DCUCI

Data Center Unified Computing Implementation

Volume 2

Version 4.0

Student Guide

Text Part Number: 97-3021-02

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Table of Contents

Volume 2

Module 6

Server Resources Implementation

Overview

Stateless computing and unified fabric are two of the cornerstone value propositions of Cisco Unified Computing System (UCS). Now that physical connectivity and administrative and operational procedures are in place, the focus of this module is logical connectivity. A blade server in the Cisco UCS B-Series is merely a compute node. To establish LAN and SAN connectivity, universal unique identifier (UUID), MAC Address, BIOS settings, and various other policies, a service profile must be created to contain all of these elements. The great benefit of abstracting policy settings and identities is that they are portable. If a blade server fails in such a way that the operating system or hypervisor can no longer operate, the service profile can simply be associated with the replacement blade. All of the elements of that server are transferred.

Module Objectives

Upon completing this module, you will be able to implement a stateless computing architecture. This ability includes being able to meet these objectives:

- Create identity and resource pools
- Create service profiles
- Create service profile templates and clone service profiles
- Manage service profiles

Lesson 1

Creating Identity and Resource Pools

Overview

Identity and resource pools are containers that facilitate consistent application of abstracted identities that are used by service profiles and service profile templates. While administrators are free to track abstracted identities in a spreadsheet or text file, it is difficult to scale in a large system.

Objectives

Upon completing this lesson, you will be able to configure identity and resource pools used by service profiles in service profile templates. This ability includes being able to meet these objectives:

- Explain the rationale for creating identity and resource pools
- Configure UUID pools
- Configure MAC pools
- Configure WWNN pools
- Configure WWPN pools
- Configure server pools
- Configure the processes to automate server pool membership based on a qualification policy
- Demonstrate an understanding of the importance of creating pools in the correct organization

Rationale for Creating Identity and Resource Pools

This topic discusses the rationale for identity and resource pools.

Pools for Mobile Service Profiles

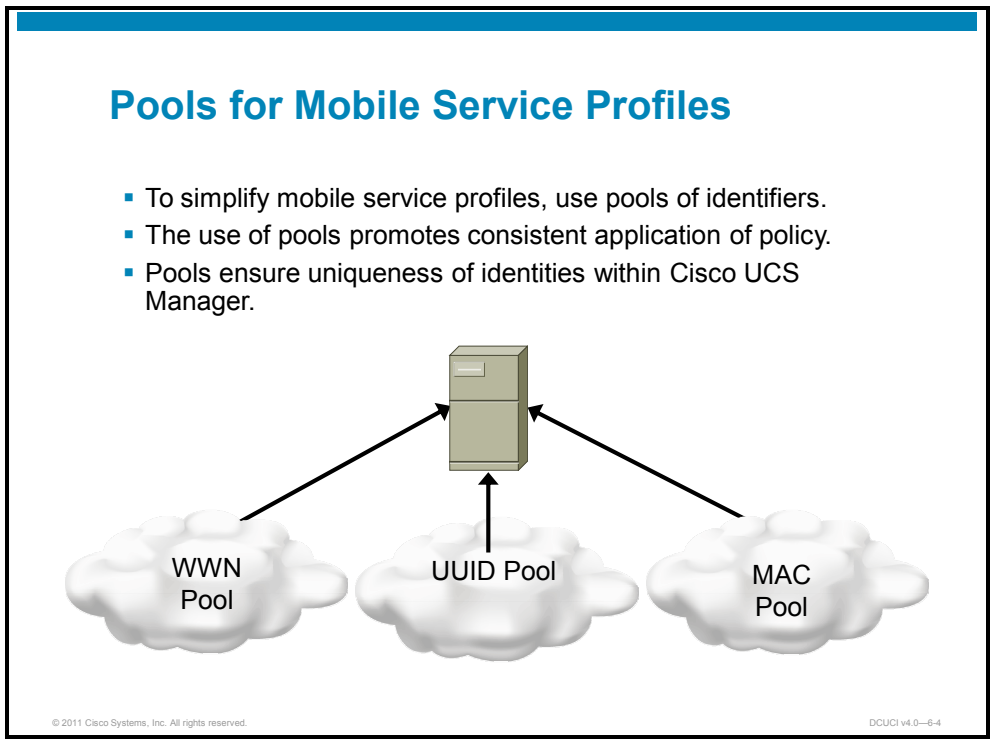

Stateless computing requires unique identity resources for universally unique identifiers (UUIDs), MAC addresses, and world wide names (WWNs) for Fibre Channel. Using pooled resources ensures consistent application of policy and reasonable assurances that identities are unique within the Cisco UCS Manager.

Logical Resource Pools

Logical Resource Pools

- Identity pools supply abstracted identities to service profiles in service profile templates.

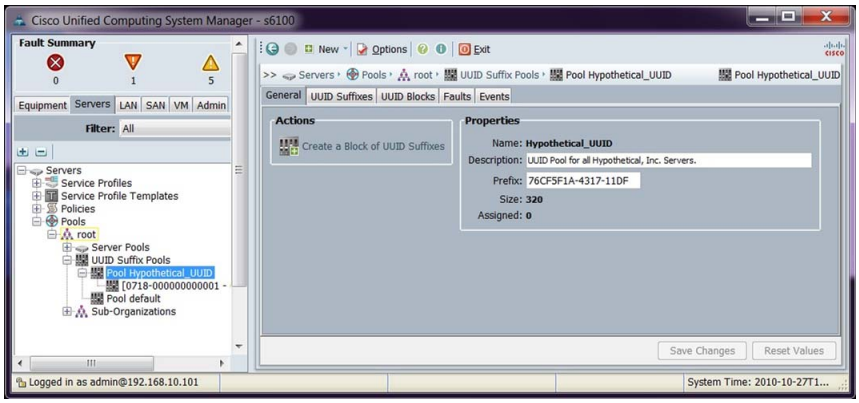

Logical resource pools provide abstracted identities that are used by service profiles in service profile templates to facilitate stateless computing.

Physical Resource Pools

Physical Resource Pools

- Physical resource pools provide blade server resources to service profiles to facilitate rapid provisioning.

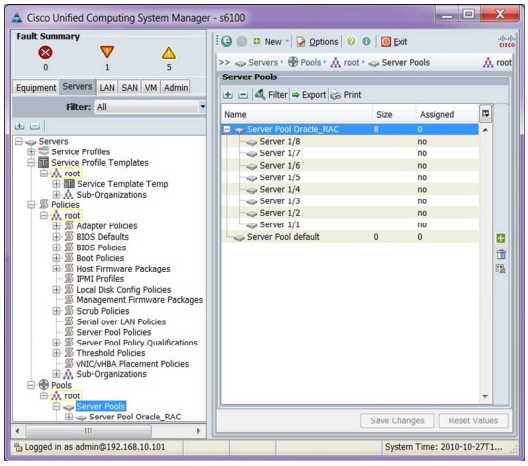

Physical resource pools are used to create groupings of blade servers that are based on arbitrary administrative criteria. These pools can be used with service profile templates to provide rapid provisioning of compute resources.

UUID Pools

This topic discusses the configuration of UUID pools.

UUID Use

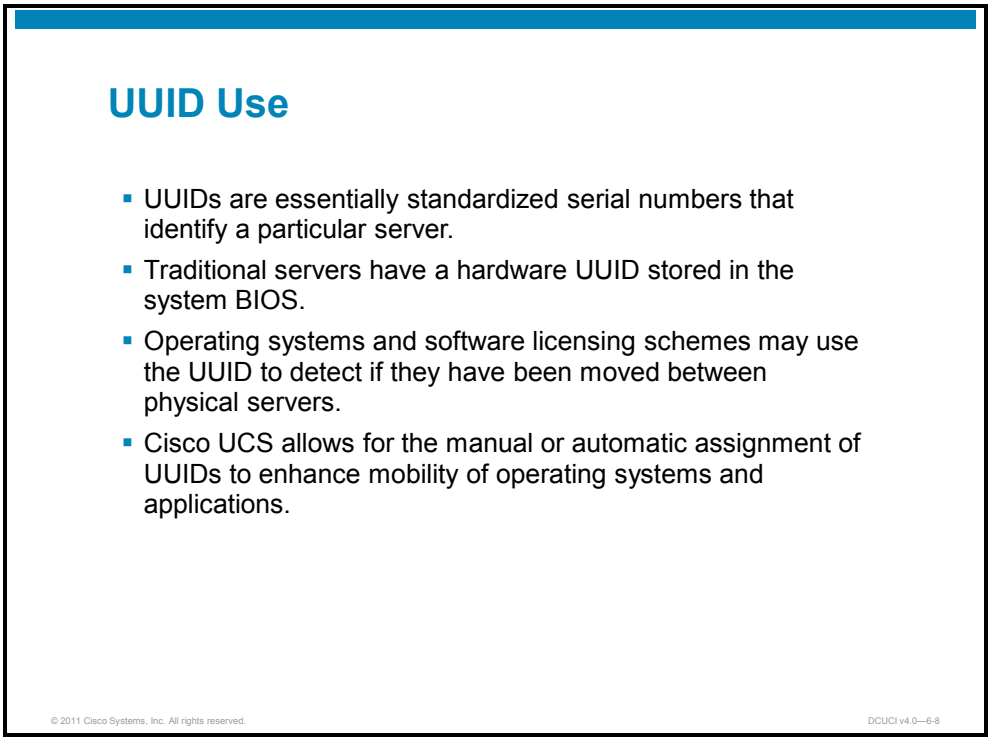

UUIDs are designed as globally unique identifiers for each compute node on a network. UUIDs are used in a number of different ways. In the context of Cisco UCS, the UUID refers to a 128-bit identifier coded into the compute node BIOS.

Operating systems, hypervisors, and applications can leverage the UUID for processes like activation, internal disk labels, and so on. Some applications may use the UUID as an internal root value propagated very tightly within data structures. Therefore, UUIDs should be locally administered in the service profile instead of derived from the BIOS. UUIDs within a service profile are mobile. If the underlying compute node fails, the service profile carries the UUID to the replacement compute node, eliminating the need for potentially time-consuming search-and-replace operations.

UUID Format

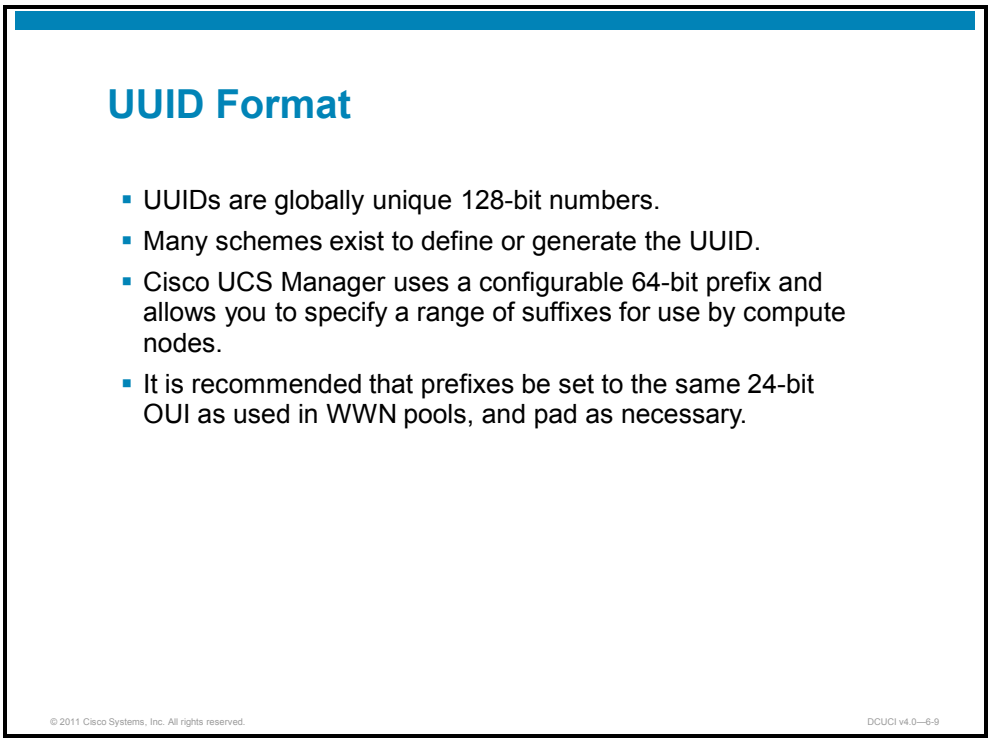

There are many schemas for deploying and formatting UUIDs. It is the responsibility of Cisco UCS to determine what values to encode in the UUID prefix and suffix.

Start the UUID Suffix Pool Wizard

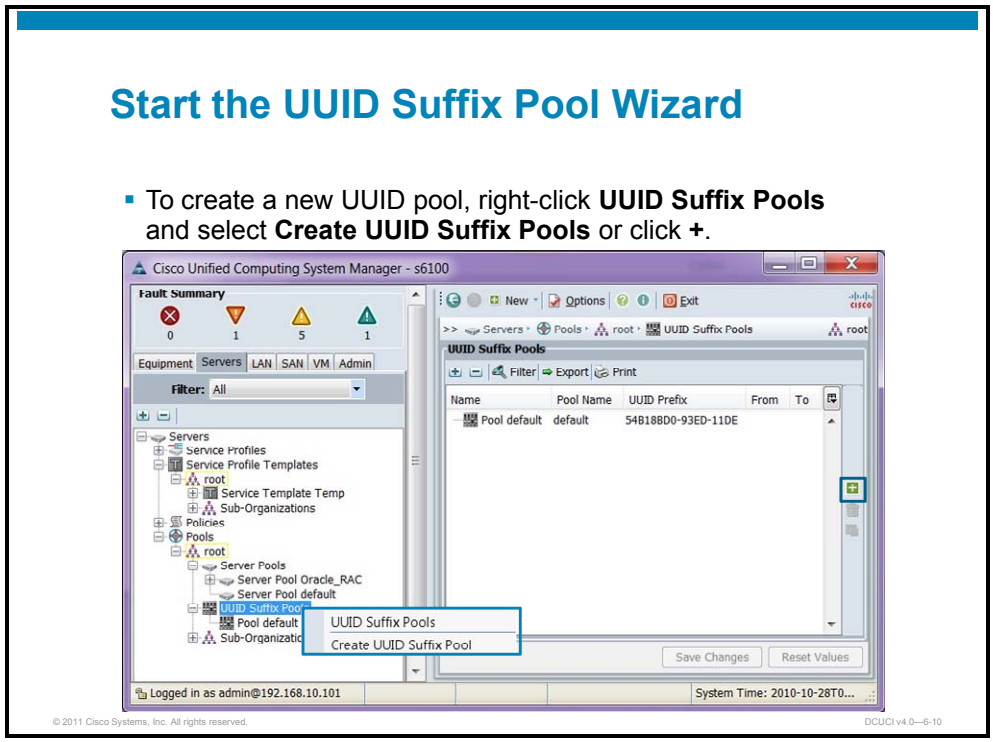

To create a UUID pool, navigate to the Servers tab in the navigation pane. Navigate to Pools and the organization in which the pool should be created. Right-click **UUID Suffix Pools** and choose **Create UUID Suffix Pool**.

Name the UUID Suffix Pool

Name the UUID Suffix Pool

- Select a unique name for the UUID pool, and, optionally, provide a description.

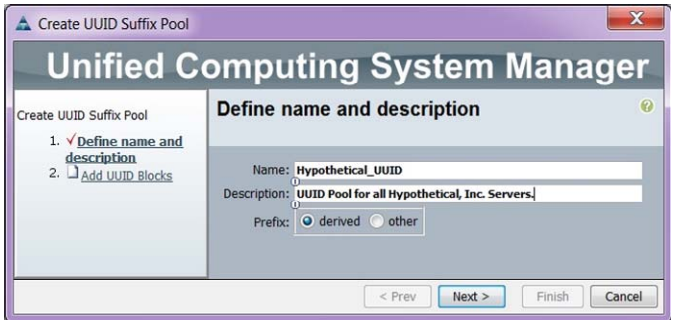

Assign a name and optional description for the pool. There are two choices for creating the UUID prefix. The prefix represents the first 8 bytes of the 16-byte value. If you select **Derived**, Cisco UCS Manager supplies the prefix. If you select **Other**, Cisco UCS Manager will prompt you to supply the first 16 bits of the UUID.

Create a Block of UUID Suffixes

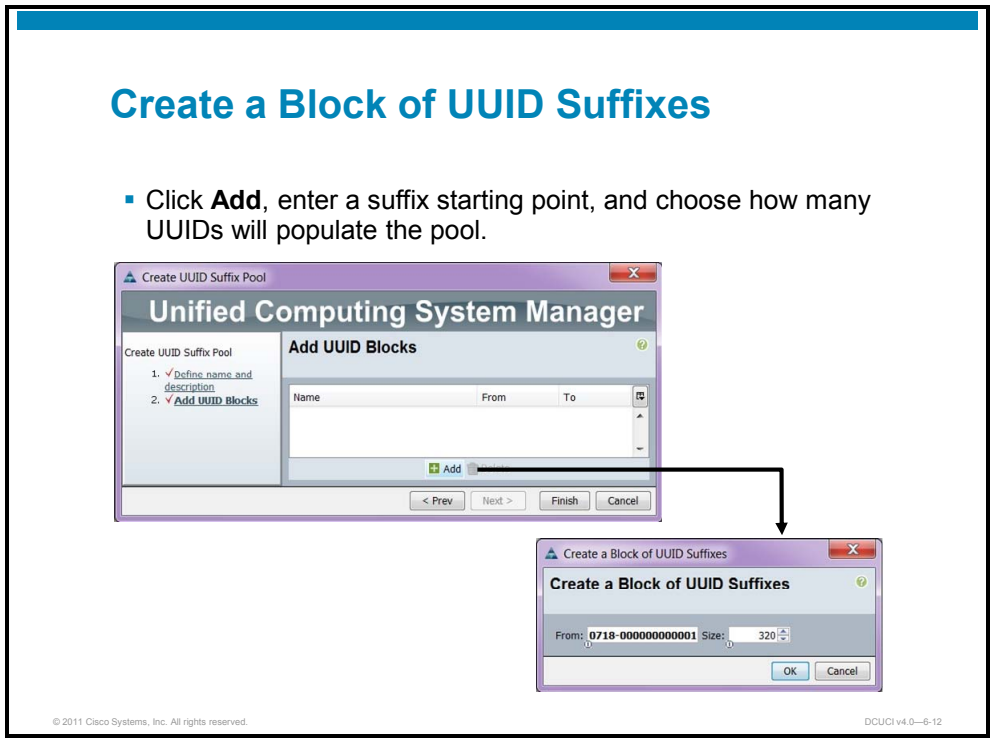

Click **Add** to create a starting point for the 16-bit UUID suffix. In the example, the first two bytes of the suffix have been changed to 0X0718. The current maximum number of compute nodes in Cisco UCS is 320. The designer has decided to preallocate all of the UUIDs that can be used.

