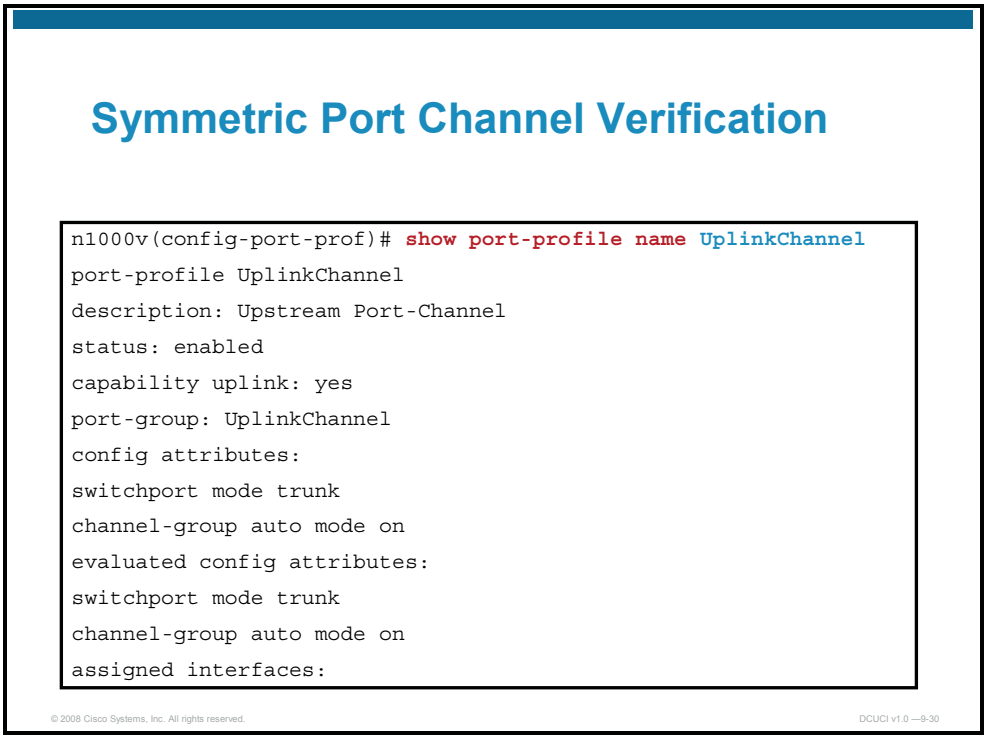

The figure illustrates the verification of the uplink port profile port channel configuration.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

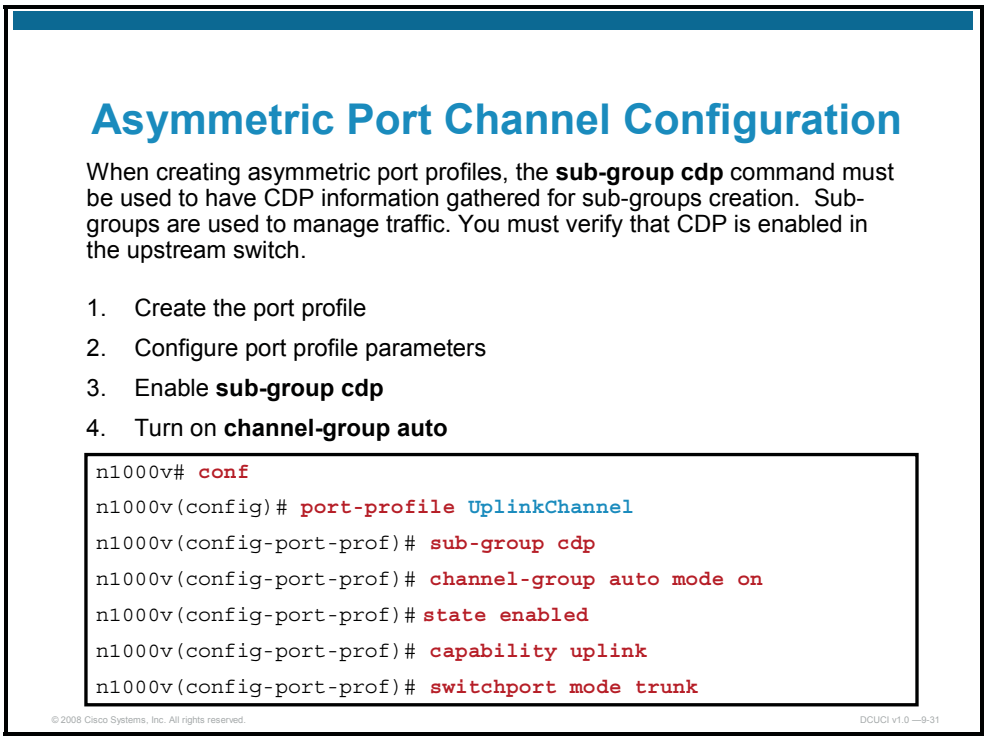

Asymmetric port channels rely on source Mac address load balancing and sub groups to prevent MAC address instability and multiple frame copies on upstream switches. Loops are automatically prevented by not forwarding frames received on uplink ports to other uplink ports.

- **MAC table instability:** In asymmetric mode, the Cisco Nexus 1000V relies on source MAC hashing to avoid MAC address instability. Source MAC hashing ensures that traffic from a source VM will always utilize the same uplink port in asymmetric configurations. This prevents upstream switches from having issues due to a source MAC being seen on multiple switches and having to make constant MAC table changes.
- **Sub-groups:** Utilizing CDP, the Cisco Nexus 1000V is able to analyze traffic and create sub-groups for the separate upstream switches in an asymmetric port channel. This allows the Cisco Nexus 1000V to prevent multiple copies of the same frame received on separate ports to be forwarded to a VM.

https://t.me/learningnets

Adding Hosts to the VSM

This topic describes the process of adding hosts to a VSM.

Add a Host

This procedure assumes that VMware Update Manager is being used or that the VEM has previously been installed on the host.

To manually install a VEM on a VMware ESX Host, refer to the Cisco VEM Install Guide.

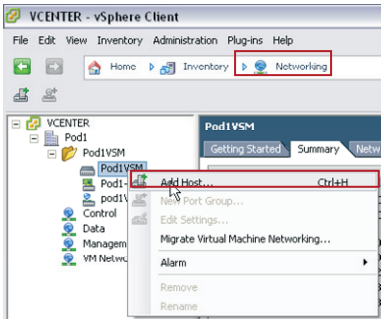

- Select the data center **Networking** view
- Right click the VSM
- Select **Add Host**

© 2008 Cisco Systems, Inc. All rights reserved.

Hosts are added to a VSM by the VMware administrator. The administrator uses the datacenter Networking view to assign hosts and their vmNICs to a Cisco Nexus 1000V controlled DVS.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

Add a Host (Cont.)

- Select the host to add
- Select the uplink NIC(s) to add
- Click **Next**

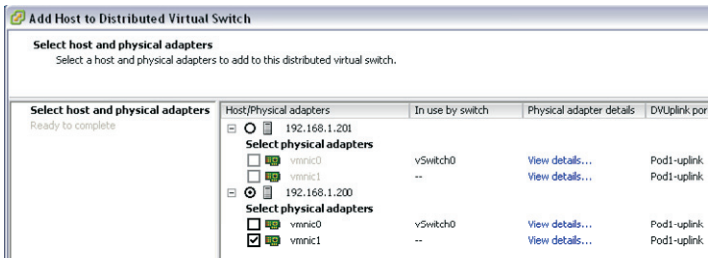

© 2008 Cisco Systems, Inc. All rights reserved.

The figure describes how to add a host to a VSM.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Add a Host (Cont.)

- Verify the configuration
- Click **Finish**
- Repeat as necessary for all hosts

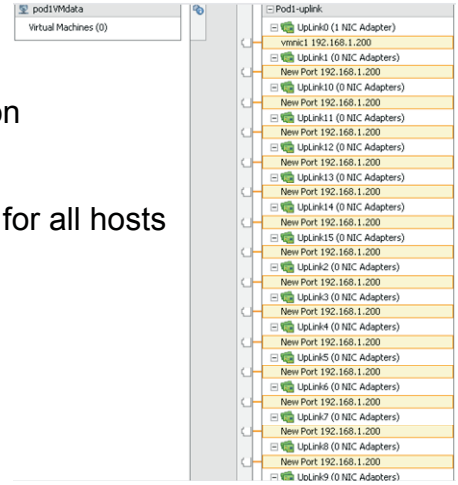

© 2008 Cisco Systems, Inc. All rights reserved.

The figure describes how to add a host to a VSM.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

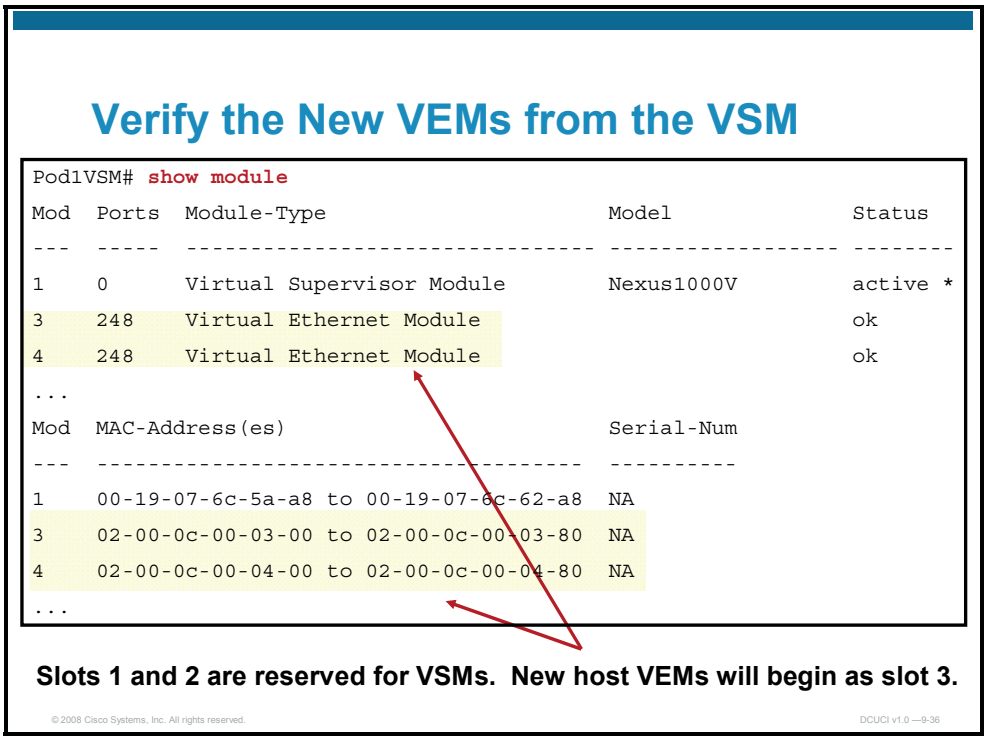

Once a host has been added and the VEM has been successfully installed within the host kerne,l it will appear as a module on the VSM CLI. This is similar to modules being added to a physical chassis. Slot 1 and 2 are reserved for VSMs.

https://t.me/learningnets

Verify the New VEMs from the VSM (Cont.)