Note	It is a best practice to only preallocate the number of identities in a given pool that are based on current and near-term forecast. Every identity resource that is allocated is a managed object in the Cisco UCS Manager database.

UUID Pool

UUID Pool

- UUIDs are now available for assignment.

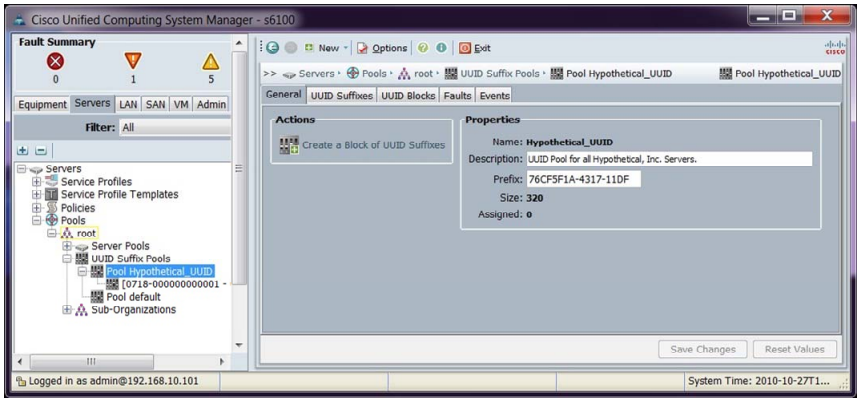

A pool of 320 UUIDs were created using the derived prefix that is combined with the 320 defined suffixes. They are immediately available for consumption by service profiles or service profile templates.

MAC Pools

This topic discusses the configuration of MAC address pools.

Start the MAC Pool Wizard

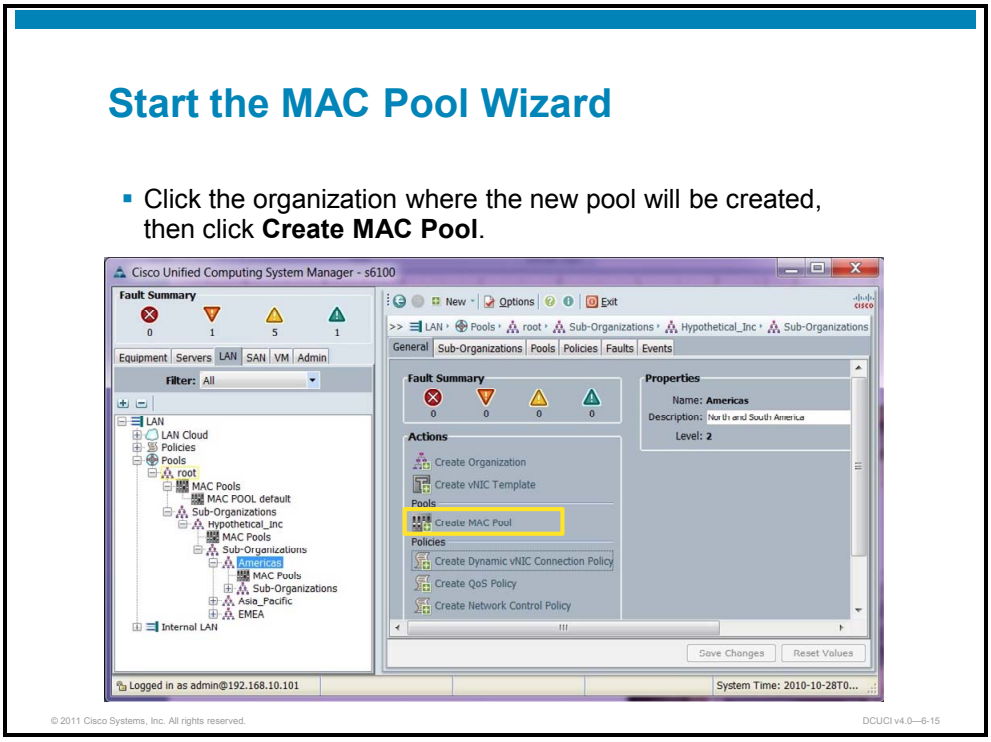

A MAC pool consists of a range of MAC addresses. Create the MAC pool and assign a name. With MAC pools, Cisco UCS administration is made easier when scaling server deployment of service profiles by prompting stakeholders to define a set of MAC addresses before actual deployment.

To create a MAC pool, navigate to the LAN tab in the navigation pane. Click the organization that the pool should be created beneath. Click the **Create MAC Pool** link to start the wizard.

Name the MAC Pool and Create a MAC Block

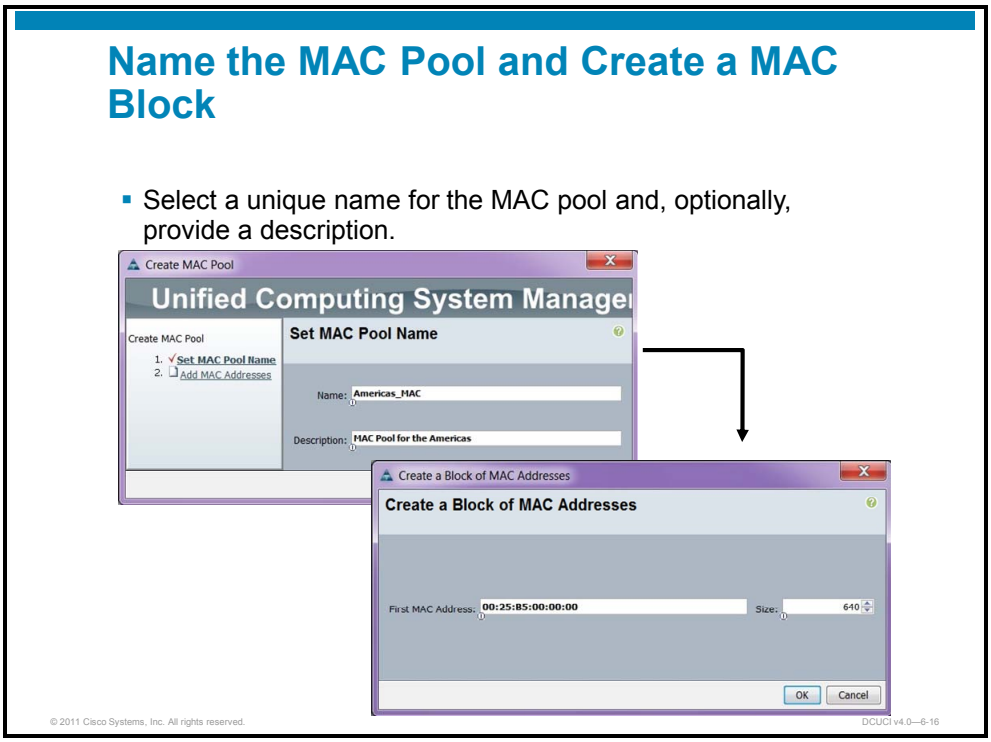

Provide the MAC pool with a unique name and, optionally, a description. Click **Next** and decide how many MAC addresses should be created in the pool. Cisco provides a three-byte Organizationally Unique Identifier (OUI) assigned by the IEEE. It is recommended that you do not modify the prefix.

Verify the MAC Block and Finish the Wizard

Verify the MAC Block and Finish the Wizard

- Verify that the MAC pool is correct or click **<Prev** to go back to creating addresses.
- Click **Finish** to complete the wizard.

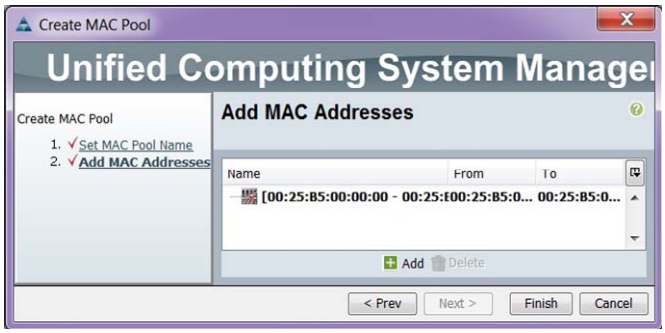

When the MAC pool has been created, there is an opportunity to verify the addresses and go back to the previous window if a mistake has been noticed. Otherwise, click **Finish** to complete the wizard.

New Americas MAC Block

New Americas MAC Block

- The Americas MAC pool is now available for assignment.

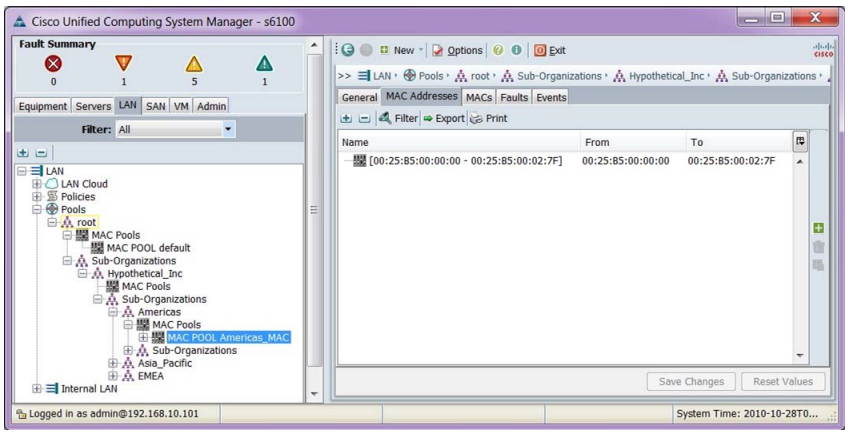

A MAC address pool was added beneath the Americas organization. These addresses are immediately available for assignment.

WWNN Pools

This topic discusses the creation and configuration of world wide node names (WWNNs).

WWN Format

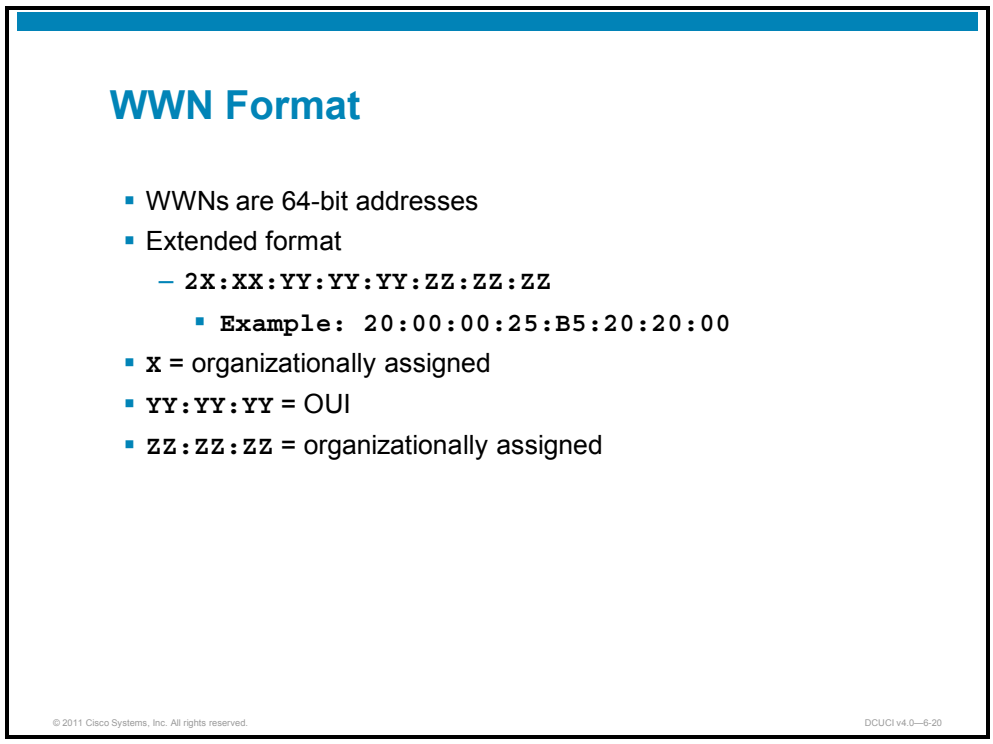

WWNs are 64-bit addresses that have many possible formats. The example that is shown is for reference only, as Cisco UCS Manager enforces a specific format in all WWN pools.

A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved.

To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, you should use the following WWN prefix for all blocks in a pool:

20:00:00:25:B5:XX:XX:XX

WWN Pool Considerations

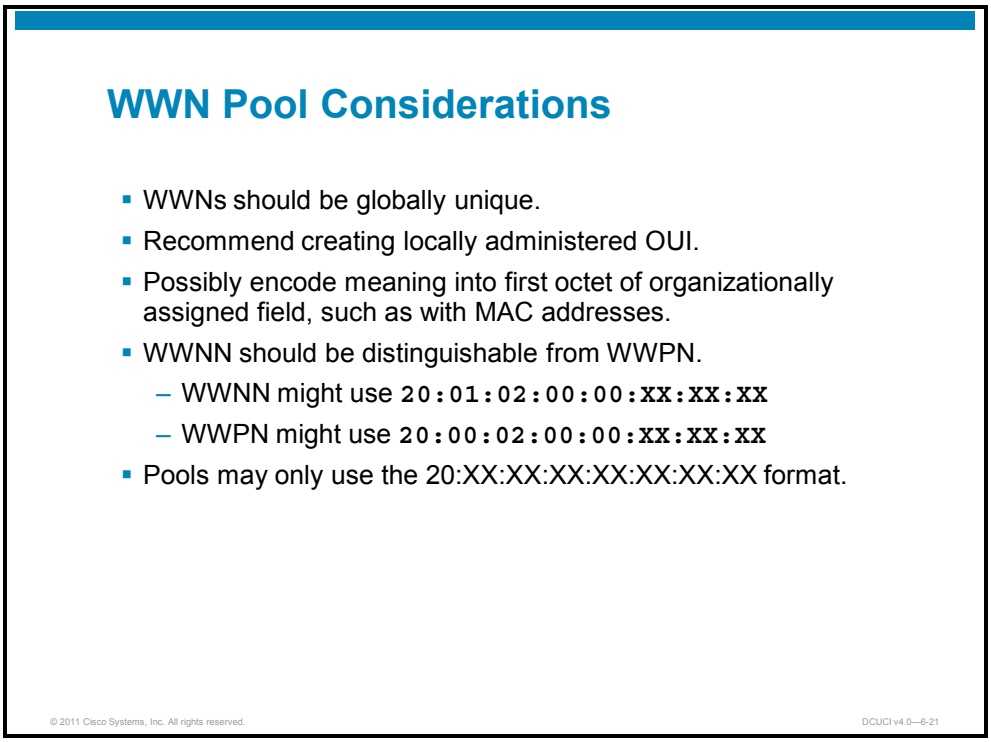

Cisco UCS Manager enforces the use of WWNNs that begin with "20." All WWN pools must begin with that value, with any remaining values that you create. In keeping with the global standards set for WWNs, it is recommended that a locally administered OUI be selected and used as the third through fifth octets. Additionally, it is useful if the WWNN and world wide port name (WWPN) values are easily distinguishable. Because the second octet of a WWN can be organizationally assigned, that octet might be used to encode meaning for the WWNN or WWPN. This convention and address block should be agreed upon by all stakeholders in the initial implementation phase of a Cisco UCS deployment.

Start the WWNN Pool Wizard

Start the WWNN Pool Wizard

- Click the organization where the new pool will be created. Then click **Create WWNN Pool**.

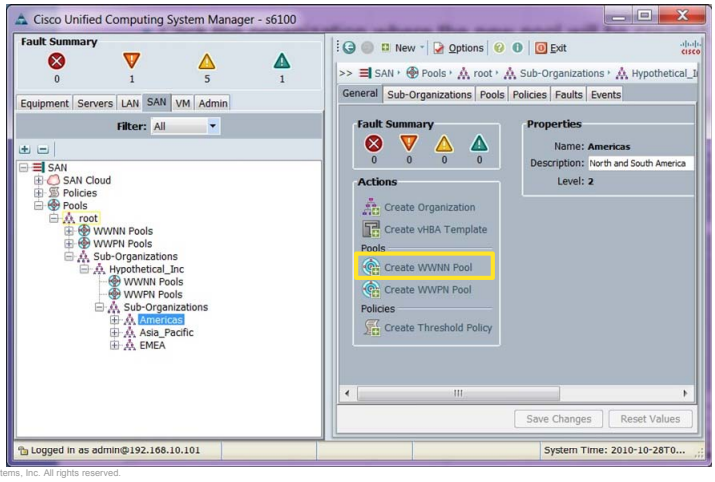

DCUCI v4.0—6-22

To create a WWNN pool, navigate to the SAN tab in the navigation pane. Click the organization that the pool should be created beneath. Click **Create WWNN Pool** to start the wizard.

Name the Pool and Create a WWNN Block

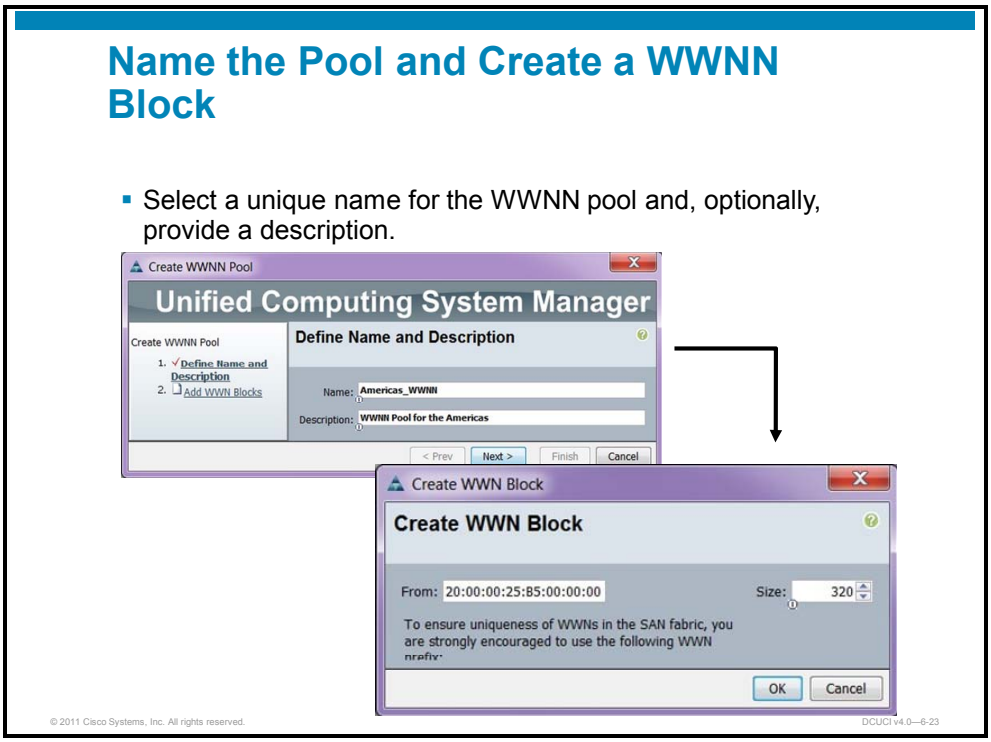

Cisco supplies the first four bytes of the WWNN prefix by combining “20” with one of the three-byte OUIs provided by Cisco. One WWNN is used by each service profile. In the example that is shown, the administrator has decided to preallocate 320 addresses.