The **show module vem map** command shows the status of all VEMs as well as the Universally Unique Identifier (UUID) of the host the VEM runs on.

```
Pod1VSM# 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 that the VEM resides on.

© 2008 Cisco Systems, Inc. All rights reserved.

The figure shows how to verify newly added VMware hosts.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Adding VMs to a VSM

This topic describes how to add VMs to a VSM.

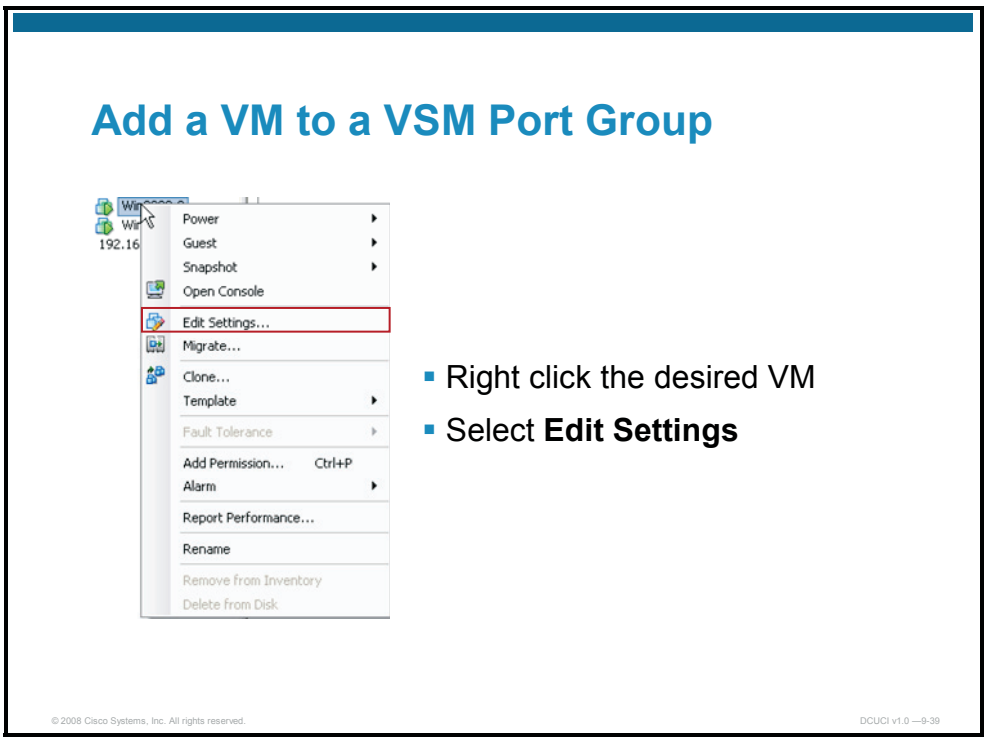

This figure shows the addition of VMs to a VSM port profile.

https://t.me/learningnets

Add a VM to a VSM Port Group (Cont.)

- Select the NIC to configure
- Change the **Network Label** to the desired VSM port profile

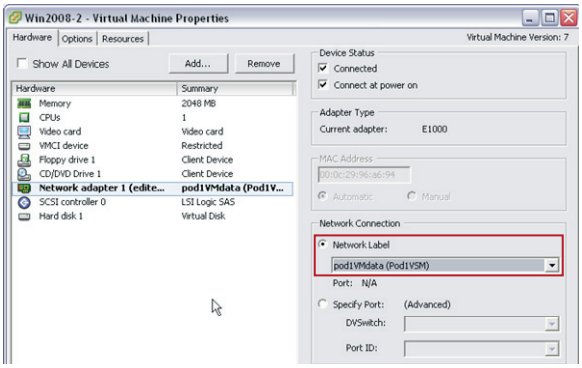

© 2008 Cisco Systems, Inc. All rights reserved.

VMs are added to Cisco Nexus 1000V controlled DVSs by the VMware administrator. After a port profile has been created and the corresponding port group has been pushed to vCenter, the port group will be available for selection as a network label within VM configuration.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

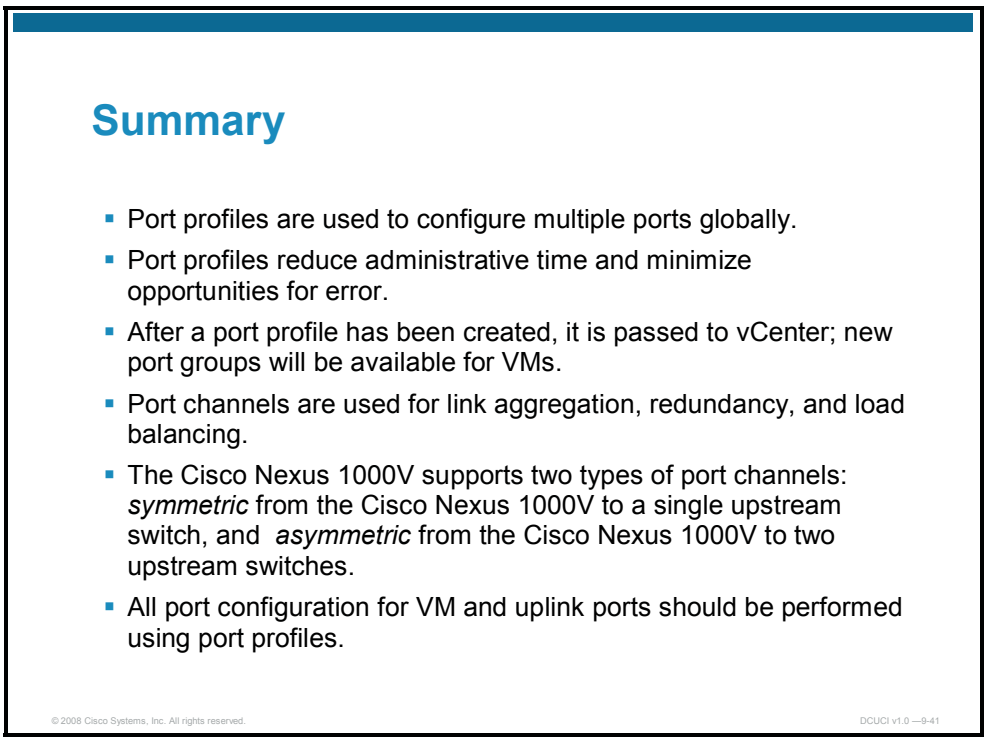

https://t.me/learningnets

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

Module Summary

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

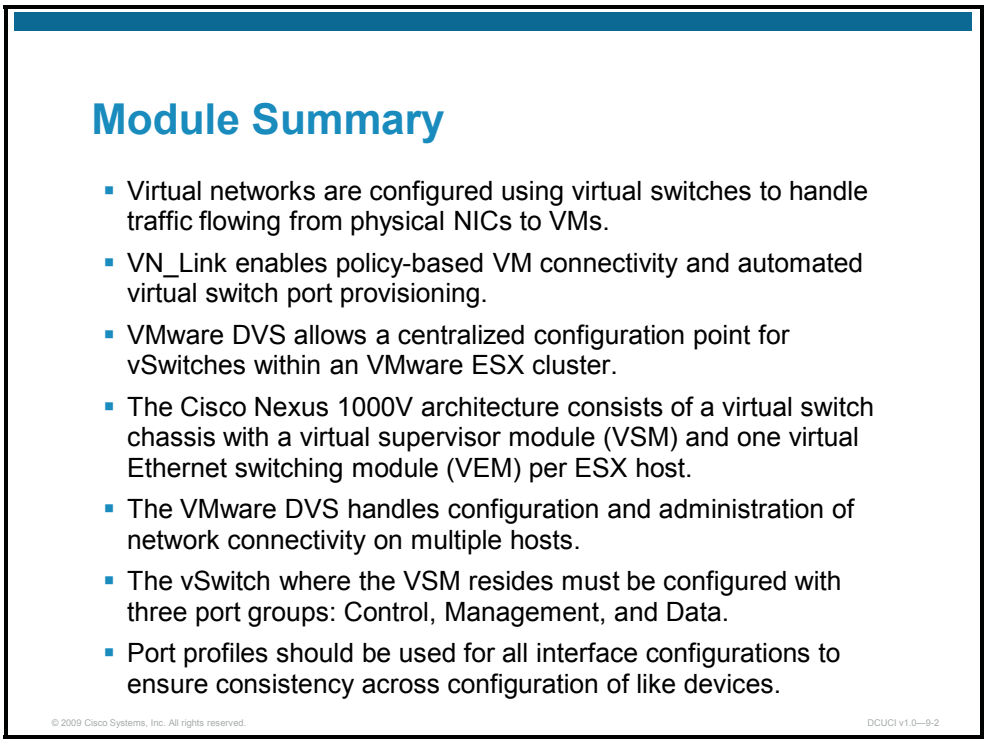

https://t.me/learningnets

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

Q1) Which of the following is a feature of the Nexus 1000V? (Source: Introduction to Nexus 1000V)

A) VM-level security and policy provisioning
B) Hardware-based virtual interfaces
C) Per-VLAN Spanning Tree (PVST)
D) Ethernet Host Virtualizer

Q2) Which VMware feature enables automated migration of virtual machines between physical hosts in response to defined thresholds? (Source: Server Virtualization Primer)

A) Distributed Resource Scheduler
B) VMotion
C) High Availability
D) vSphere Client

Q3) In traditional VMware deployments, virtual switches are configured by which group? (Source: VMware Ethernet Networking)

A) Network Administrators
B) Operating System Administrators
C) VMware Administrators
D) Security Administrators

Q4) Which property can a Cisco Nexus 1000V Port Profile define? (Source: Nexus 1000V Architecture)

A) OSPF and EIGRP routes
B) VSAN membership
C) CoS and DSCP markings
D) EtherChannel membership

Q5) Which feature did VMware's vNetwork Distributed Switch (DVS) introduce that did not exist in traditional virtual switches? (Source: Configure VMware Switching)

A) Egress Rate Limiting
B) Ingress Rate Limiting
C) VLAN Trunking
D) Port Groups

Q6) Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) can be configured for which type of high availability? (Source Installing and Configuring the Nexus 1000V)

A) Active/Active
B) Primary/Subordinate
C) Primary/Standby
D) Active/Standby

https://t.me/learningnets

Q7) What are the two types of Cisco Nexus 1000V Port Profiles? (Choose two.) (Source: Configuring Port Profiles)

A) VM profiles and uplink profiles
B) ACL profiles and QoS profiles
C) LAN profiles and SAN profiles
D) Forwarding profiles and host profiles

© 2009 Firefly Communications, LLC
https://t.me/learningnets

Module Self-Check Answer Key

Q1) A

Q2) A

Q3) C

Q4) C

Q5) B

Q6) D

Q7) A

https://t.me/learningnets

© 2009 Firefly Communications, LLC
https://t.me/learningnets

Appendix

Data Center Unified Computing Implementation

Overview

This appendix provides additional information on the command line interfaces (CLIs) supported by the Cisco Unified Computing System (Cisco UCS), and defines many acronyms used within the course.

Module Objectives

Upon completing this module, you will be able to provide additional information on command line interfaces and acronyms used within the course. This includes being able to meet these objectives:

- Explore the unified computing user interfaces

https://t.me/learningnets

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

Appendix

Exploring the Unified Computing User Interfaces

Overview

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

Objectives