Verify the WWNN Block and Finish the Wizard

Verify the WWNN Block and Finish the Wizard

- Verify that the WWNN Pool is correct or click **<Prev** to go back to the addresses.
- Click **Finish** to complete the wizard.

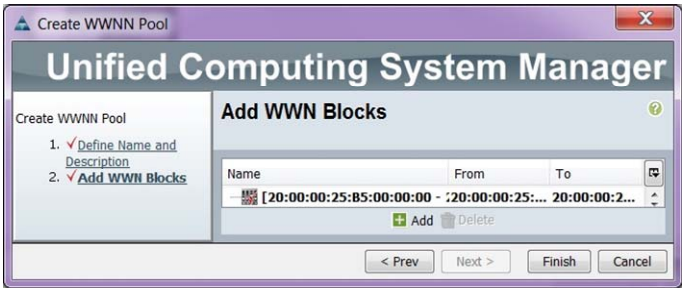

After the WWNN pool has been created, there is an opportunity to verify the addresses and go back to the previous window if you notice a mistake. Otherwise, click **Finish** to complete the wizard.

New Americas WWNN Block

New Americas WWNN Block

- The Americas WWPN Pool is now available for assignment.

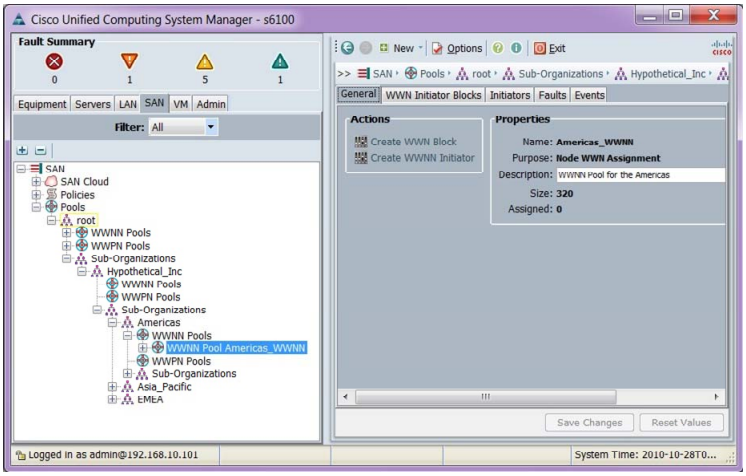

A WWNN pool was added beneath the Americas organization. These addresses are immediately available for assignment.

WWPN Pools

This topic discusses the creation of WWPN pools.

Start the WWPN Pool Wizard

To create a WWPN pool, navigate to the SAN tab in the navigation pane. Click the organization that the pool should be created beneath. Click **Create WWPN Pool** to start the wizard.

Name the Pool and Create a WWPN Block

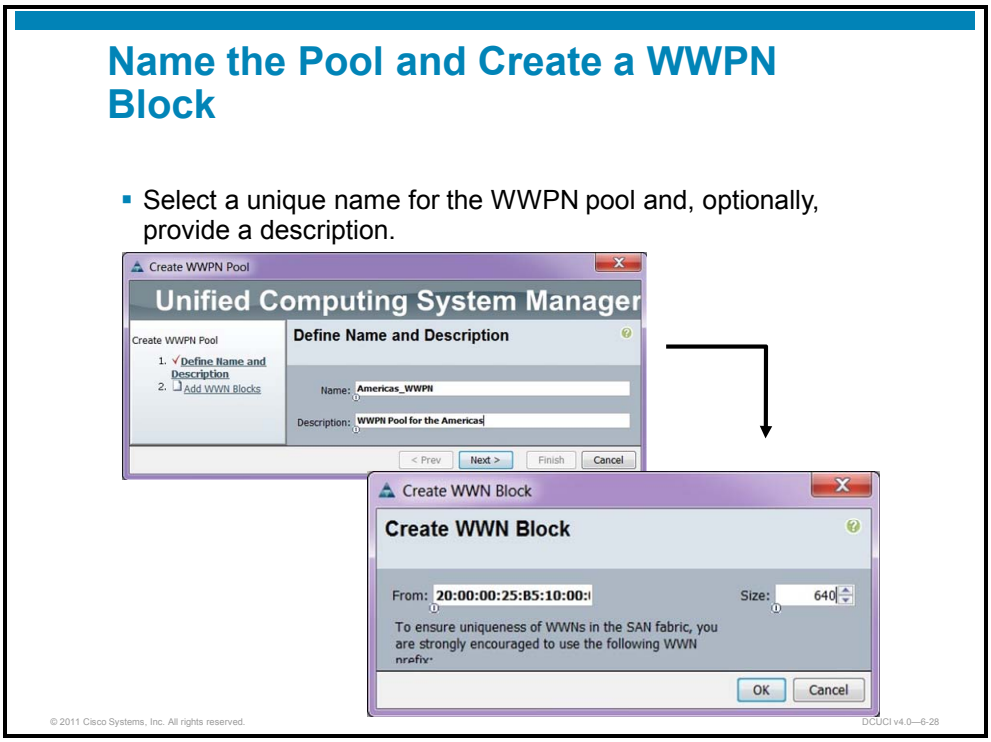

Cisco supplies the first four bytes of the WWPN prefix by combining “20” with one of the three-byte OUIs provided by Cisco. One WWNN is used by each service profile. In this example, the administrator decided to preallocate 640 addresses. One WWPN is required for each virtual host bus adapter (HBA).

Verify the WWPN Block and Finish the Wizard

Verify the WWPN Block and Finish the Wizard

- Verify that the WWPN pool is correct or click **<Prev** to go back to create addresses.
- Click **Finish** to complete the wizard.

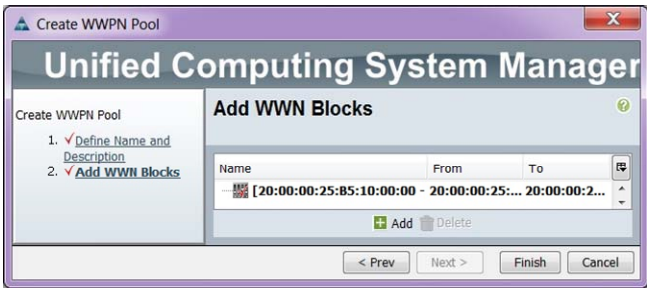

After the WWPN pool has been created, there is an opportunity to verify the addresses and go back to the previous window if you notice a mistake. Otherwise, click **Finish** to complete the wizard.

New Americas WWPN Block

New Americas WWPN Block

- The Americas WWPN pool is now available for assignment.

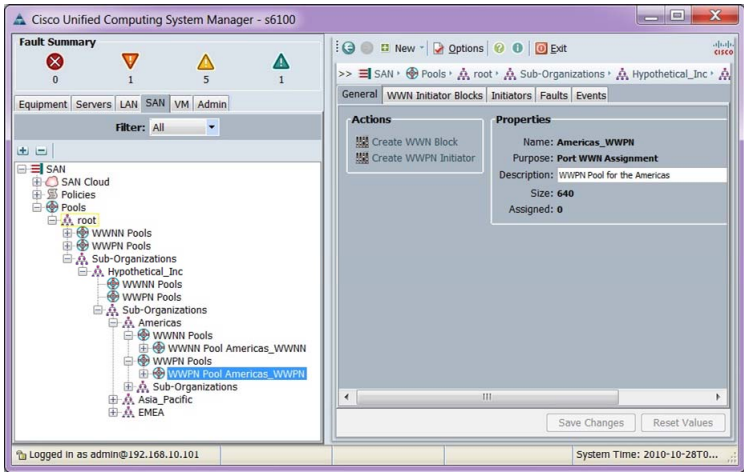

A WWPN pool was added beneath the Americas organization. These addresses are immediately available for assignment.

Server Pools

This topic discusses the creation of server pools.

Server Pools

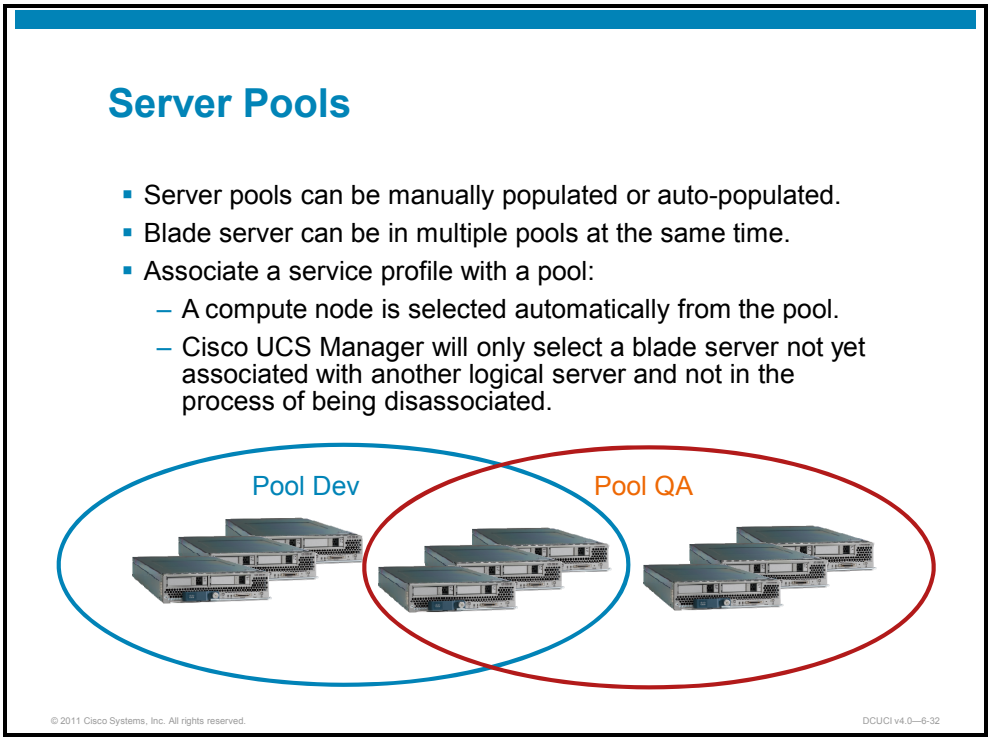

Pools can be manually populated or auto-populated using a server pool policy.

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

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

Start the Server Pool Wizard

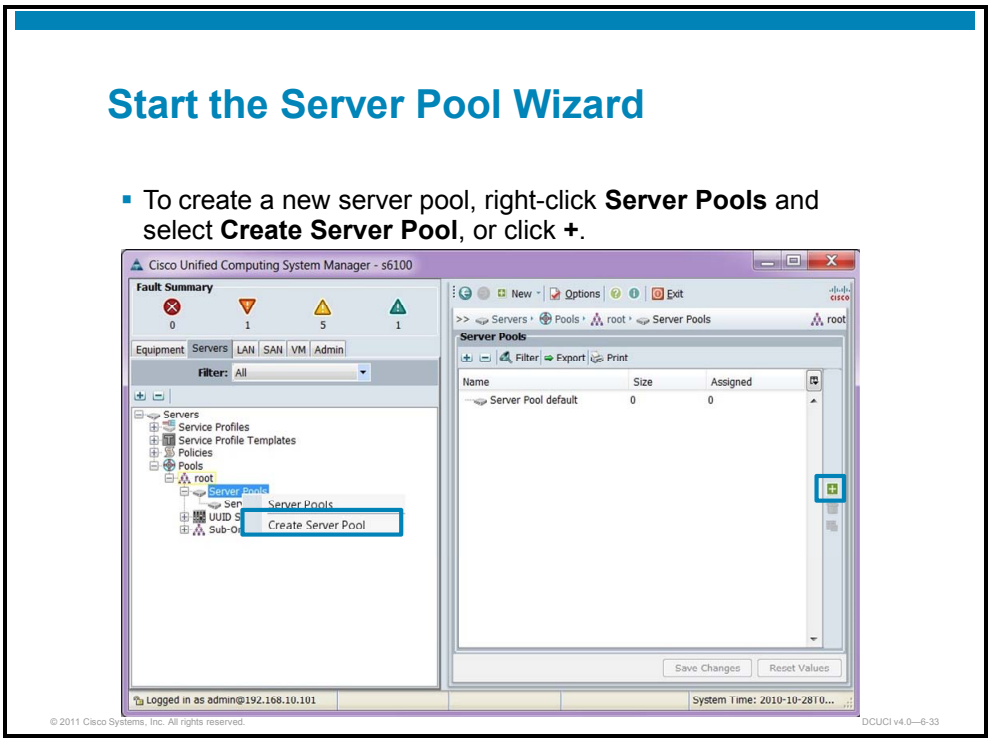

Server pools are configured under the LAN tab in the navigation pane. To create a new server pool, right-click **Server Pools** and select **Create Server Pool**, or click the plus sign.

Name the Server Pool

Name the Server Pool

- Select a unique name for the server pool and, optionally, provide a description.

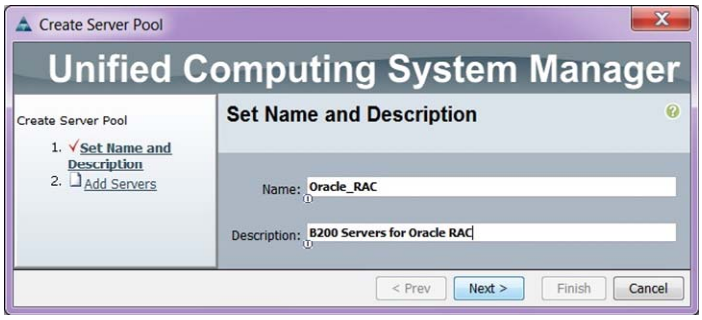

Enter a unique name and, optionally, a description for the new server pool and click **Next**.

Select Server Pool Members

Select Server Pool Members

- Select servers and click the **>>** button to add them to the pool.

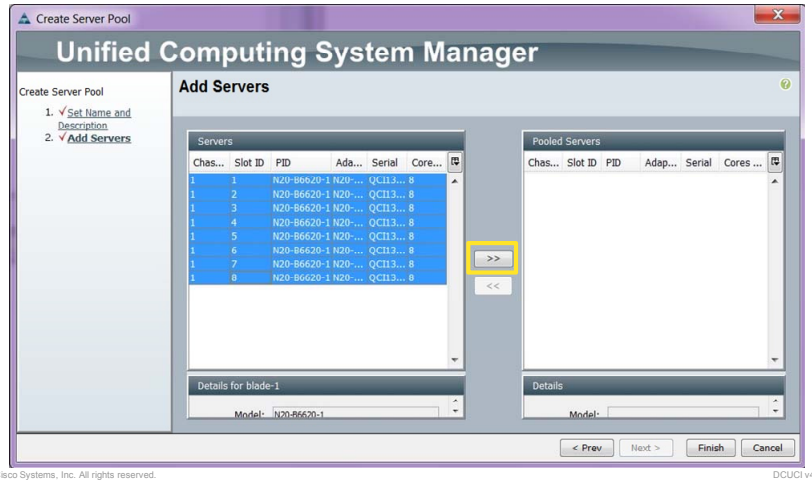

Use the mouse to populate the new server pool. Hold down the shift key to select a range of servers. Click the >> button to move the selected servers into the pool.

New Server Pool Members

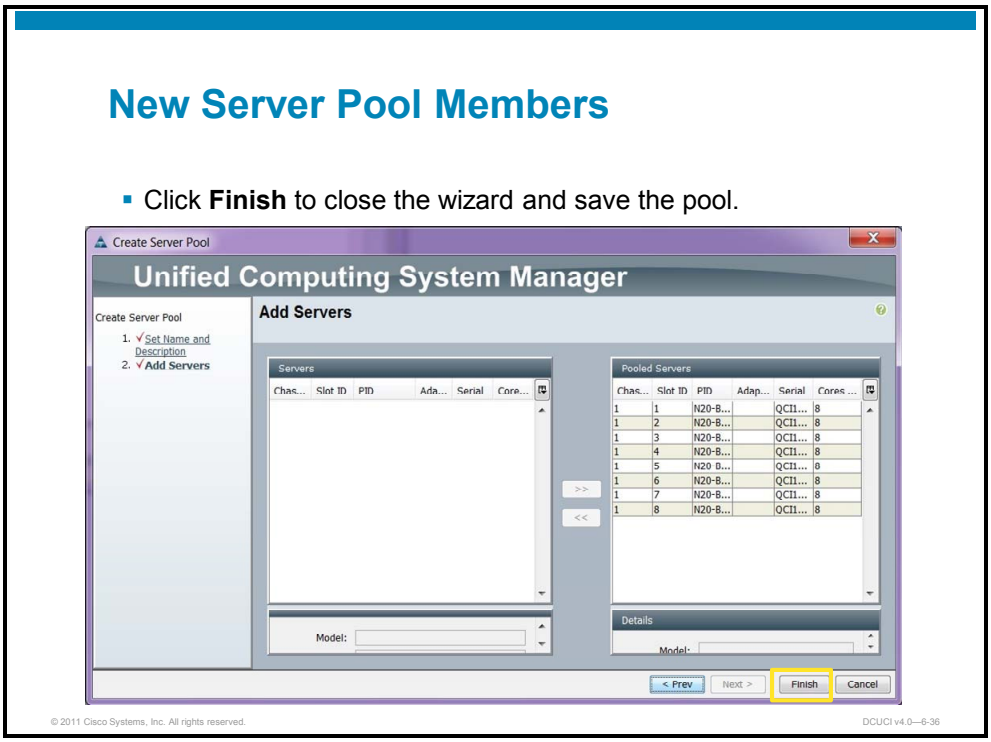

Verify that the desired servers are members of the new pool. Click **Finish** to complete the wizard.