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

- Describe the UCS command-line interfaces
- Describe the connect commands
- Describe various acronyms used within Cisco UCS

https://t.me/learningnets

Command-Line Interface

This topic describes the command-line interfaces available in the Cisco Unified Computing System (Cisco UCS).

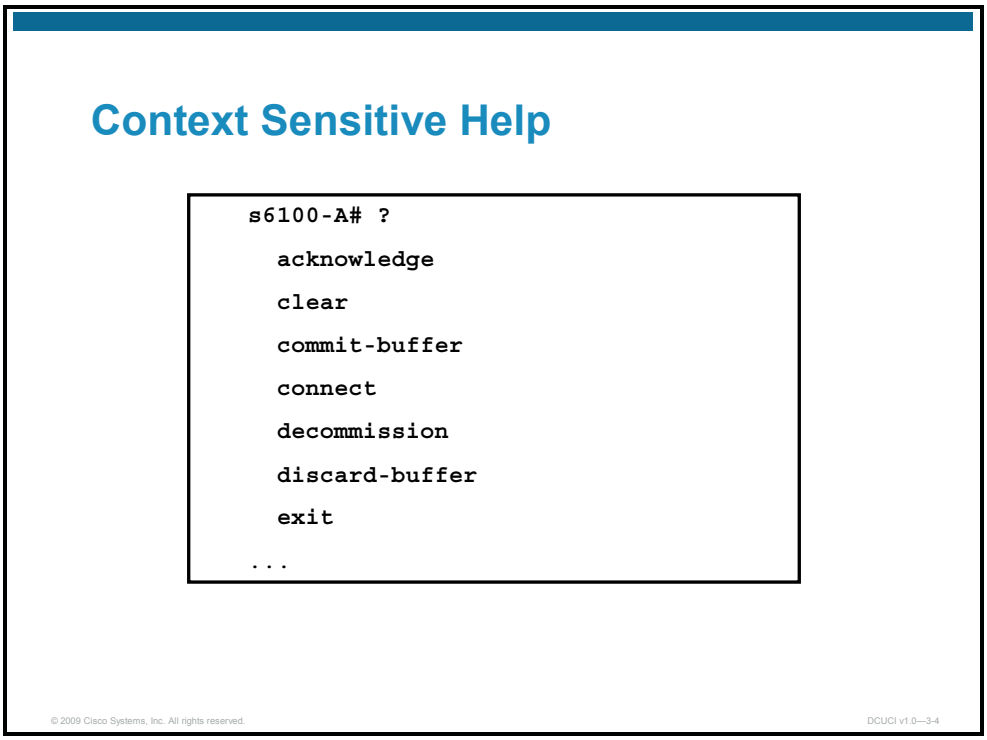

The question mark is familiar to most Cisco command-line interface (CLI) users. It shows all of the commands available within the present context. The figure shows some of the most common commands.

`UCS-6100-A# ?`

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

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

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

https://t.me/learningnets

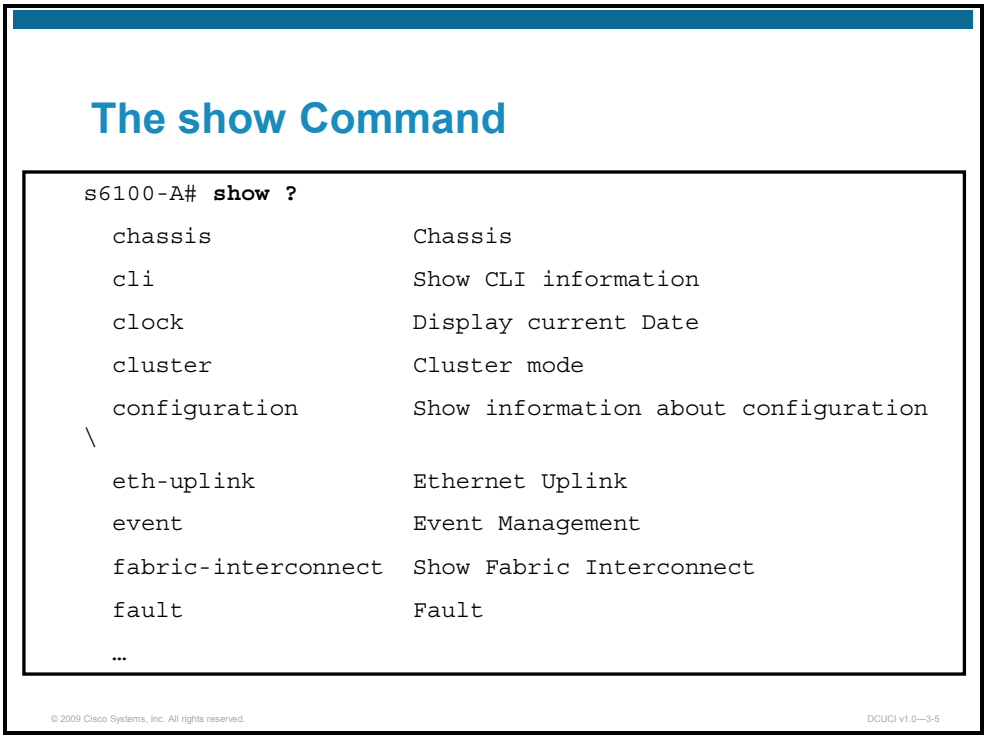

The **show** command, followed by a question mark, displays the attributes or parameters available for viewing.

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

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

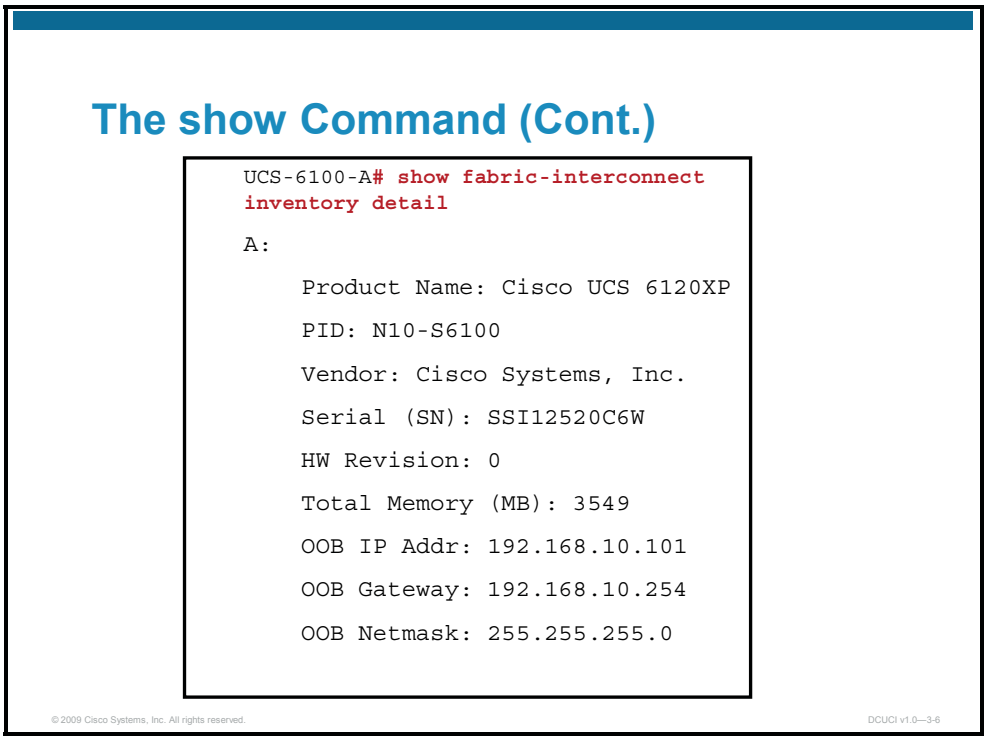

This example illustrates the **show** command with the arguments of *fabric-interconnect inventory detail*. The resulting display shows the inventory components of the switch. This is one example of the **show** command.

https://t.me/learningnets

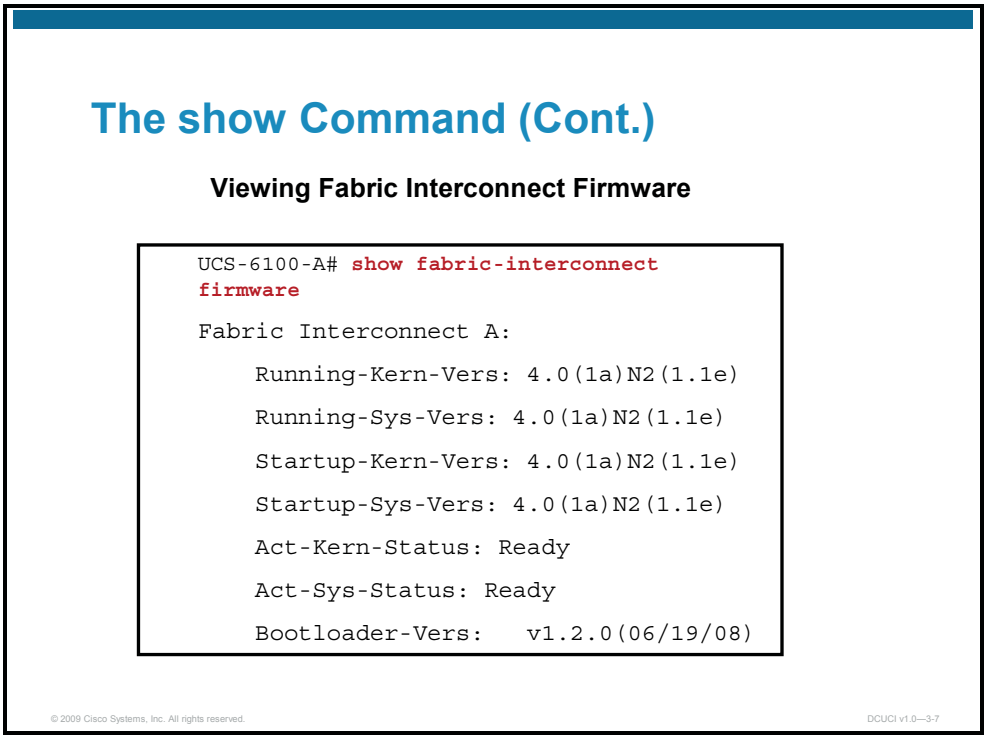

This figure shows the use of the **show** command to view switch firmware versions.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

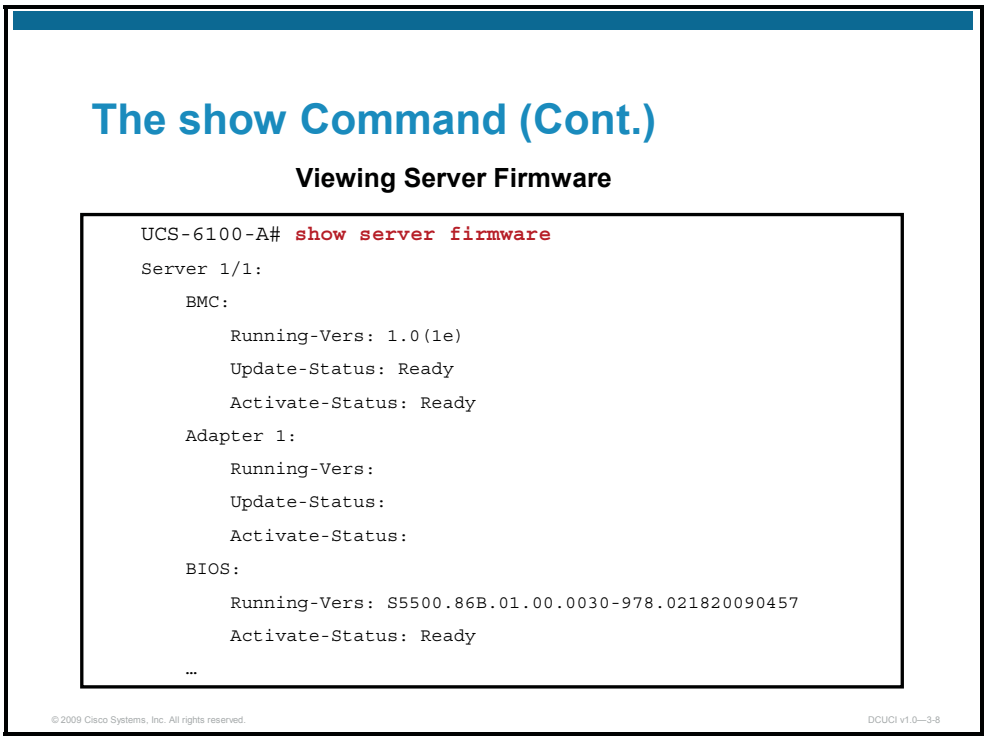

This figure shows sample output from the **show server firmware** command.

https://t.me/learningnets

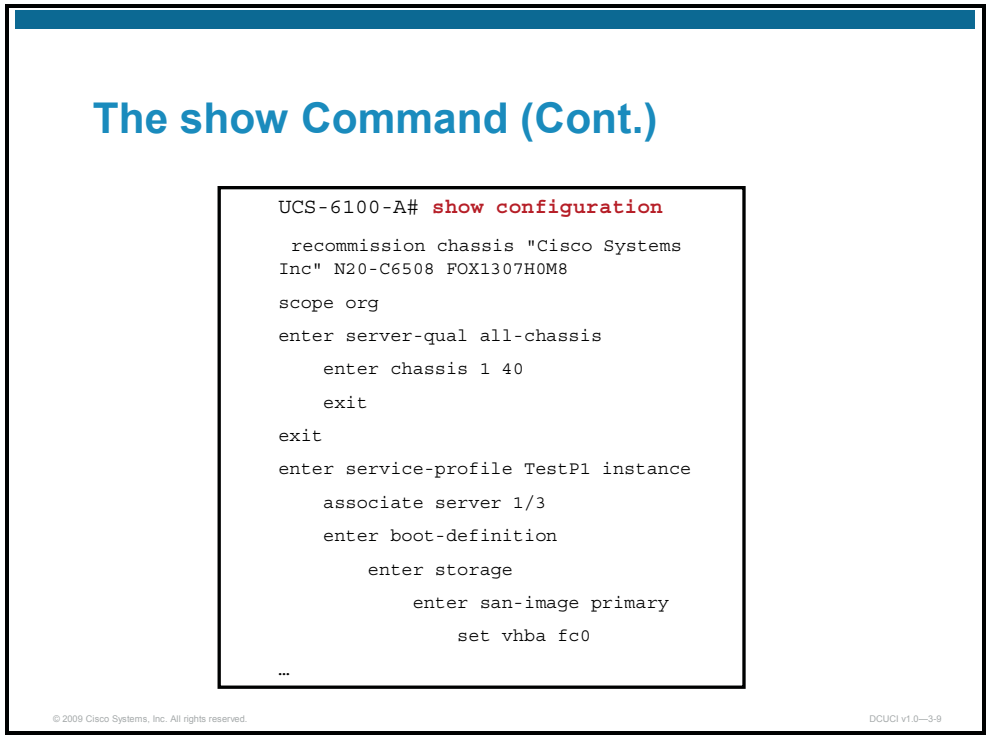

The **show configuration** command displays the commands to rebuild the configuration. Note that the **show configuration** command pulls its configuration information from the configuration database. This command, as well as other commands, can be run in scripted fashion using the **run-script** command.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

The scope Command

```
UCS-6100-A# scope ?
 adapter               Mezzanine Adapter
  chassis               Chassis
  eth-server            Ethernet Server Domain
  eth-uplink            Ethernet Uplink
  fabric-interconnect   Fabric Interconnect
  fc-uplink             FC Uplink
  firmware              Firmware
...
```

© 2009 Cisco Systems, Inc. All rights reserved.

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

https://t.me/learningnets

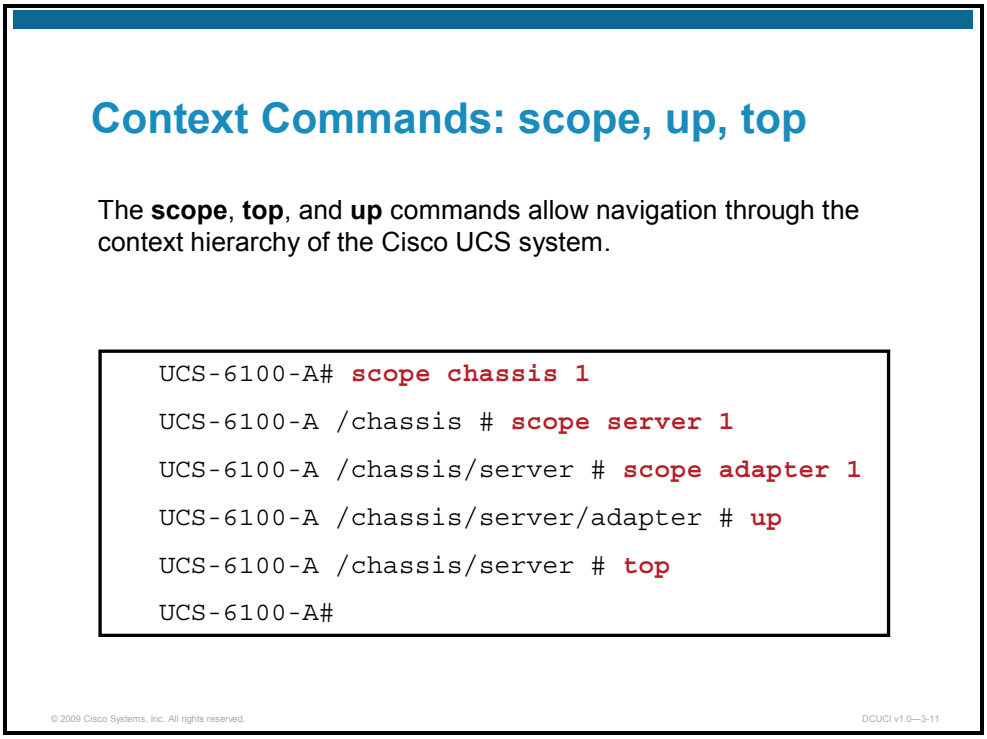

The figure shows the use of the **scope**, **up**, and **top** commands. Levels are traversed using the **scope** command. The **up** command is used to go up one level, and the **top** command is used to return to the root level.

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

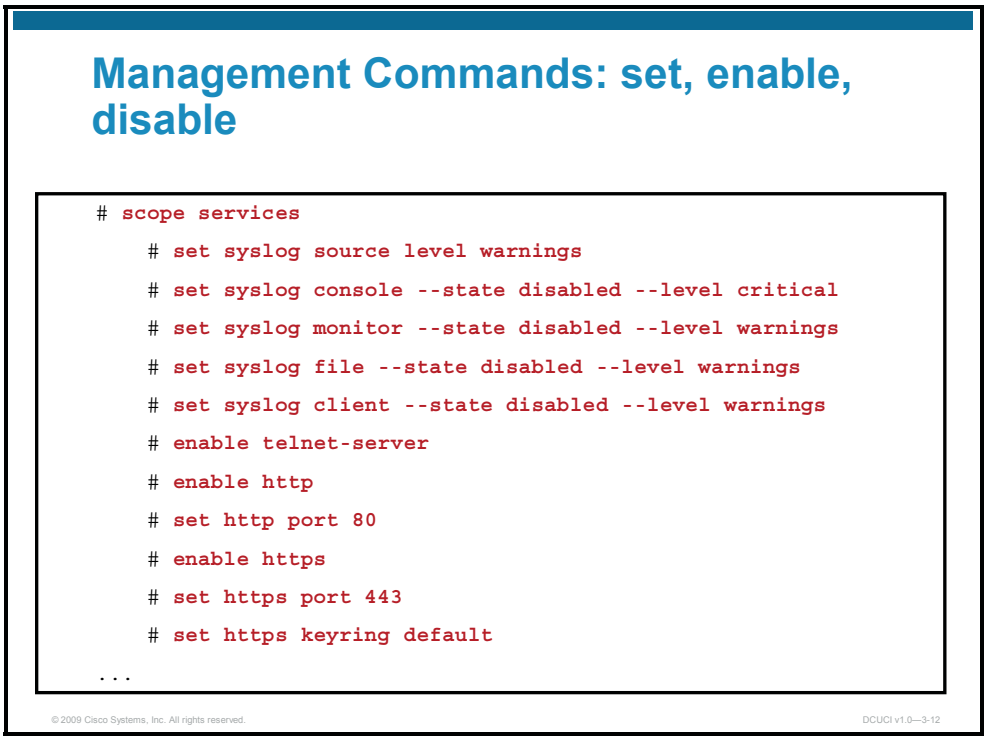

Other Management Commands

This example shows the available system scope services for enabling or disabling services and ports using **enable**, **disable**, and **set** commands.

https://t.me/learningnets

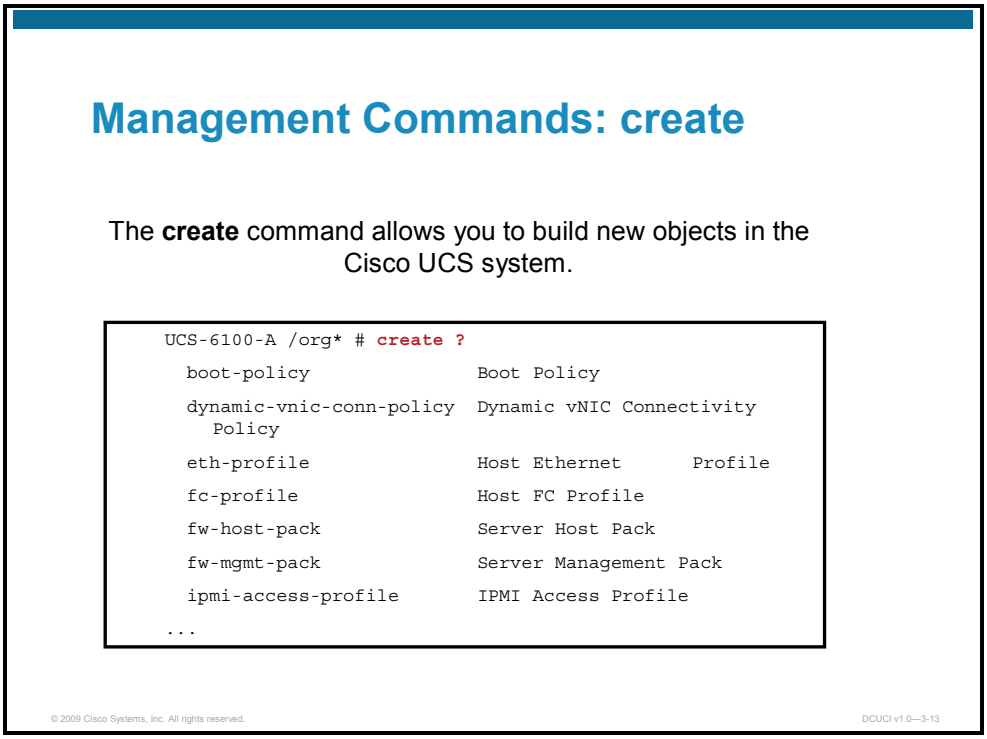

This figure shows the options available with the **create** command.

https://t.me/learningnets

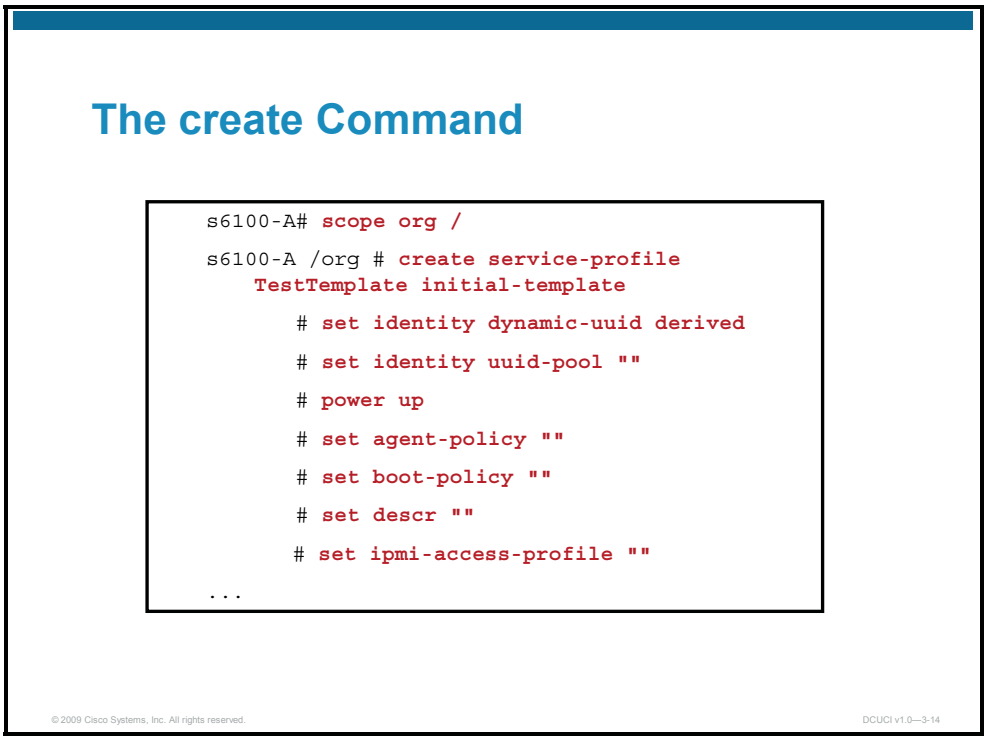

Use the **create** command to create new objects in the Cisco UCS database.

https://t.me/learningnets

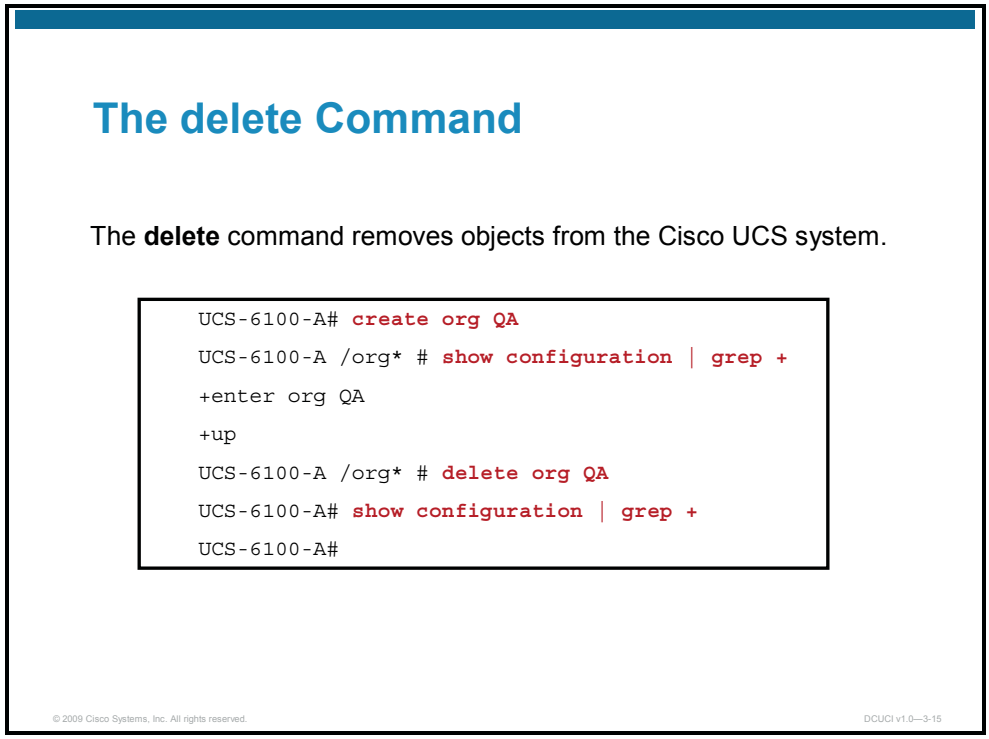

Use the **delete** command to delete an object that was created by the **create** command.