Server Pool Created

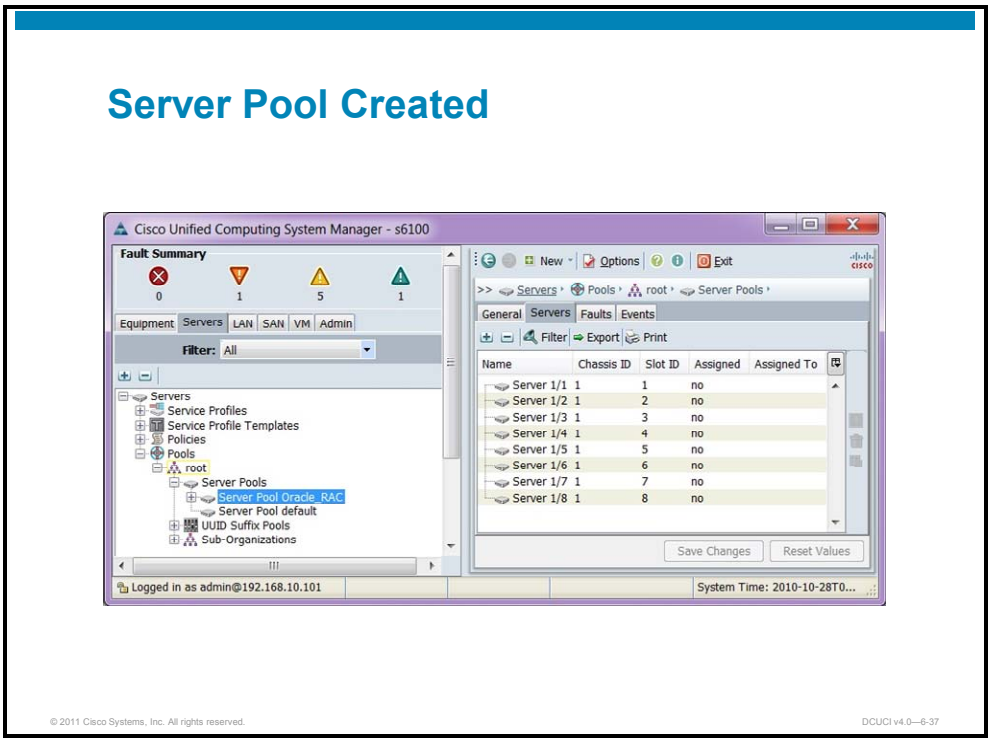

The content pane displays the IDs of the servers added to the pool. It includes information about whether the server has been assigned. If assigned, a link to the service profile of the server will display in the Assigned To column.

Automating Server Pool Membership Based on a Qualification Policy

This topic discusses how to configure a qualification policy to automatically add new blade servers to a given pool.

Creating an Auto-Populating Pool

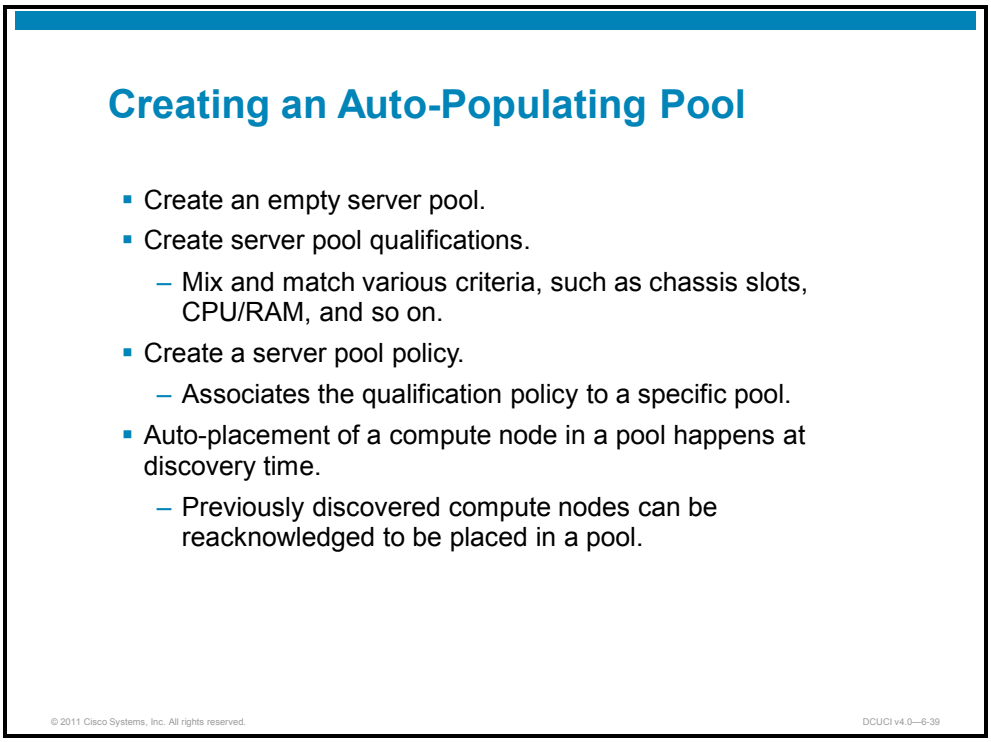

The auto-population feature lets you do the following:

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

Note	Server pool auto-population only happens as individual blade servers are discovered. If you want auto-population to apply to existing compute node servers, they need to be reacknowledged.

Create an Empty Server Pool for SAP-Qualified Servers

Create an Empty Server Pool for SAP-Qualified Servers

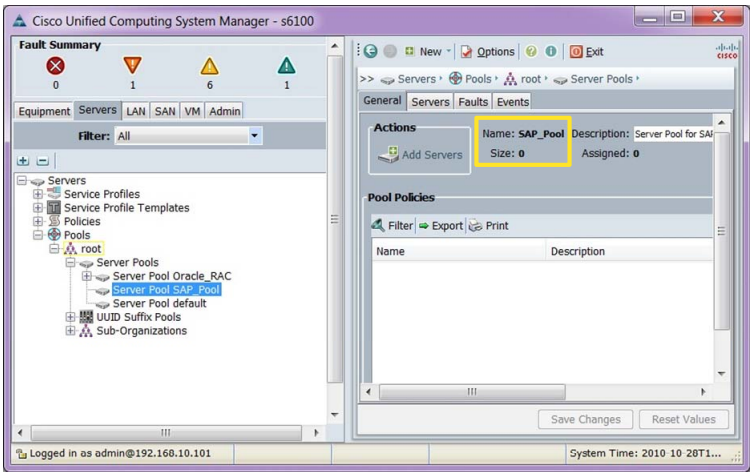

Create an empty server pool where servers matching a qualification criteria will be placed.

Create a New Qualification Policy

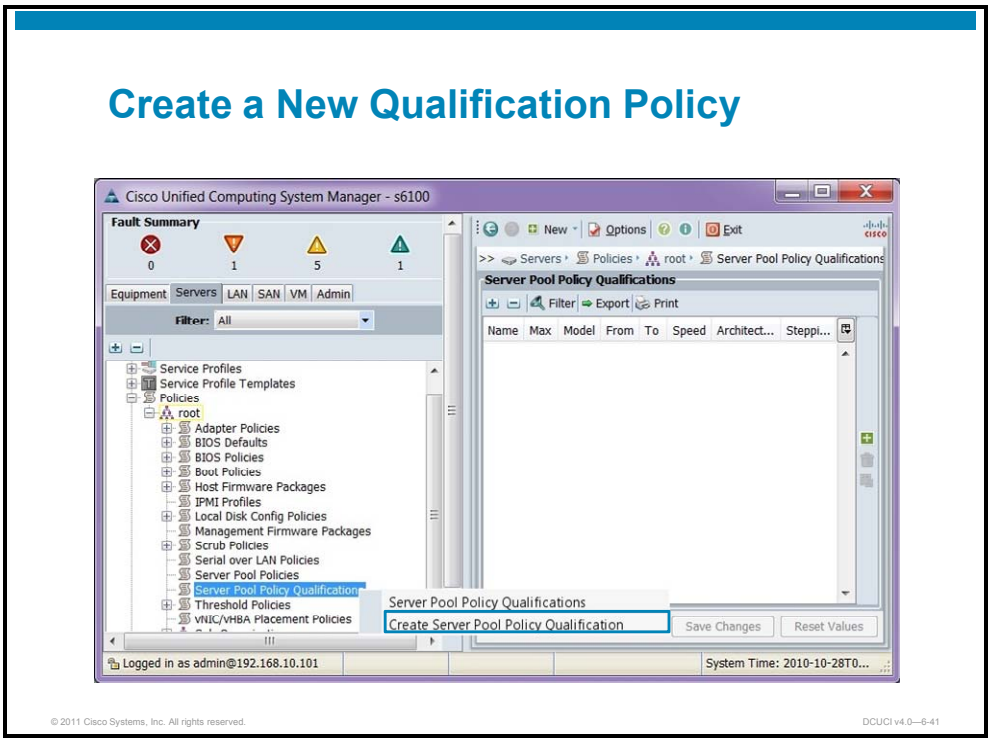

In the Servers tab of the navigation pane, expand the Policies member of the navigation tree. Right-click the **Server Policy Pool Qualification** policy and select **Create Server Pool Policy Qualification**.

Server Selection Criteria

Server Selection Criteria

- Choose server selection criteria from the Actions box.

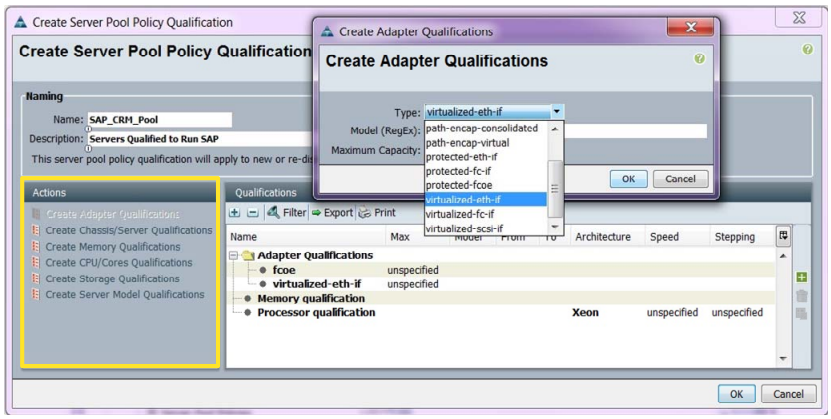

In the Actions section of the policy wizard, there is a list of six categories that can be used for selection criteria. The policy SAP_CRM_POOL is assigned to name this policy. In the example, requirements for CPU memory and mezzanine card are specified.

Create Server Pool Policy

Create a Server Pool Policy

- Map the qualifications to a server pool.

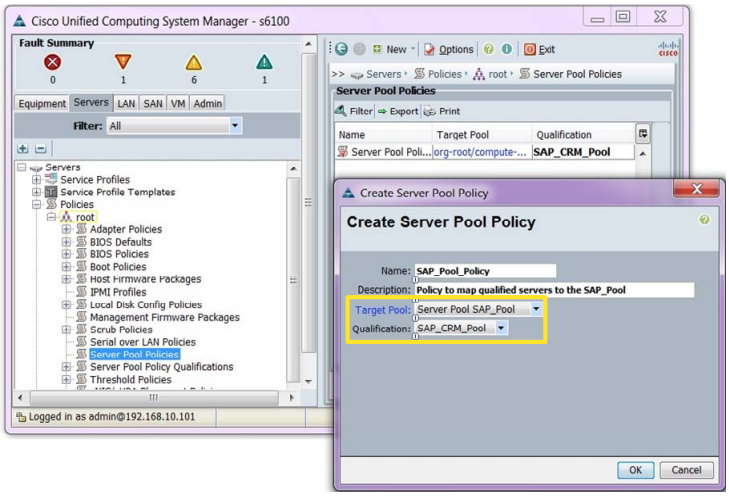

The next step is to create a server pool policy. The purpose of this policy is to map a qualification policy to the empty pool that was created earlier.

Servers Auto-Populated After Discovery

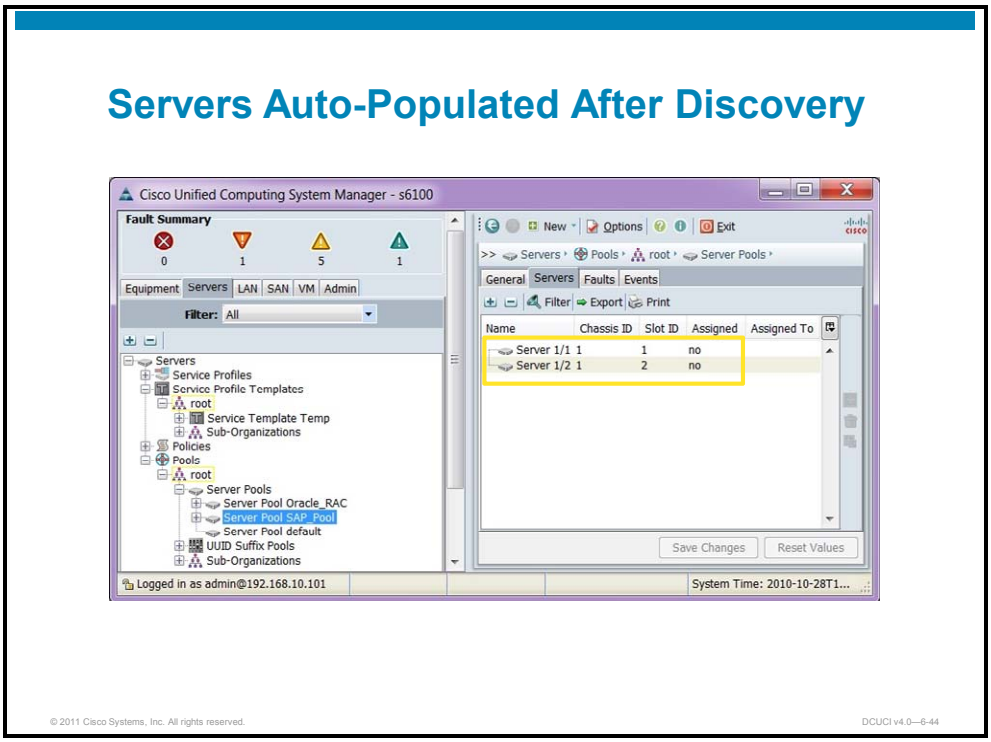

The example shows that two servers were automatically added to the pool SAP_Pool. As was previously discussed, servers that have already been discovered will not automatically be matched against a qualification policy. Both servers in the example were reacknowledged. After discovery was completed, they were added to the SAP_Pool server pool. A new blade server that is inserted into a chassis will be evaluated by all qualification policies that have been configured.

Importance of Creating Pools in the Correct Organization

This topic discusses restrictions that are imposed on objects that are created in an organization.

Objects Created in an Organization Cannot Move

Objects Created in an Organization Cannot Move

- It is possible to copy a pool, policy, or service profile, but not to paste it.

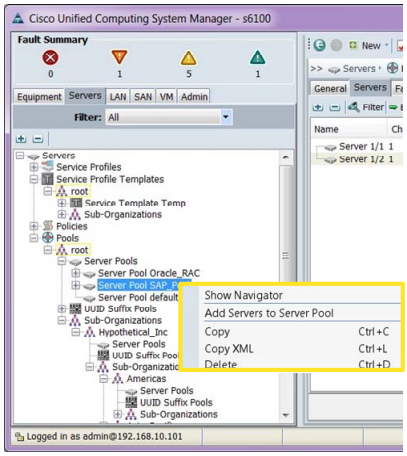

DCUCI v4.0—6-46

Important lessons in IT are sometimes learned in the most difficult way possible. One such example would be creating many service profiles, templates, policies, pools, and thresholds under the root organization. If it is later decided to create suborganizations, it is not possible to move any of the objects that are created under root to another organization. It is also not possible to move an object from one nonroot organization to another. When you right-click on a policy object or template, notice that there is no option for “cut.” There is the tantalizing option of “copy,” but if one right-clicks on a different organization, it is apparent that there is no option for “paste.”

The only way to remedy the situation is to delete and re-create every object in its appropriate organization.

Summary

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

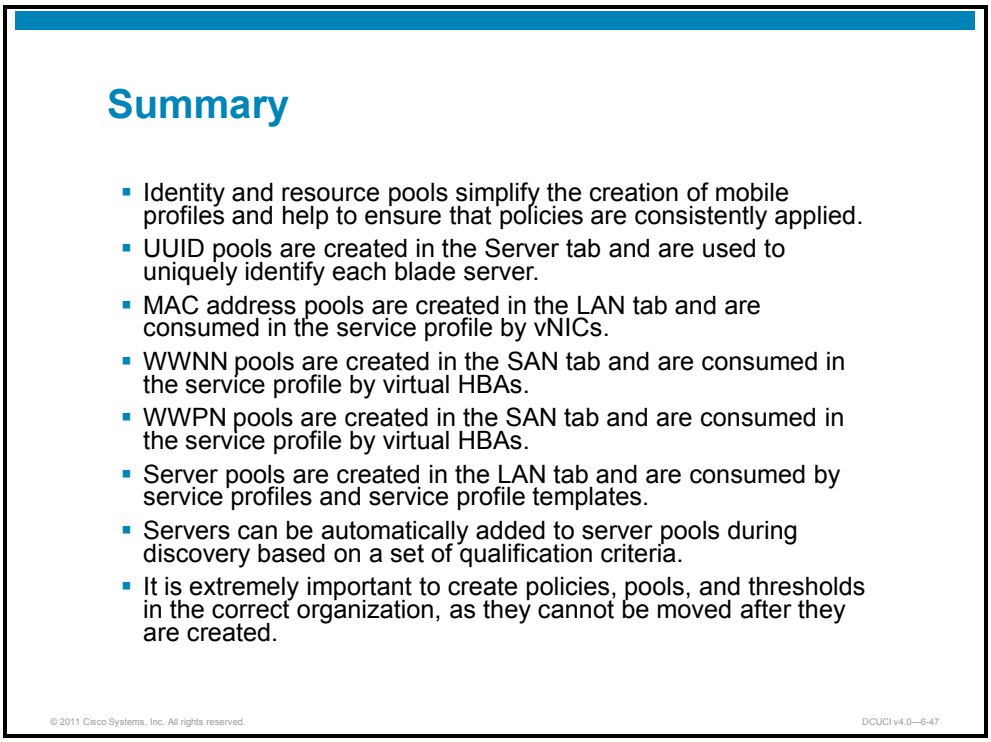

Lesson 2

Creating Service Profiles

Overview

Service profiles are the structural elements of the compute node. Service profiles contain all of the identity resources and policies that are required to operate a blade server. The service profile makes stateless computing possible by abstracting identity and policy information from the physical server. If the compute node that a service profile is associated with becomes unavailable, the service profile is simply reassociated with another compatible compute node. When the operating system or hypervisor reboots on the replacement compute node, it believes that it is running on the same hardware.

Objectives

Upon completing this lesson, you will be able to configure service profiles and service profile templates. This ability includes being able to meet these objectives:

- Describe the rationale for and benefits of service profiles
- Configure a BIOS policy to enable virtualization features
- Configure an adapter policy to enable RSS and set the failback timer for fabric failover
- Create a QoS system class and allow all Ethernet traffic to use jumbo frames up to an MTU of 9216
- Configure IPMI and SoL policies
- Configure a scrub policy for local disks and BIOS
- Differentiate between the features available in the simple service profile wizard and the expert wizard
- Start the service profile expert wizard
- Configure the service profile to take its UUID from a pool
- Configure a vHBA for two fabrics and have the service profile take its assignment of WWNNs and WWPNs from a pool
- Configure a vNIC for two fabrics and have the service profile take its assignment of MAC addresses from a pool
- Configure vNIC and vHBA placement on full-slot blades

- Configure the binding of a vHBA to a Fibre Channel boot target
- Configure server assignment
- Differentiate between required components and optional components of the service profile definition

Benefits of Service Profiles

This topic discusses the benefits of and rationale for creating service profiles.

Service Profile Basics

Service Profile Basics

- Service profiles contain identity and state information for a logical server.
- LAN and SAN connectivity to a compute node is unavailable without a service profile associated with it.
- Every compute node needs its own unique service profile.

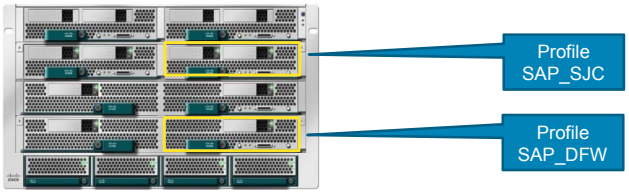

Stateless computing requires unique identity resources for universally unique identifiers (UUIDs), MAC addresses, and world wide names (WWNs) for Fibre Channel. Using pooled resources ensures consistent application of policy and reasonable assurances that identities are unique within the Cisco UCS Manager.

Service Profile Contents

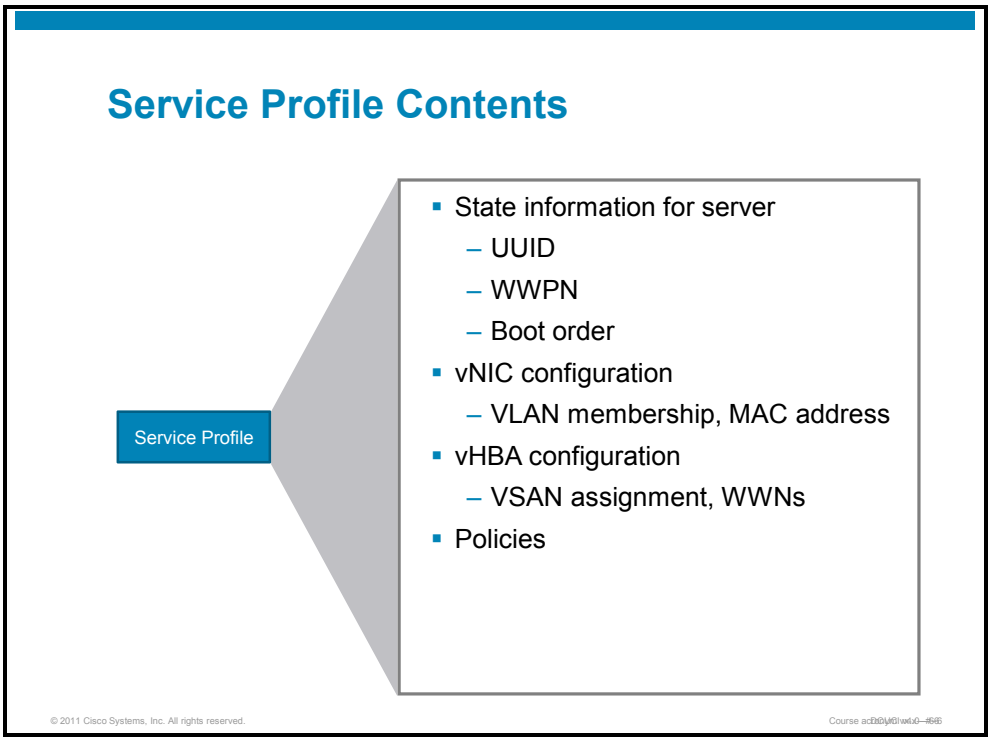

The service profile represents a logical view of a server without any ties to a specific physical device. The profile object contains all the elements of server function. This identity contains the unique information for that server, including MAC, world wide port name (WWPN), UUID, boot order, and so on. Each profile can only be associated with a single blade server at any given time, and every blade server requires a unique service profile.

Service profiles facilitate server mobility. Mobility is the ability to transfer server identity seamlessly between compute nodes in such a way that the underlying operating system or hypervisor does not detect any change in server hardware.

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

Configuration of a BIOS Policy

This topic describes the configuration of BIOS policies that can be applied to service profiles in order to enable virtualization features.

Locate BIOS Policies

Locate BIOS Policies

- Expand policies in the Server tab and locate BIOS Policies.
- Right-click on the organization where the new policy is to be created.

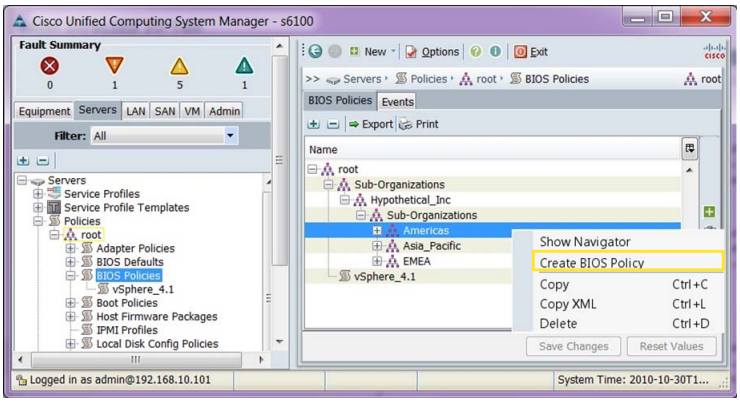

In the Server tab of the navigation pane, expand the policies to locate BIOS Policies. Right-click on the organization in the content pane where the new policy is to be created.

Name New BIOS Policy

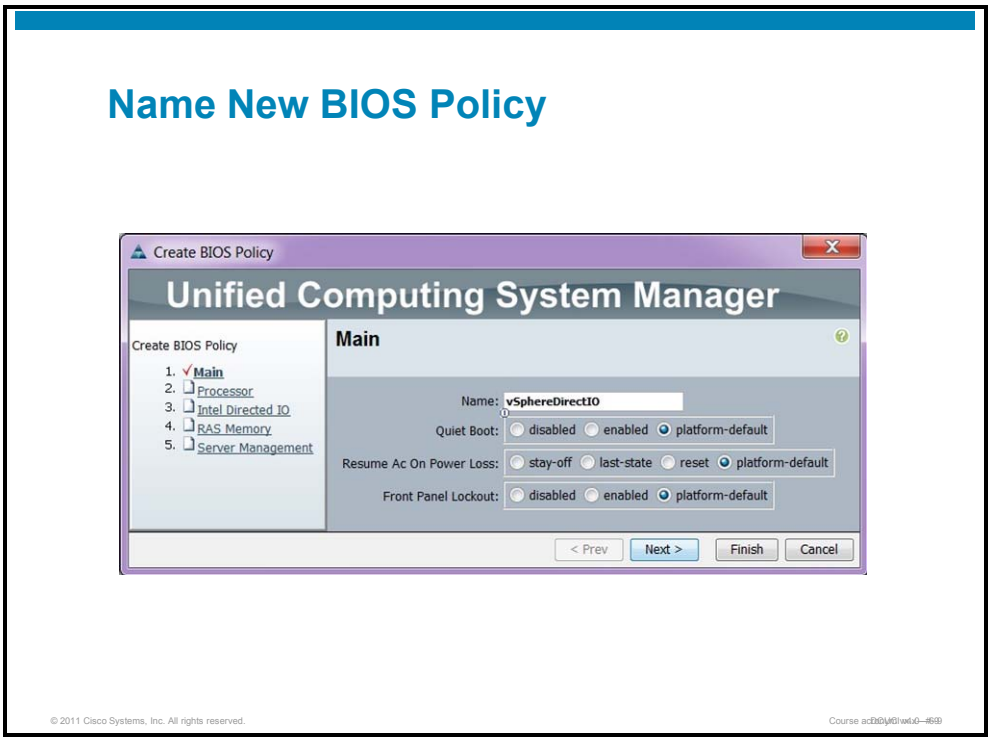

In earlier Cisco UCS Manager versions, the only BIOS setting that could be modified was boot order. Beginning in Cisco UCS Manager version 1.3, control of nearly every BIOS setting is available as a policy that can be assigned to a service profile.

The BIOS policy requires a name. There are three options for each BIOS setting:

- **Disabled:** The setting or attribute is disabled.
- **Enabled:** The setting or attribute is enabled.
- **Platform-default:** The setting or attribute will remain at factory default.

Enable CPU Performance Features

Enable CPU Performance Features

- Enable performance-enhancing CPU features and virtualization support.

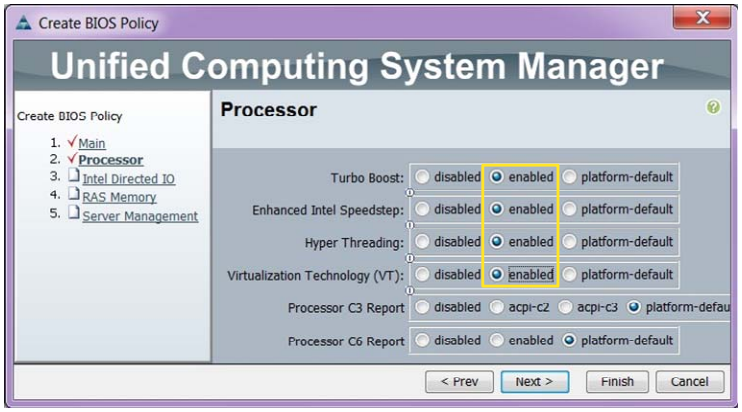

The settings on the example shown in the figure enable Intel performance features of the CPU. The service profile of this BIOS policy will be applied and can take complete and direct advantage of the CPU features.

Feature	Description
Turbo Boost	Allows dynamic overclocking of the CPU if the operating system requests.
Enhanced Intel Speedstep	This feature allows the CPU to lower clock frequency to lower power consumption during periods of low demand.
Hyper Threading	This is an Intel proprietary technique that can improve instruction parallelization by presenting an operating system or hypervisor with two virtual CPUs for each physical core.
Virtualization Technology	This feature enables Intel virtualization-specific CPU instructions to improve hypervisor performance.

Enable Intel Direct I/O Support

Enable Intel Direct I/O Support

- Enable CPU support for acceleration of VMware DirectPath I/O.

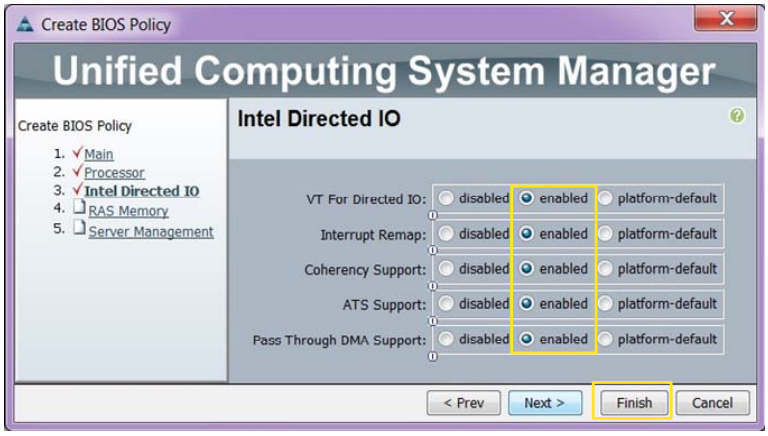

Intel Virtualization Technology for Directed I/O (VT-d) options can accelerate I/O operations in virtualized environments where a physical interface is mapped directly to a virtual machine, bypassing the hypervisor. VMware vSphere can benefit from enabling these options. Refer to operating system or hypervisor documentation for guidance on the appropriate settings for these options.

BIOS Policy “vSphereDirectIO” Available

BIOS Policy “vSphereDirectIO” Available

- The new BIOS policy is available for assignment.

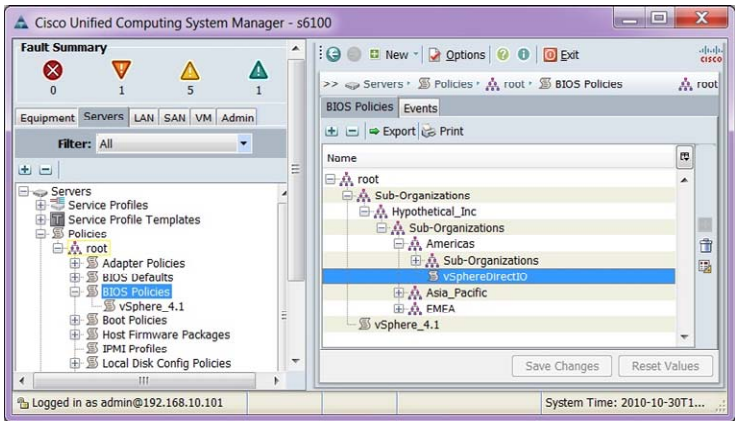

Click **Finish** in the wizard to save the new BIOS policy. The new policy is now available for assignment to a service profile. When a service profile that references this policy is associated to a blade server, there is no need for manual configuration of the BIOS at power-on self-test (POST) time. This new capability can greatly simplify and accelerate the pace of server provisioning.

Configuration of an Adapter Policy

This topic describes the configuration of adapter policies that can be applied to service profiles in order to enable RSS and set the failback timer for fabric failover.

Locate Adapter Policies

Locate Adapter Policies

- Expand Policies in the Server tab, and locate Adapter Policies.
- Right-click on the organization where the new policy is to be created.

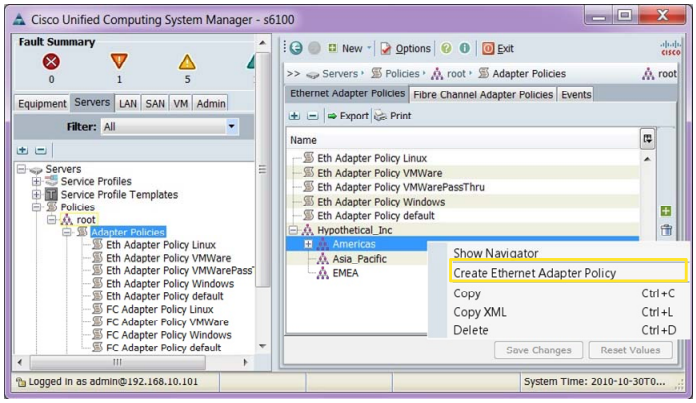

In the Server tab of the navigation pane, expand Policies to locate Adapter Policies. Right-click on the organization in the content pane where the new policy is to be created.

Create a New Adapter Policy

Create a New Adapter Policy

- Enable **Receive Side Scaling (RSS)** and lower the Failback Timeout to 2 seconds.

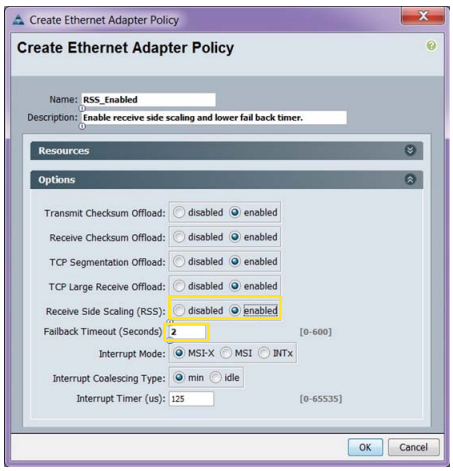

Receive-side scaling (RSS) relieves the single-server bottleneck that occurs in multicore CPU systems. TCP packets that are received without RSS being enabled are only processed by a single core. By enabling RSS, received packets are processed on all cores. RSS should generally be enabled on any server with multicore CPUs. This adapter feature is disabled by default.

The failback timer determines how long an adapter should wait to fail back to its original fabric if a fabric failover event has occurred. As an example, if fabric interconnect A became unavailable, servers would failover to their backup connections on the fabric interconnect B. The 2-second timer that is employed here would apply as soon as fabric interconnect A becomes available.

Adapter Policy “RSS_Enabled” Available

Adapter Policy “RSS_Enabled” Available

- The new adapter policy is available for assignment.

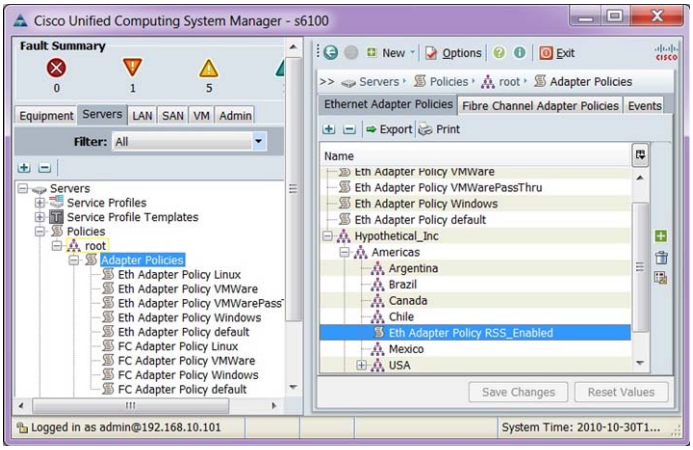

Click **Finish** in the wizard to save the new adapter policy. The new policy is now available for assignment to a service profile.