This example shows the creation and deletion of the *org* object. This example also shows how to display the object to be created in the command buffer.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

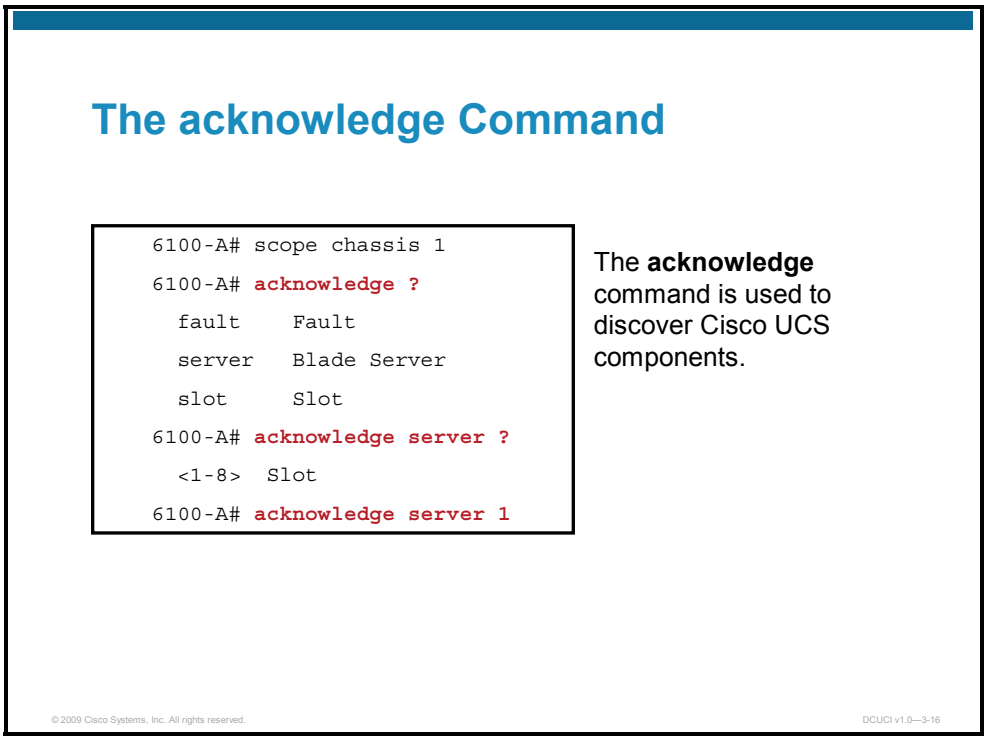

The **acknowledge** command is used to rediscover objects within the system. The individual actions associated with each object will vary depending on the object. The **acknowledge** command also reestablishes network connectivity.

https://t.me/learningnets

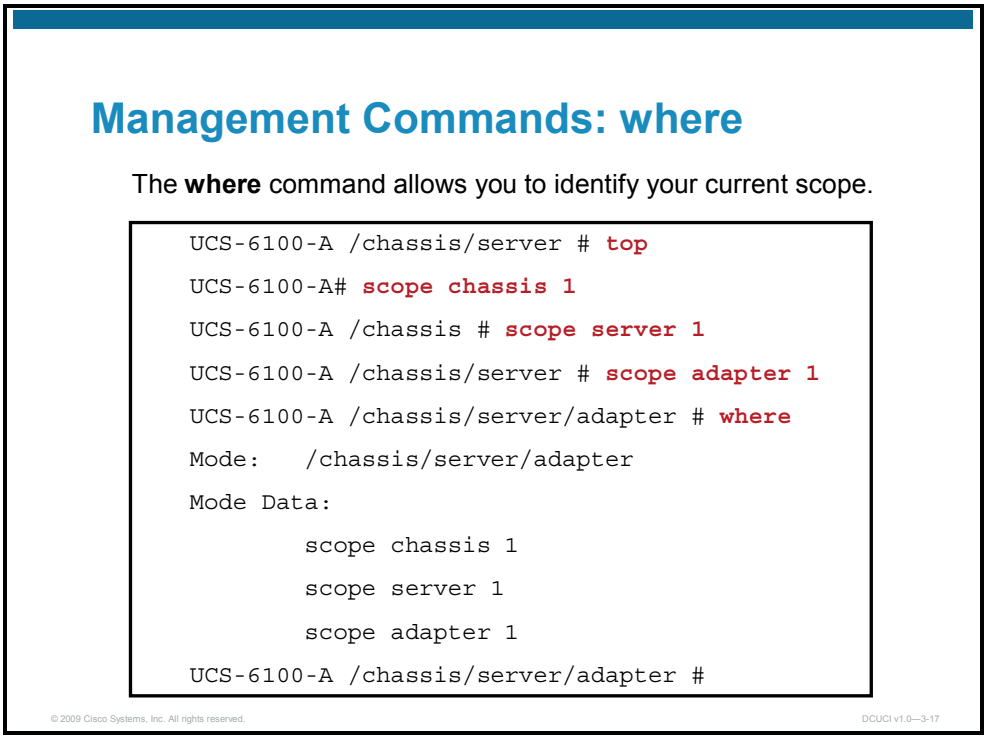

The **where** command shows the current scope and all scopes associated with the current view. The **where** command also displays the mode (the exec mode in this case).

https://t.me/learningnets

The connect Commands

This topic describes the **connect** commands.

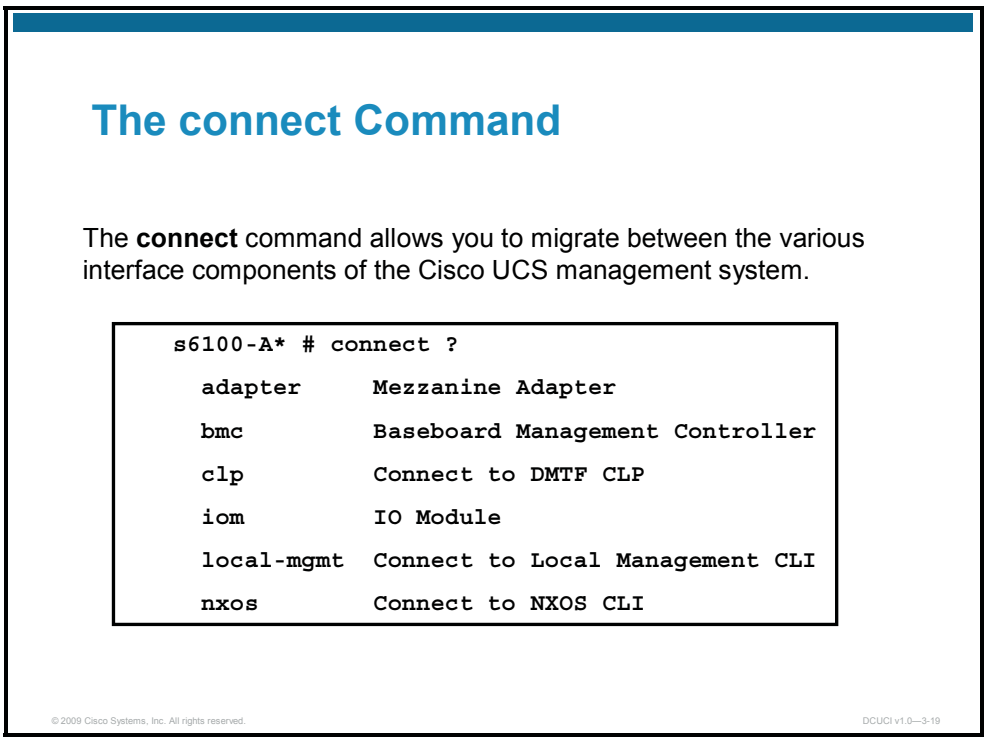

When connecting to the Cisco UCS system with Telnet or Secure Shell (SSH), you are placed in the default UCSM shell. Enter **connect ?** to display connect options.

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

https://t.me/learningnets

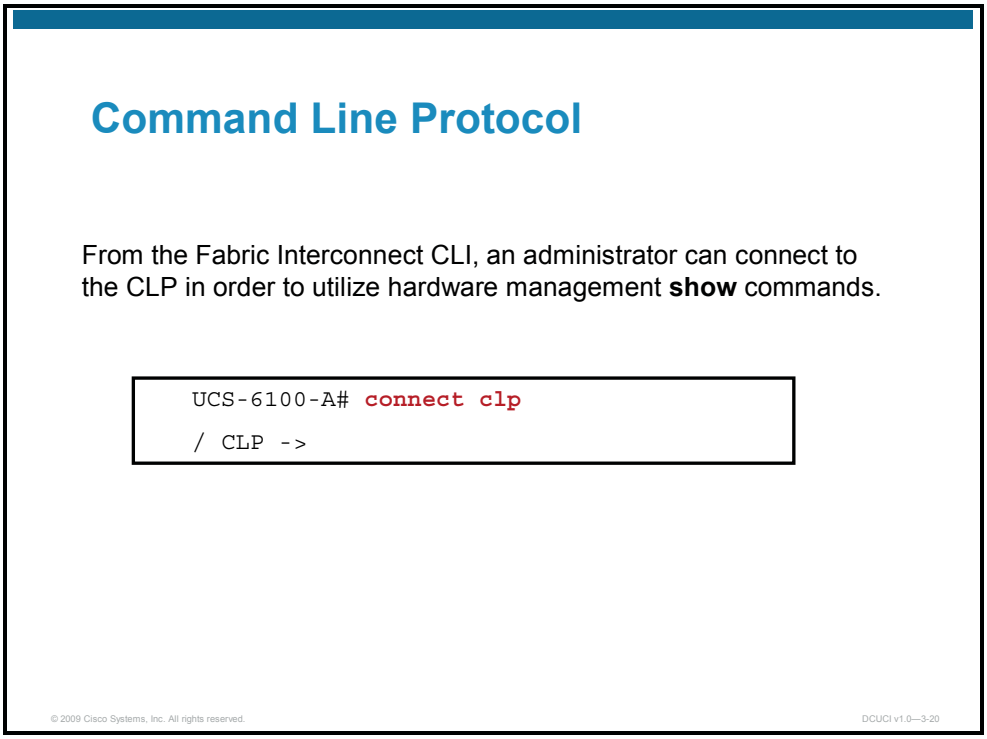

The Command Line Protocol (CLP) uses the smash CLP format. At present, it is mostly limited to **show** commands.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

Command Line Protocol (Cont.)