Create a QoS System Class

This topic describes the two-step process to modify QoS system classes and apply them to adapter policies in order to allow all Ethernet traffic to use jumbo frames.

Open LAN Uplinks Manager

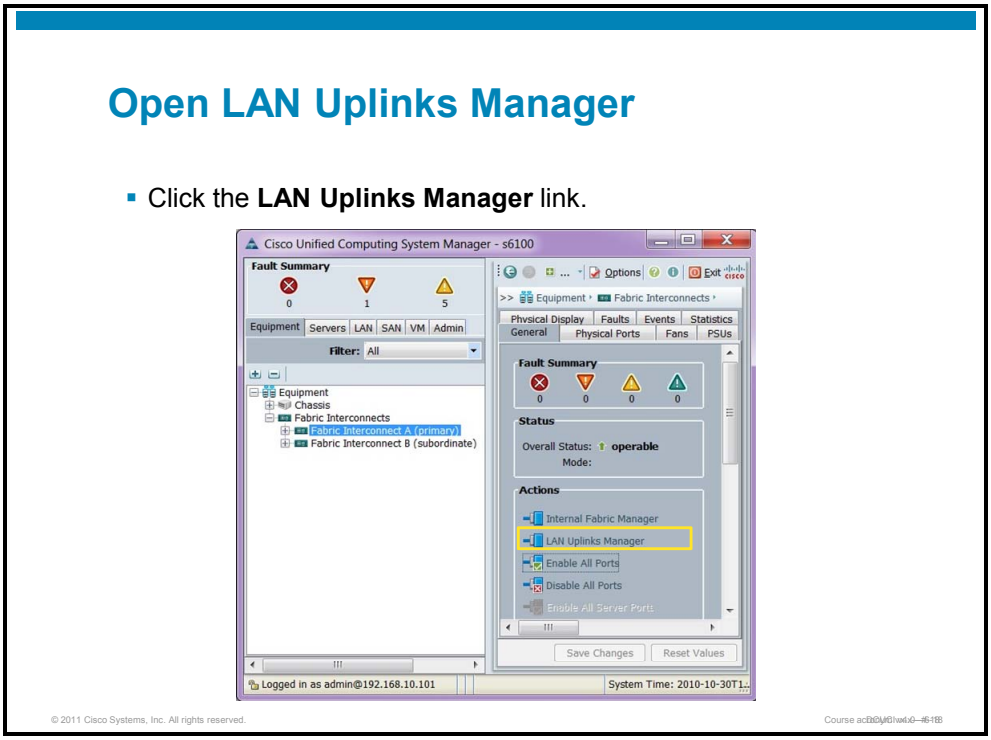

You can create a quality of service (QoS) policy in two steps. To access a QoS system class, the LAN Uplinks Manager must be opened.

From the Equipment tab of the navigation pane, select one of the fabric interconnects. In the content pane, click the **LAN Uplinks Manager** link.

Select QoS Tab in Content Pane

Select QoS Tab in Content Pane

- Click the **QoS** tab to access QoS system classes.

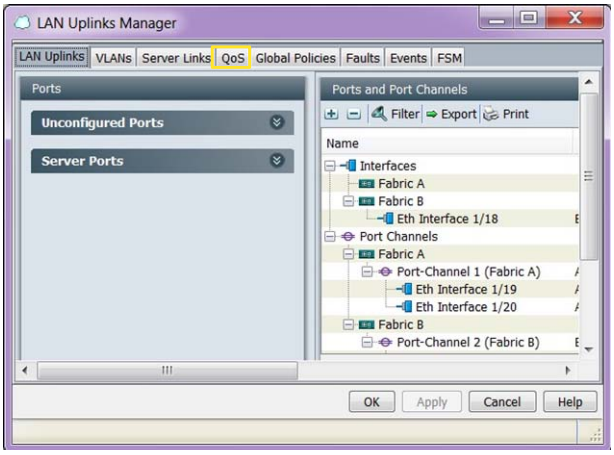

Click the **QoS** tab in the content pane to open the dialog box to access QoS and modify system classes.

Modify a System QoS Class

Modify a System QoS Class

- Enable the Gold class to allow Ethernet jumbo frames.

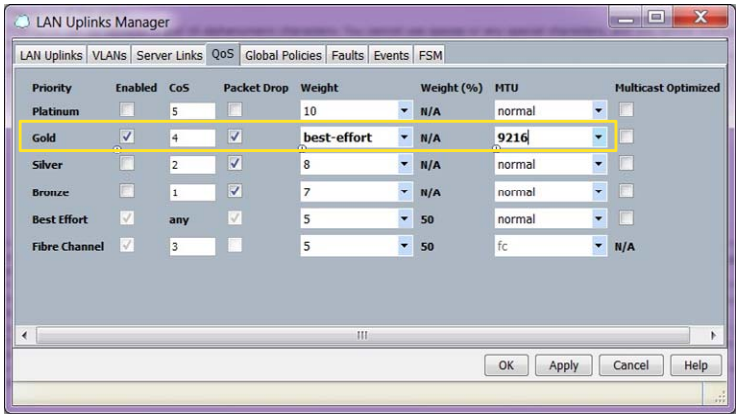

In the example, the Cisco UCS administrator has enabled the Gold QoS system class. In addition, the administrator is configuring this as a member of the "drop class" and setting the relative weighting to "best-effort." These two parameters ensure that, from the perspective of priority flow control and enhanced transmission selection, this traffic receives no special handling. The goal of this policy is limited to enabling jumbo frame support for every virtual network interface card (vNIC) that the policy is applied to.

Note	Disabling drop class and setting a weighting percentage other than best-effort will affect the performance of service profiles that do not include the adapter policy that references this system class.

Locate QoS Policies

Locate QoS Policies

- Expand Policies in the LAN tab, and locate QoS Policies.
- Select the organization where the new policy is to be created and click the **Create QoS Policy** link in the content pane.

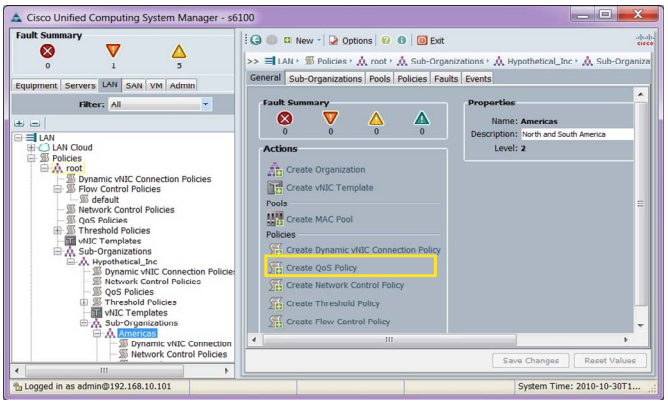

Select the LAN tab of the navigation pane and expand Policies to locate the QoS Policies item in the tree. Select the organization where the policy will reside and click the **Create QoS Policy** link.

Create a New QoS Policy

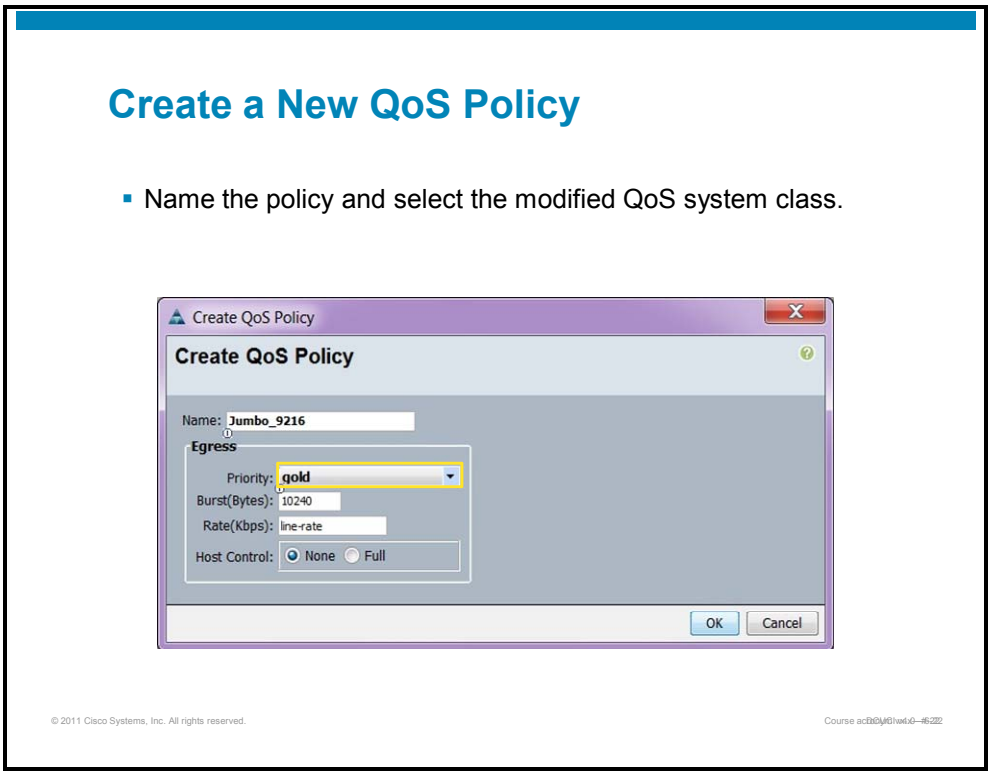

Name the new policy and select the QoS system class "**gold**" from the priority drop-down list. Because the goal of this policy is simply to enable jumbo frame support, leave the burst and rate options at their defaults.

The Host Control option allows you to determine whether this QoS policy can be modified or overridden by the administrator of the hypervisor or operating system. The default is None, which acts as a lockout. Only a Cisco UCS administrator with sufficient privileges in this organization can modify or delete a policy.

QoS Policy Enabled

QoS Policy Enabled

- The new QoS policy is available for assignment.

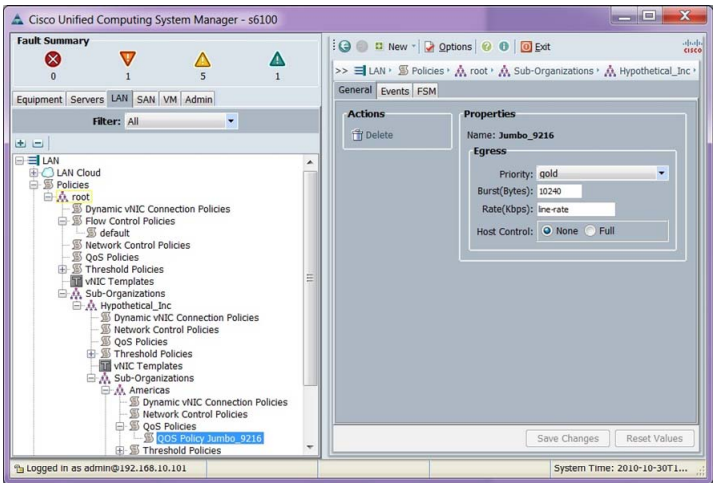

The QoS Policy Jumbo_9216 is now available to be assigned to a vNIC in a service profile.

Configuration of IPMI and SoL Policies

This topic discusses the creation of Intelligent Platform Management Interface (IPMI) and Serial over LAN (SoL) policies that can be consumed by service profiles.

Locate IPMI Policies

Locate IPMI Policies

- Select the organization where the new policy is to be created.
- Click the **Create IPMI Profile** link.

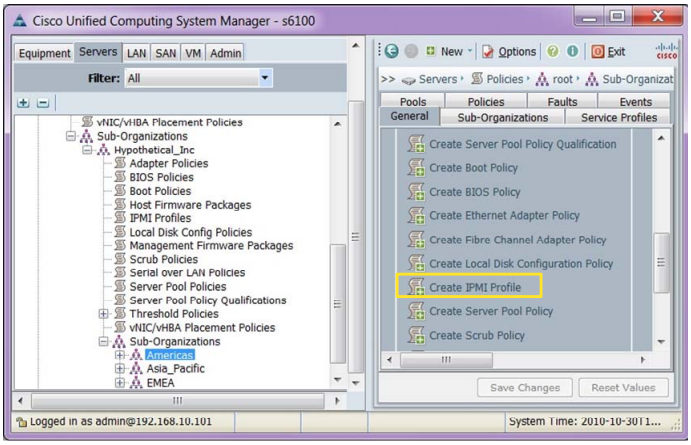

Click the **Server** tab and select the organization where the new IPMI policy will be created. Click the **Create IPMI Profile** link to start the policy wizard.

Create a New IPMI Policy

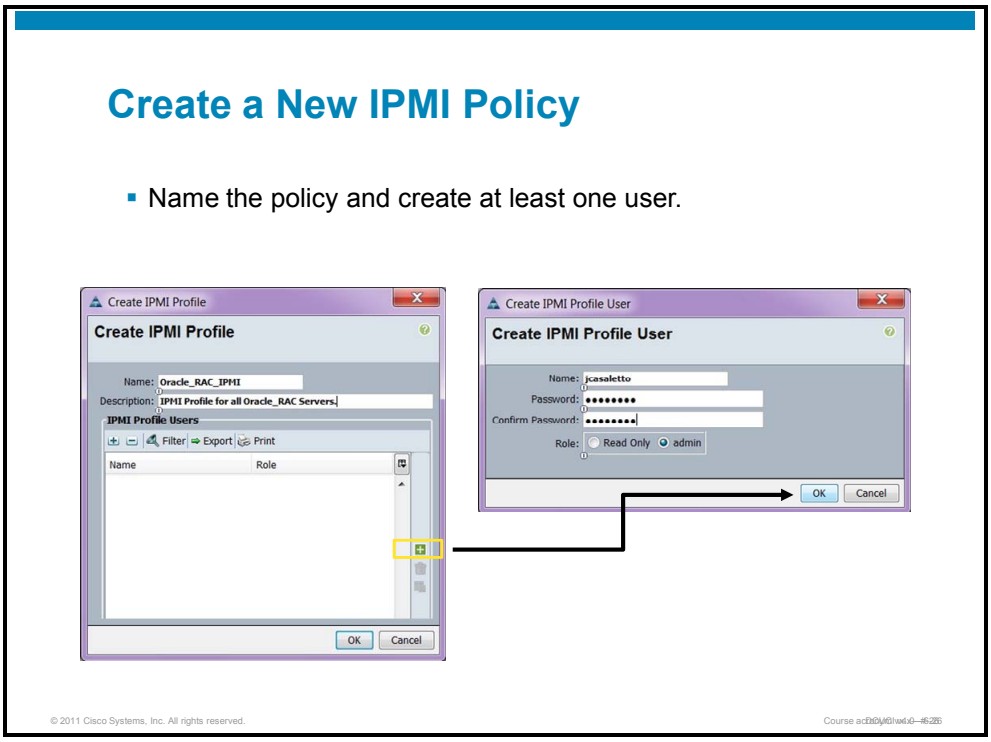

Name the new policy and create at least one user. The role can be either admin or read-only. Read-only users may query the IPMI system that is provided by the Cisco Integrated Management Controller for the status of any IPMI sensor. Admin users can access sensors and additionally perform power control operations of the server.

IPMI Policy “Oracle_RAC_IPMI” Available

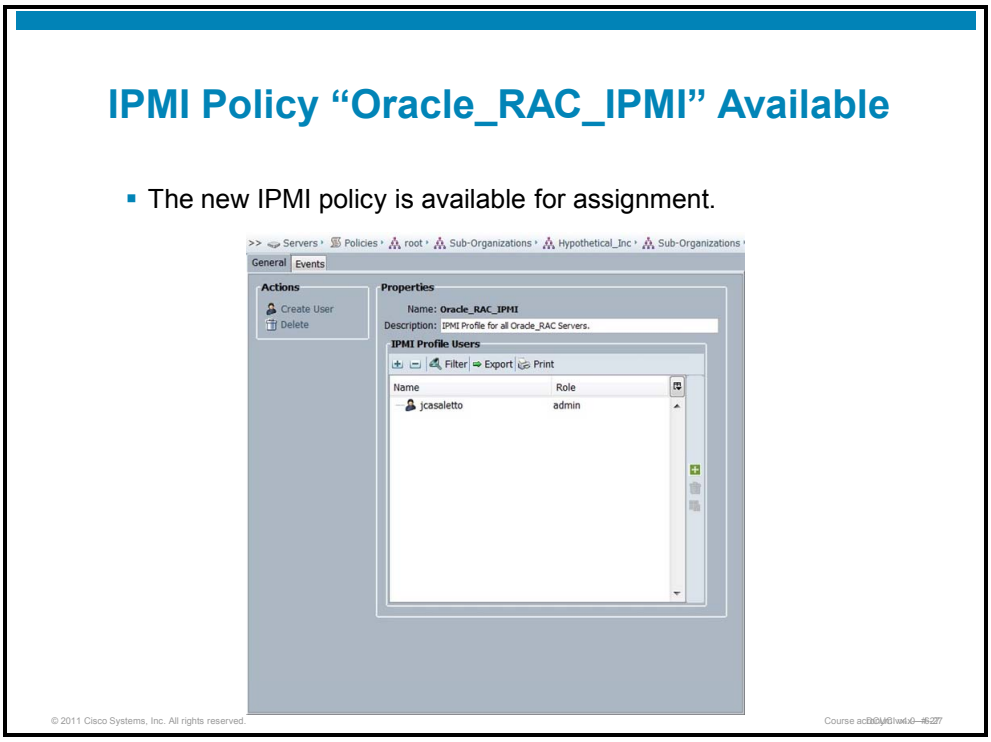

The Oracle_RAC_IPMI policy is now available for assignment to a service profile. User jcasalleto is the sole administrative user.

Note	The users that are created in IPMI policies do not count against the 40-user limit on the local user authentication database.

Locate SoL Policies

Locate SoL Policies

- Expand Policies in the Server tab and locate SoL Policies.
- Right-click on the organization where the new policy is to be created.

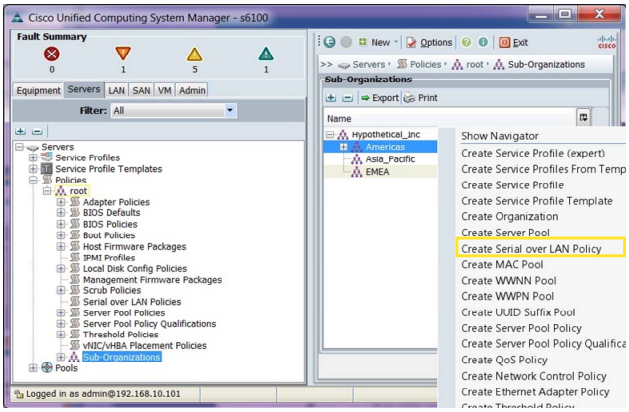

Expand the Server tab in the navigation pane and locate Serial Over LAN policies. In the content pane, right-click on the organization where the policy should be created.