```
UCS-6100-A# connect clp
/ CLP -> show -d targets -l all
COMMAND COMPLETED - execution time: 00:00:22:607618
   Targets:
      hdwr1
      modular1
/admin1/hdwr1
   Targets:
      chassispkg1
/admin1/hdwr1/chassispkg1/fanmodpkg1
   Targets:
      fanpkg1
...
```

© 2009 Cisco Systems, Inc. All rights reserved.

This **show** command lists all targets that can show or display their status.

https://t.me/learningnets

Command Line Protocol (Cont.)

```
/ CLP -> show -d properties /admin1/modular1/tempsensor8
COMMAND COMPLETED - execution time: 00:00:00:51358
/admin1/modular1/tempsensor8
   Properties:
      SystemName = sys-machine/chassis-1
      SystemCreationClassName = Nuova:94-23817-1
      DeviceID = sys-machine/chassis-1/fan-module-1-
8/tempsensor
      CreationClassName = Nuova:97-23908-8
      PollingInterval = 0
      CurrentState = Normal
      PossibleStates = Normal,……..
```

© 2009 Cisco Systems, Inc. All rights reserved.

This **show** command shows the properties of a single target.

 © 2009 Cisco Systems, Inc.

https://t.me/learningnets

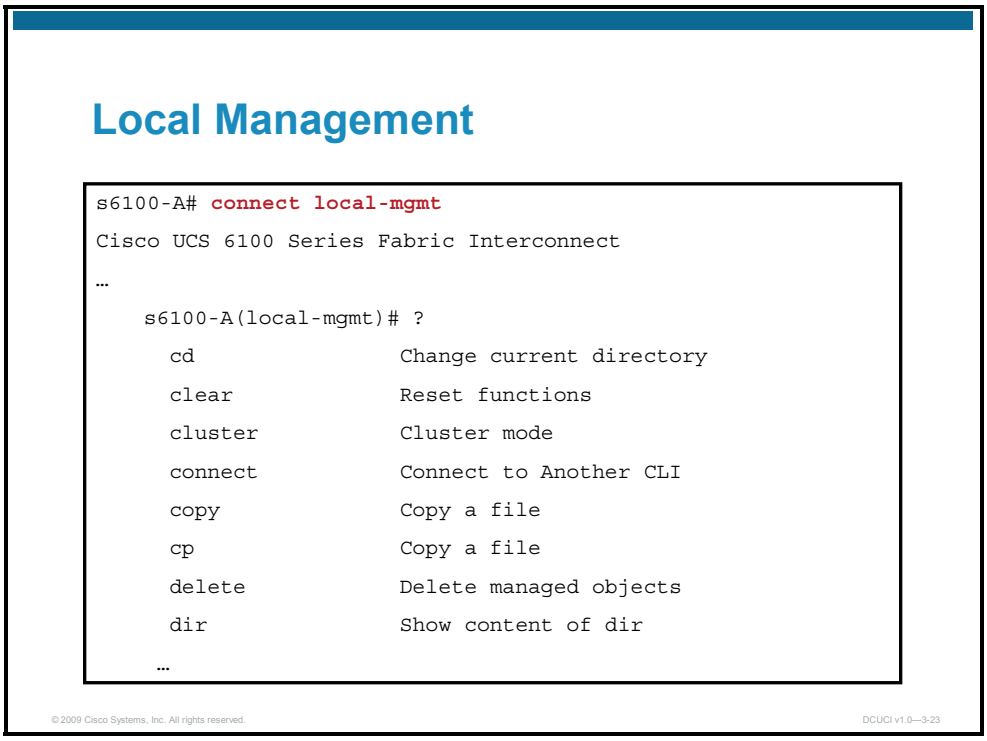

The figure shows the **connect local-mgt** command to attach to the local management CLI. The local management (local-mgmt) CLI contains the **reboot** and **cluster** management commands, as well as directory navigation and file manipulation commands.