Create a New SoL Policy

Create a New SoL Policy

- Provide a name and optional description for the policy.
- Set administrative state and serial baud rate.

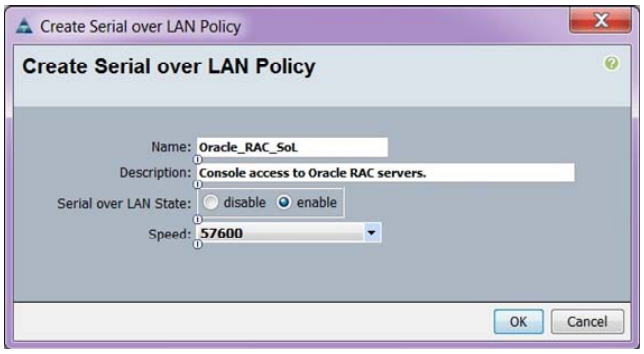

Name the new SoL policy, set the administrative state to **Enabled**, and set the serial baud rate so that the connection will communicate. SoL connections use UDP port 623.

SoL Policy “Oracle_RAC_SoL” Available

SoL Policy “Oracle_RAC_SoL” Available

- The new SoL policy is available for assignment.

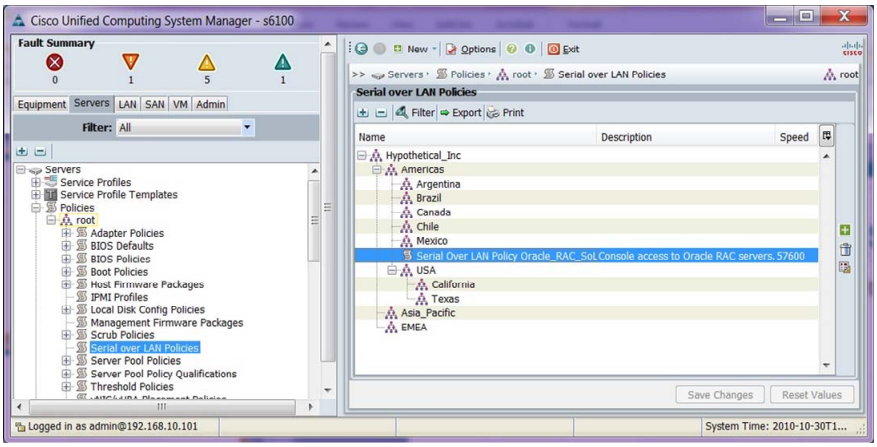

The new SoL policy is available for assignment to a service policy.

Configuration of a Scrub Policy for Local Disks and BIOS

This topic describes the configuration of scrub policies for local hard disks in the server BIOS.

Locate Scrub Policies

Locate Disk Scrub Policies

- Expand Policies in the Server tab and locate Scrub Policies.
- Click **Create Scrub Policy** in the organization where the new policy is to be created.

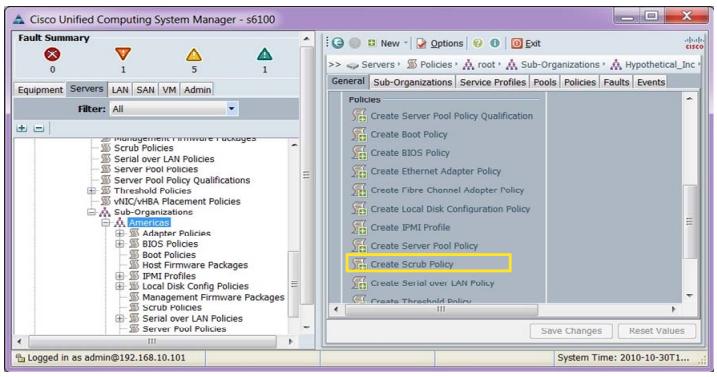

Expand Policies in the Server tab and select the organization where the policy will be created. Click **Create Scrub Policy** to begin the wizard.

Create a Disk Scrub Policy

Create a Disk Scrub Policy

- Provide a name and optional description for the policy.
- Decide whether to scrub only disks, or disks and BIOS settings.

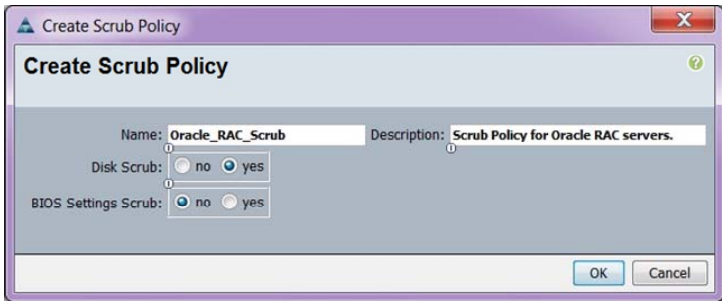

Name the policy and, optionally, provide a description. Previous versions of Cisco UCS Manager did not include the BIOS scrubbing capability of the new scrub policies. Cisco UCS Manager version 1.3 introduces this new combined scrub policy.

When disk scrub is set to Yes, local disk drives will be completely erased upon disassociation of a service profile.

When the BIOS settings scrub policy is set to Yes, all BIOS settings will revert to factory defaults upon disassociation of a service profile.

Scrub Policy “Oracle_RAC_Scrub” Available

Scrub Policy “Oracle_RAC_Scrub” Available

- The new scrub policy is available for assignment.

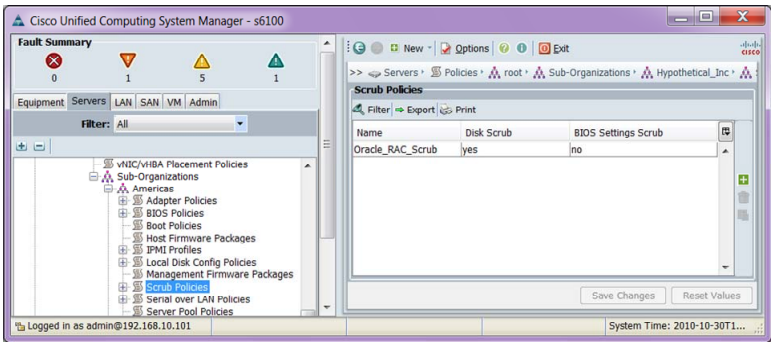

The new scrub policy is available for assignment to service profiles.

Simple vs. Expert Service Profile Wizards

This topic discusses the differences between the simple and expert service profile wizards.

Service Profile Wizards

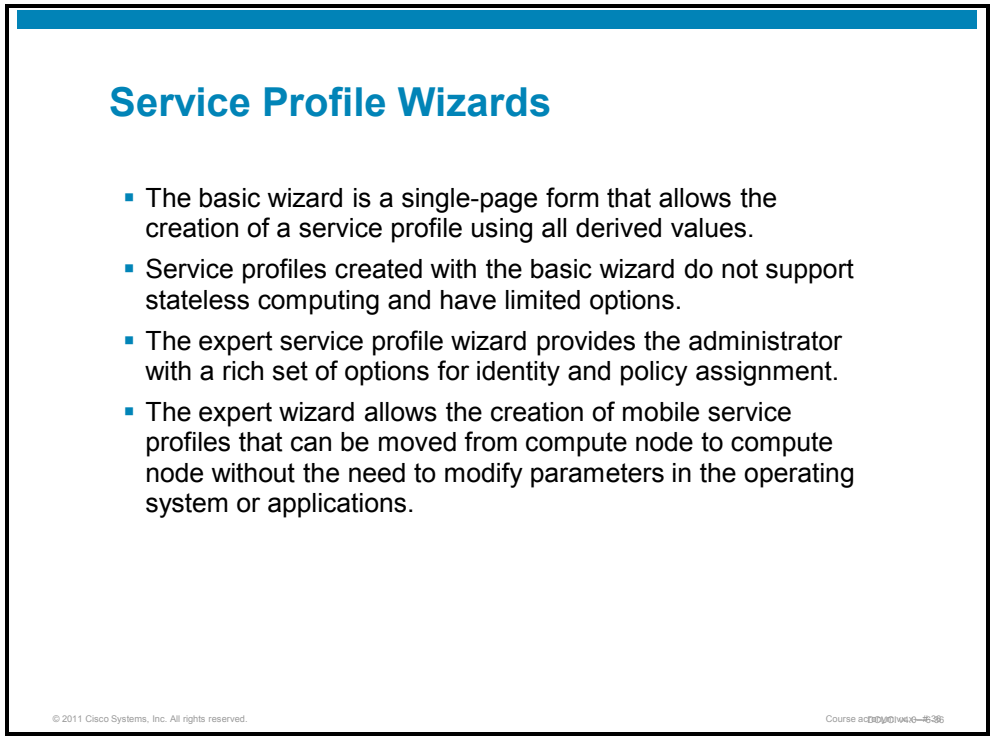

The primary difference between the simple service profile wizard and the expert service profile wizard is the scope of tools available to manipulate within the wizard. The simple wizard provides a fast, single-page form for the rapid provisioning of a blade server using all derived identity values. The expert service profile wizard allows for the granular configuration of policies, identities, and thresholds.

Comparison of Service Profile Wizards

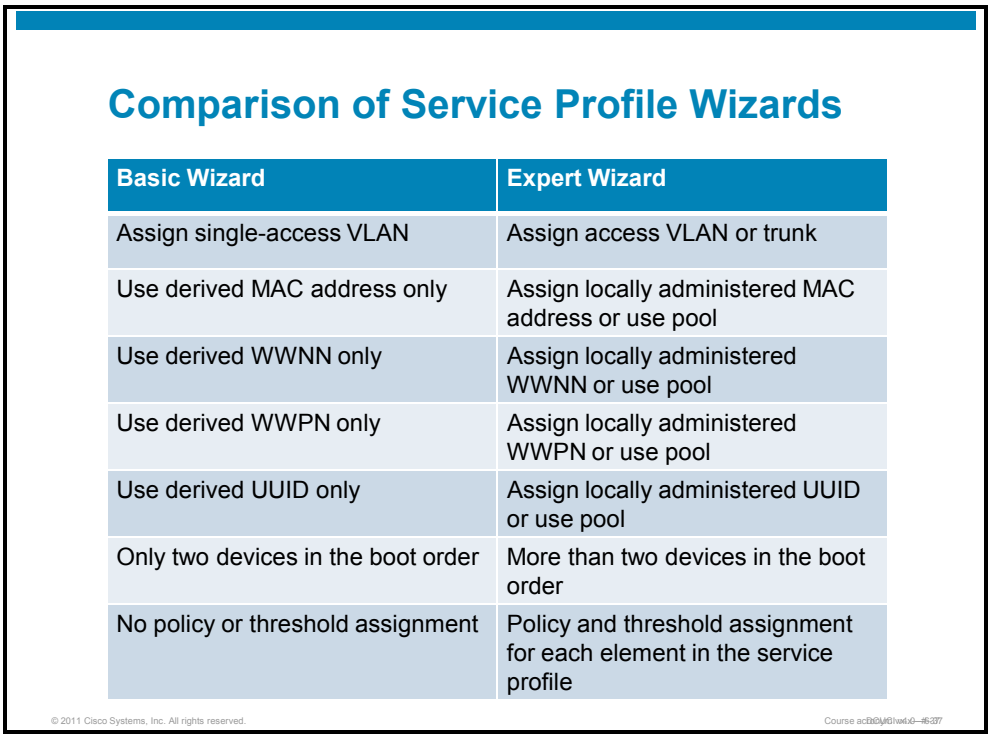

The table summarizes the most important differences between the simple and expert service profile wizards.

Service Profile Expert Wizard

This topic discusses the configuration of the service profile using the expert service profile wizard.

Launch Expert Service Profile Wizards

Launch Expert Service Profile Wizards

- Select the organization where the new service profile should be created.

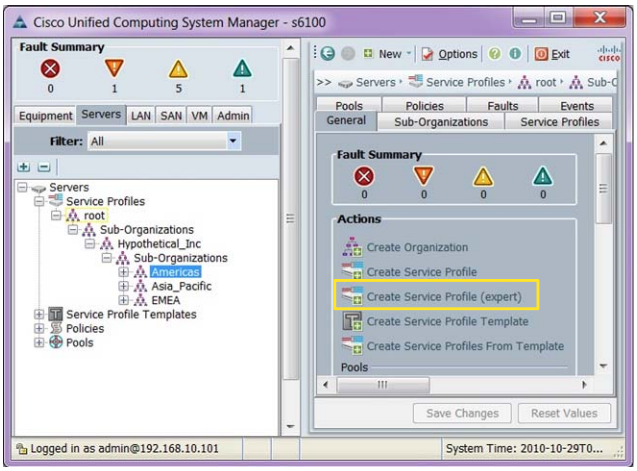

From the Server tab in the navigation pane, select the organization for which a new service profile will be created. In the content pane, click the link for **Create Service Profile (expert)**.

Name the Service Profile

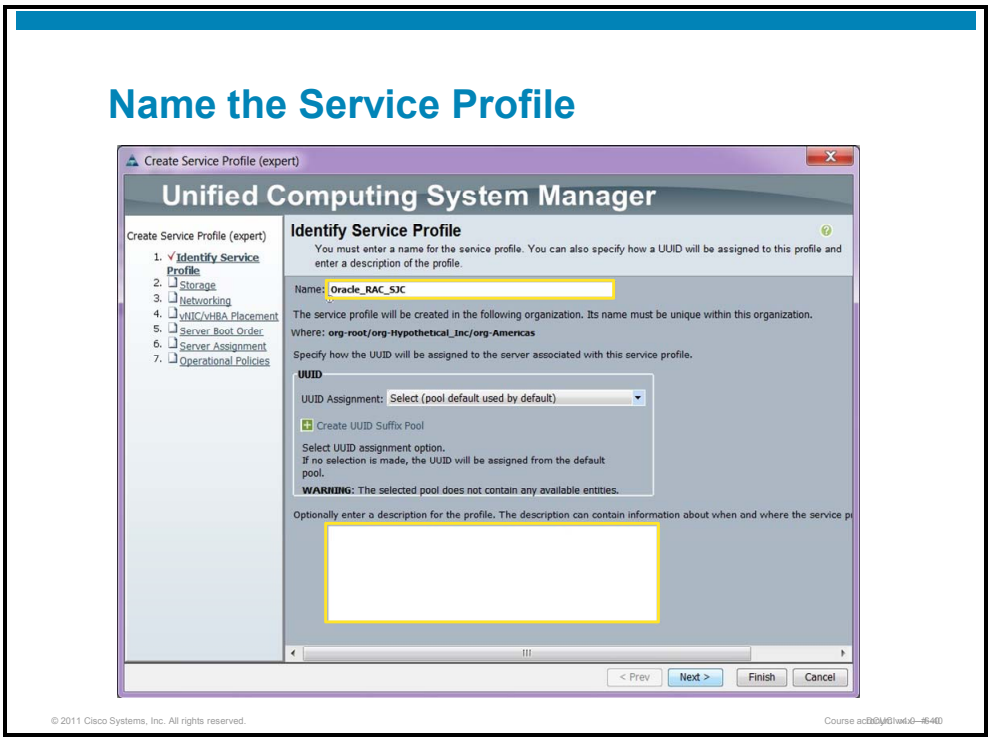

Next, name the service profile and, optionally, provide a description. The name must not exceed 16 characters and may not include special characters or spaces.

Configure the Service Profile to Take Its UUID from a Pool

This topic discusses how to assign UUIDs from a pool.

Create UUID Pool in Service Profile Wizard

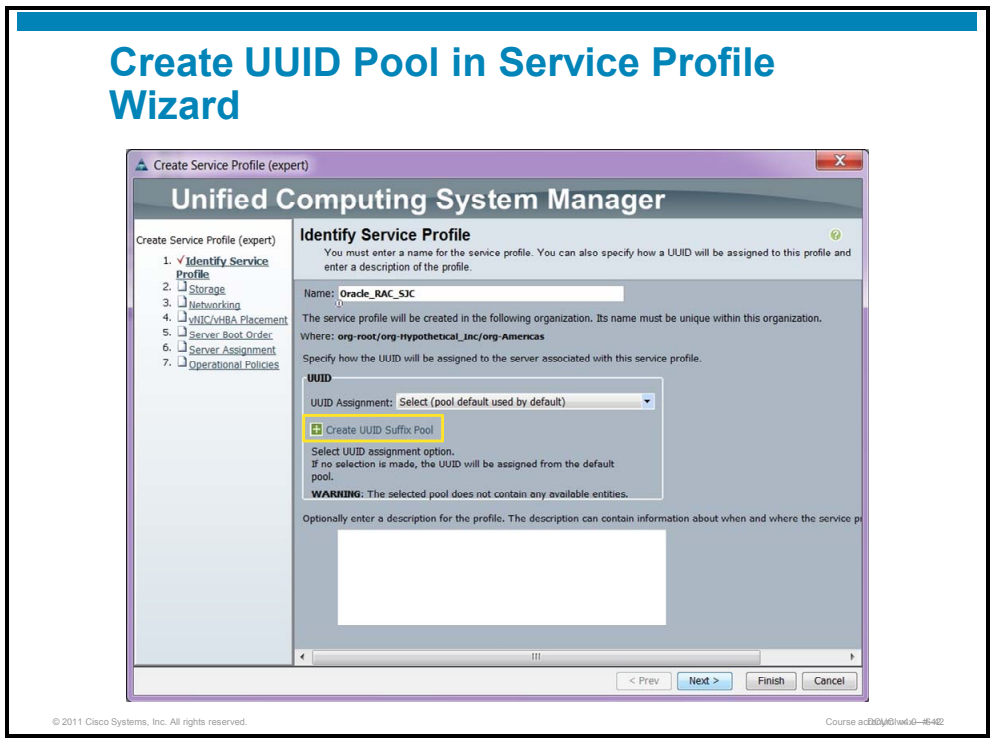

If you begin the service profile wizard and realize that you have forgotten to first create a UUID pool, you can create the new pool from within the wizard. Some very useful capabilities are available throughout the expert wizard for pooled values.