https://t.me/learningnets

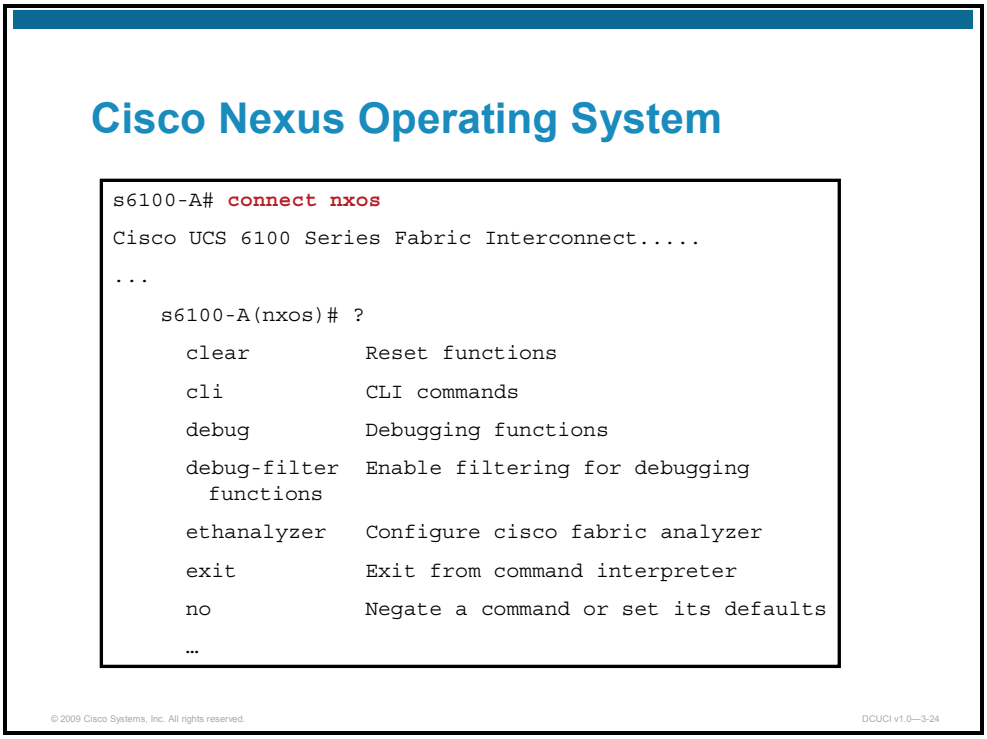

The figure shows a subset of commands available through the Cisco Nexus Operating System (Cisco NX-OS) CLI. An example of a command found in the Cisco NX-OS CLI is **ethanalyzer**, a Cisco fabric analyzer. Use **show ?** for an additional list of commands.

© 2009 Cisco Systems, Inc.

https://t.me/learningnets

UCS Acronyms

This topic describes various acronyms used within Cisco UCS.

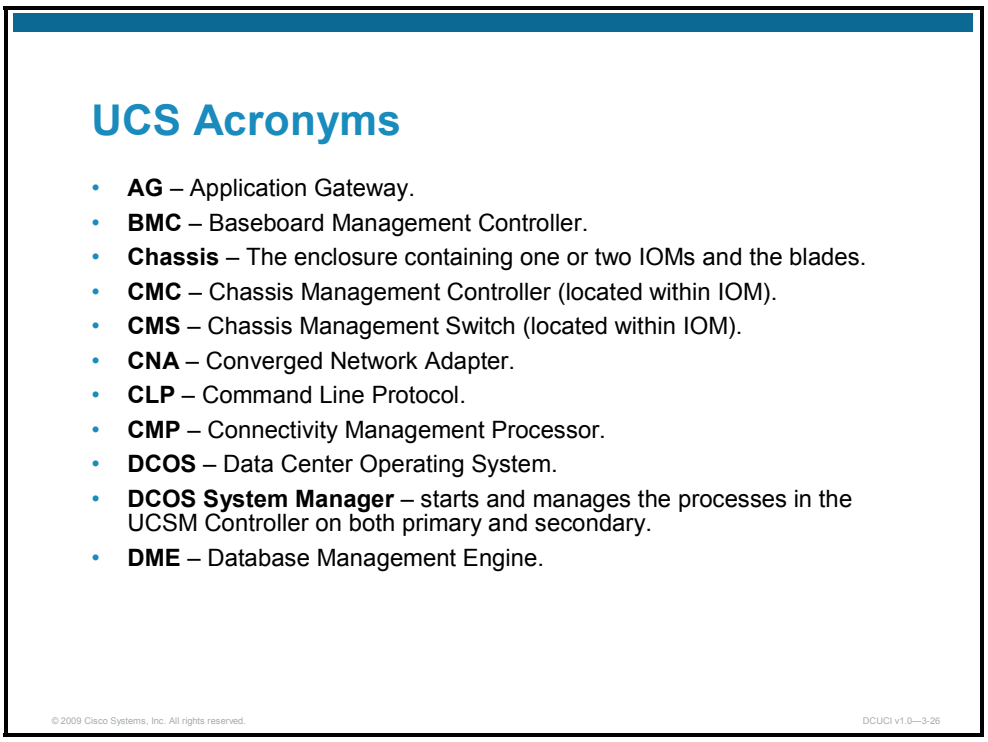

The figure lists various acronyms used within the course.

https://t.me/learningnets

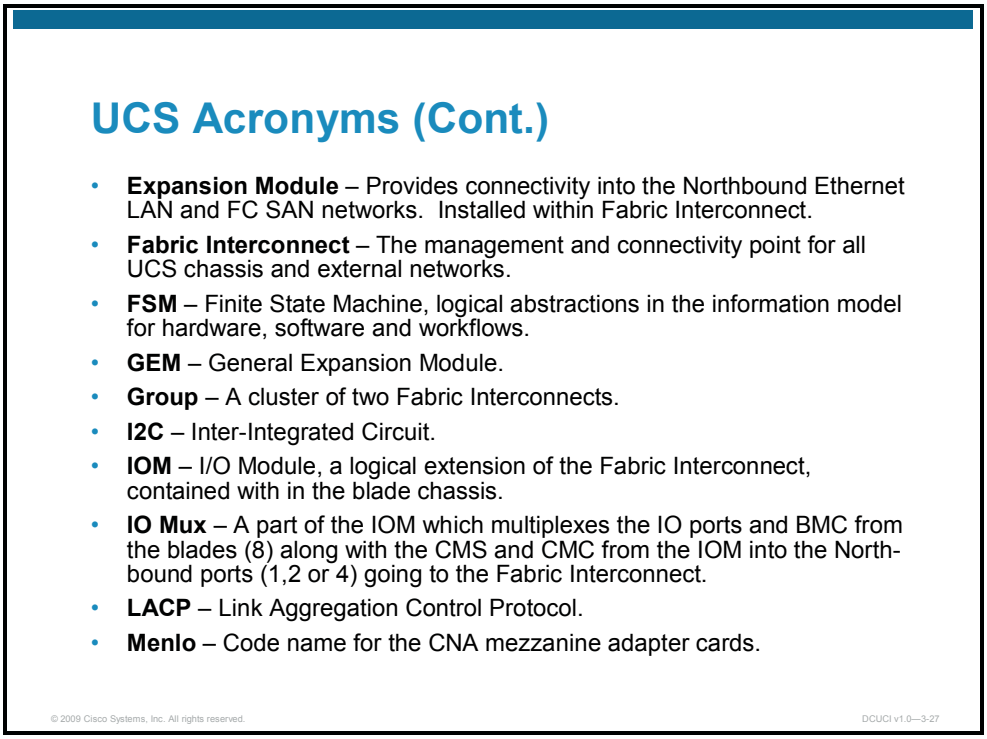

The figure lists various acronyms used within the course (continued).

https://t.me/learningnets

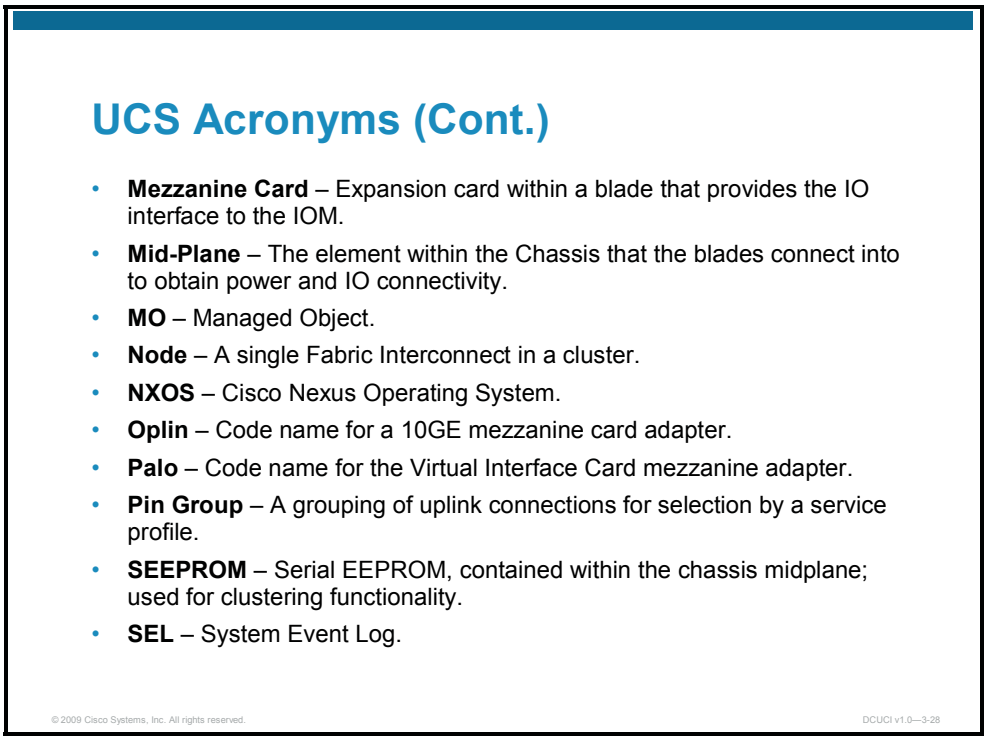

The figure lists various acronyms used within the course (continued).

https://t.me/learningnets

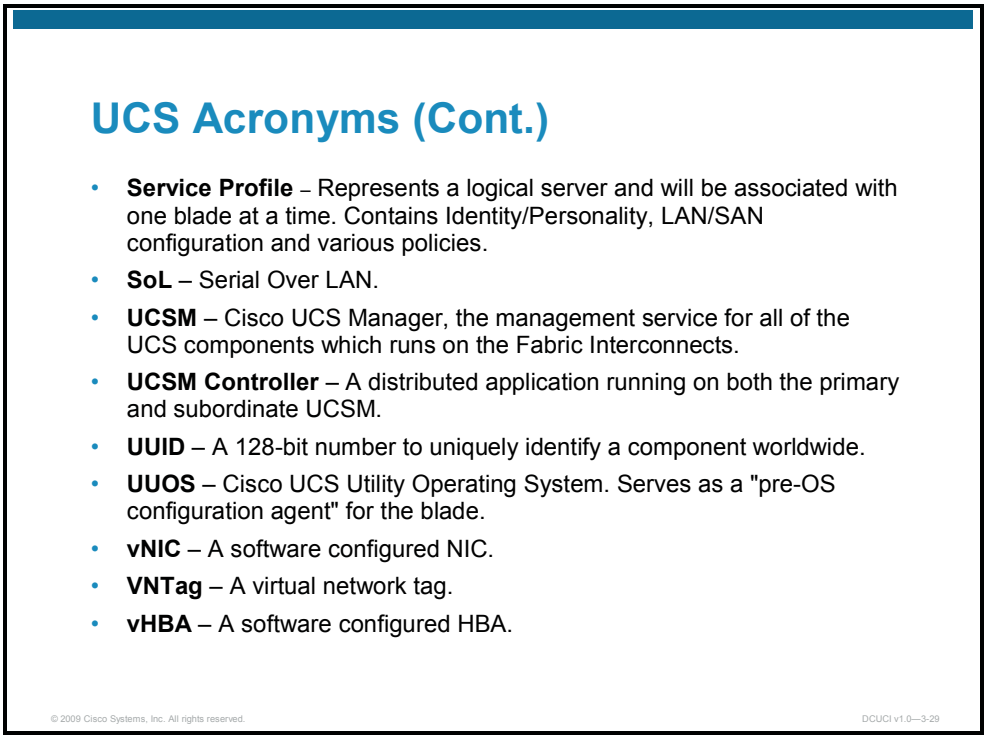

The figure lists various acronyms used within the course (continued).

https://t.me/learningnets

Summary

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

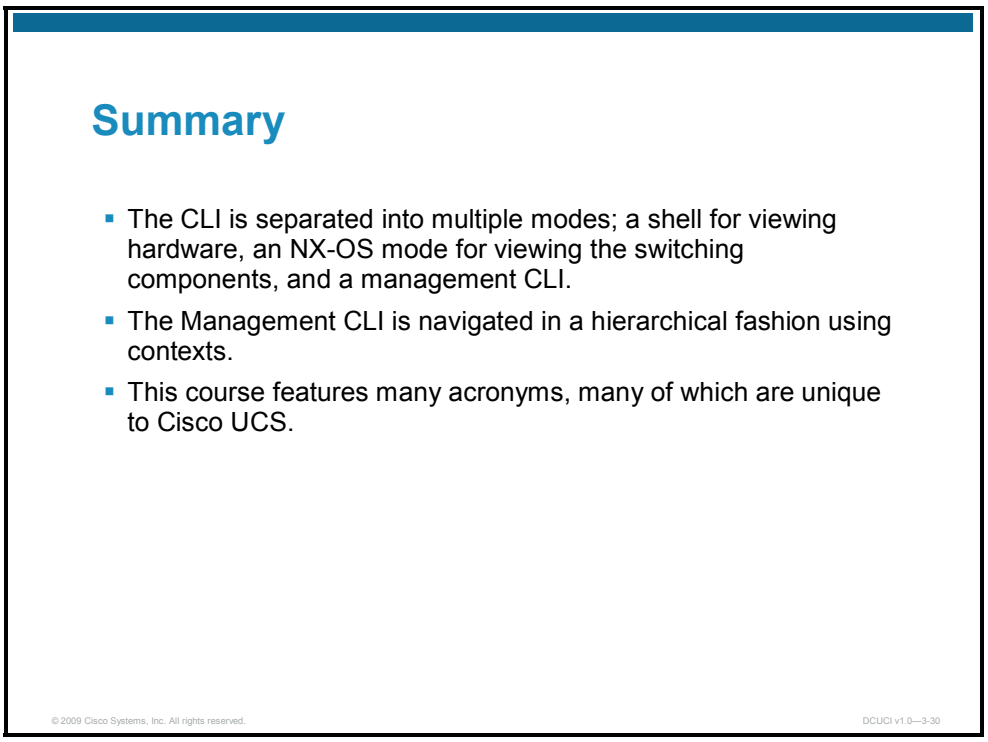

https://t.me/learningnets

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