Assign UUID from a Pool

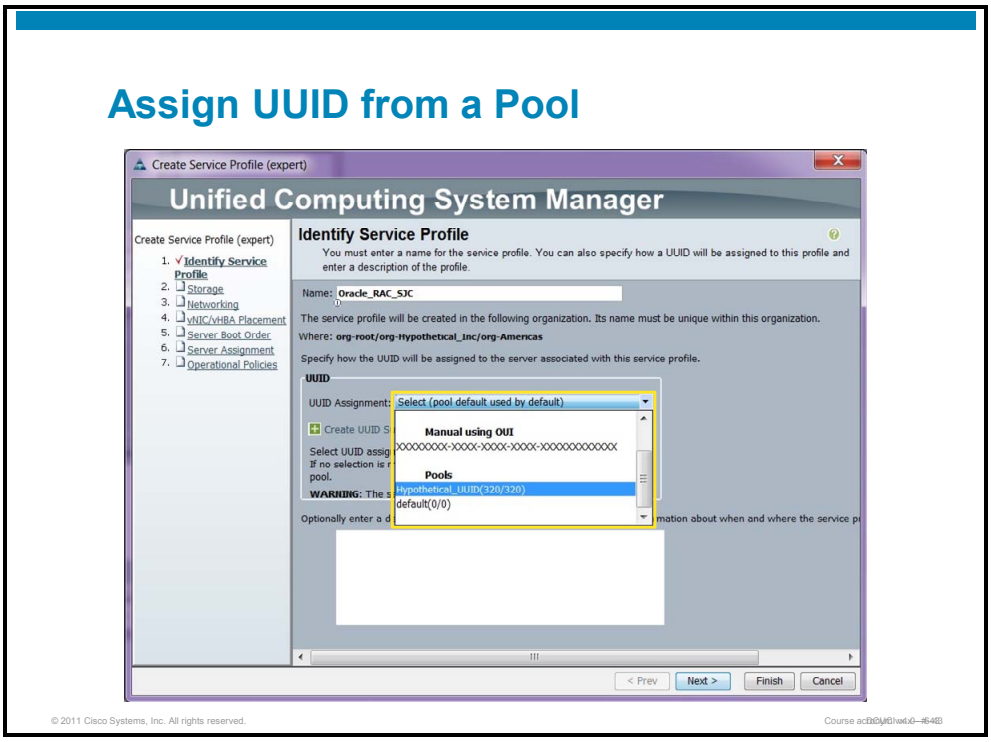

You can select the UUID pool directly from the drop-down list. After the assignment is made, click **Next** to continue the wizard.

Configuration of vHBAs

This topic discusses the creation of virtual host bus adapters (vHBAs) that derive their identity from pools.

Create WWNN Pool Within the Service Profile

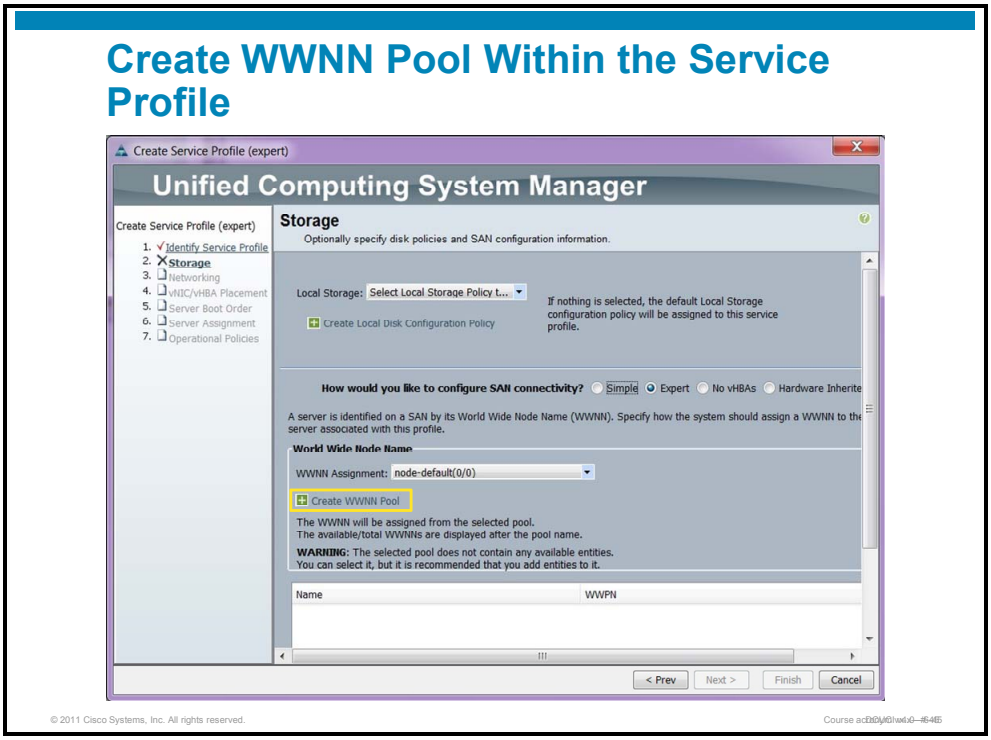

The figure shows an example of creating a pool from within a wizard.

Assign WWNN from a Pool

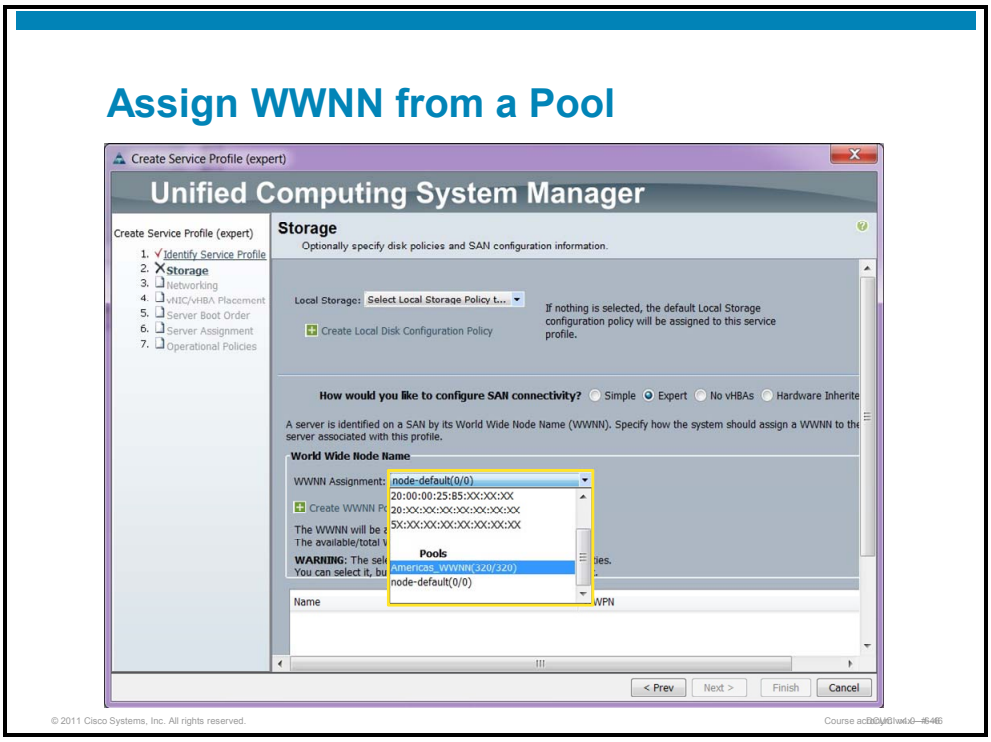

Since Americas_WWNN pool was defined and populated with values, it can be directly referenced from the drop-down list. The WWNN refers to the mezzanine card and there is only one WWNN per service profile unless there are two mezzanine cards that are populated in a full-slot blade. Click **Next** to continue the wizard.

Begin Creation of vHBAs

Begin Creation of vHBAs

- Click the **+ Add** button to create the first vHBA.

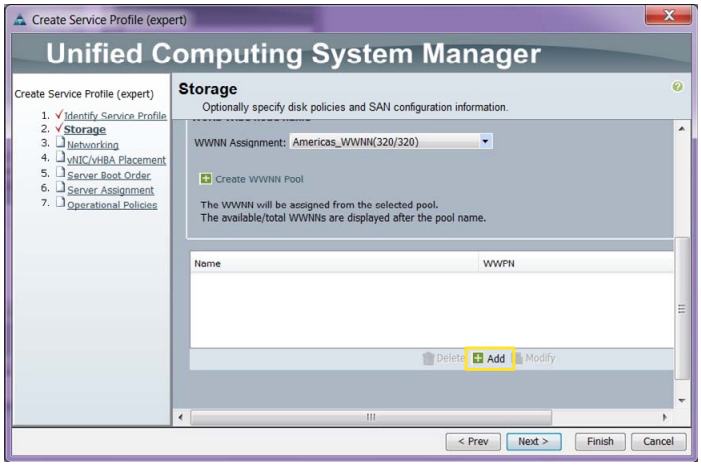

Click the plus sign to open a dialog box to create the vHBA for fabric A.

Create vHBA for Fabric A

Create vHBA for Fabric A

- Define the configuration parameters for fabric A.

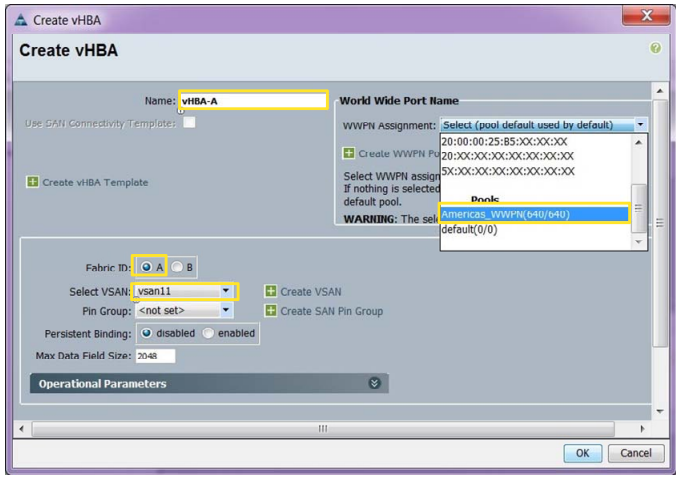

The vHBA requires a name, WWPN assignment, VSAN assignment, and fabric assignment. The example creates a new vHBA named vHBA-A. It will pull its WWPN assignment from Americas_WWPN pool, is a member of VSAN 11, and is associated with fabric A.

Create vHBA for Fabric B

Create vHBA for Fabric B

- Define the configuration parameters for fabric B.

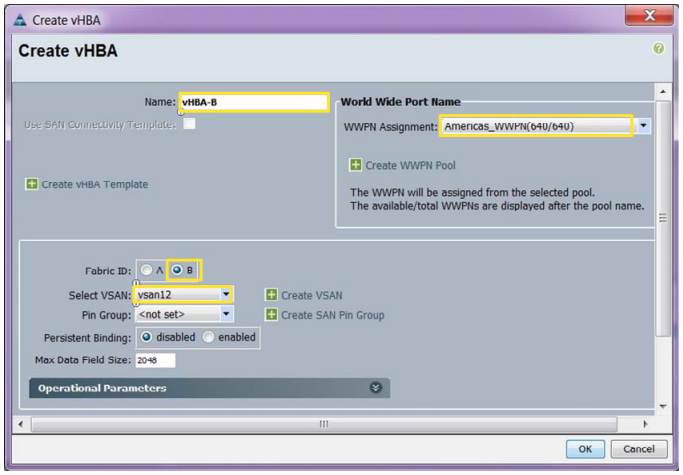

Repeat the steps for creating the vHBA on fabric A to create the vHBA on fabric B. Be certain to select a different VSAN for fabric B. Click **Next** to continue the wizard.

Configuration of vNICs

This topic discusses the configuration of virtual network adapters in the assignment of identity from pools.

Simple View for Networking

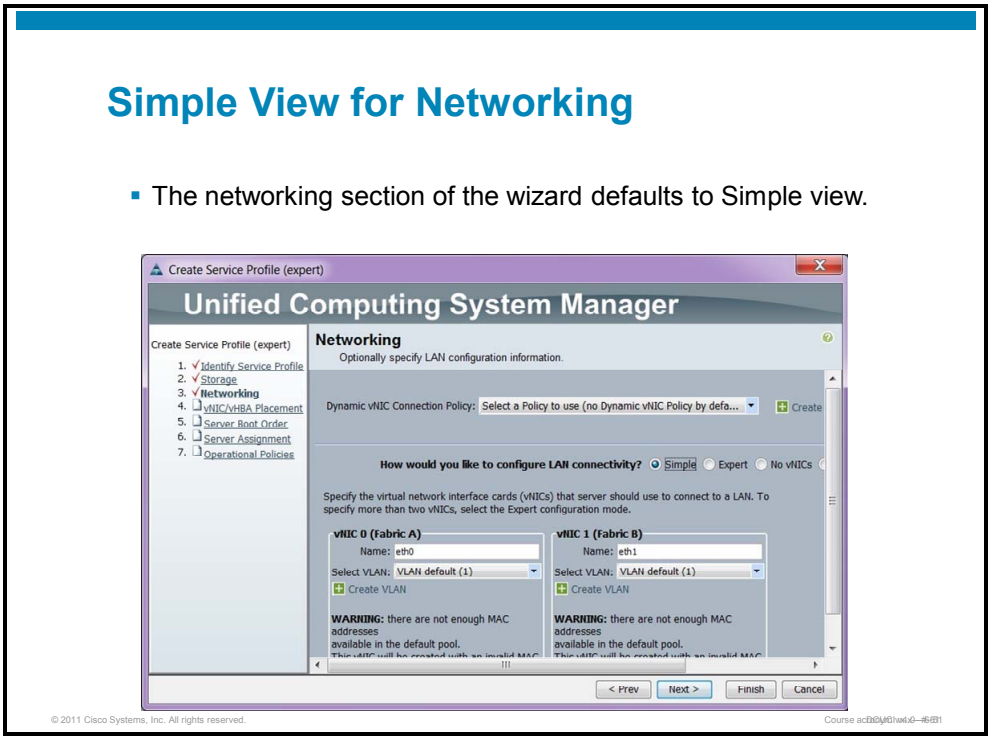

The expert wizard defaults to simple view on the opening page of networking configuration. This view limits you to selecting derived MAC addresses and a single VLAN.

Note	The MK81-KR virtualization adapter does not have burned-in MAC addresses. Pooled or manual address assignment is required on this mezzanine adapter.

Switch to Expert View for Networking

Switch to Expert View for Networking

- Click the **Expert** radio button for greatest control.

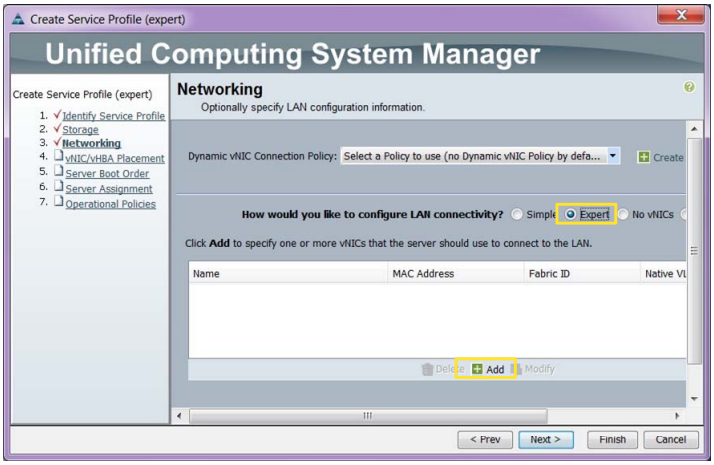

Click the **Expert** radio button to reveal the complete suite of networking configuration tools available within the expert wizard. Click the **+Add** button to open the dialog box and define a new virtual network card.

Define Properties of vNIC for Fabric A

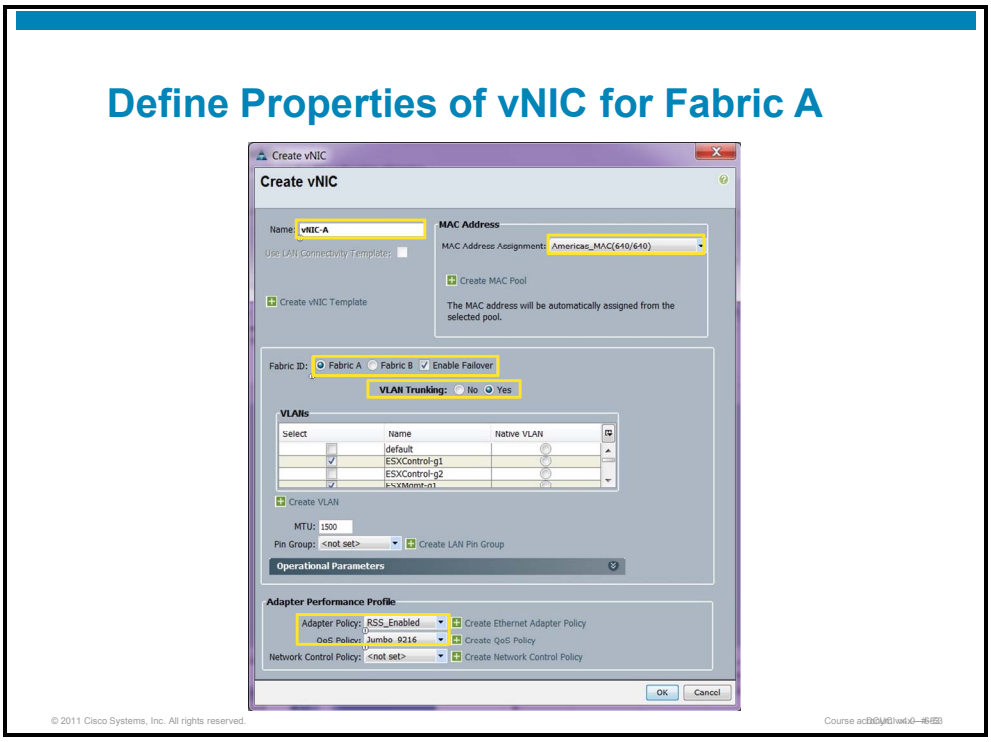

The new vNIC requires the following configuration elements:

- Name
- MAC assignment
- Fabric assignment and failover
- Configuration of access to the LAN or trunk
- Policies

This vNIC, named vNIC-A, will have its MAC address assignment from the Americas_MAC identity pool. Because the mezzanine card that this vNIC will be associated with is an MK81-KR, it is configured for hardware-based fabric failover.

A VLAN trunk with two tagged VLANs will be provided to the hypervisor.

The adapter and QoS policies enable RSS, reduced failback window, and jumbo frame support.

Define Properties of vNIC for Fabric B

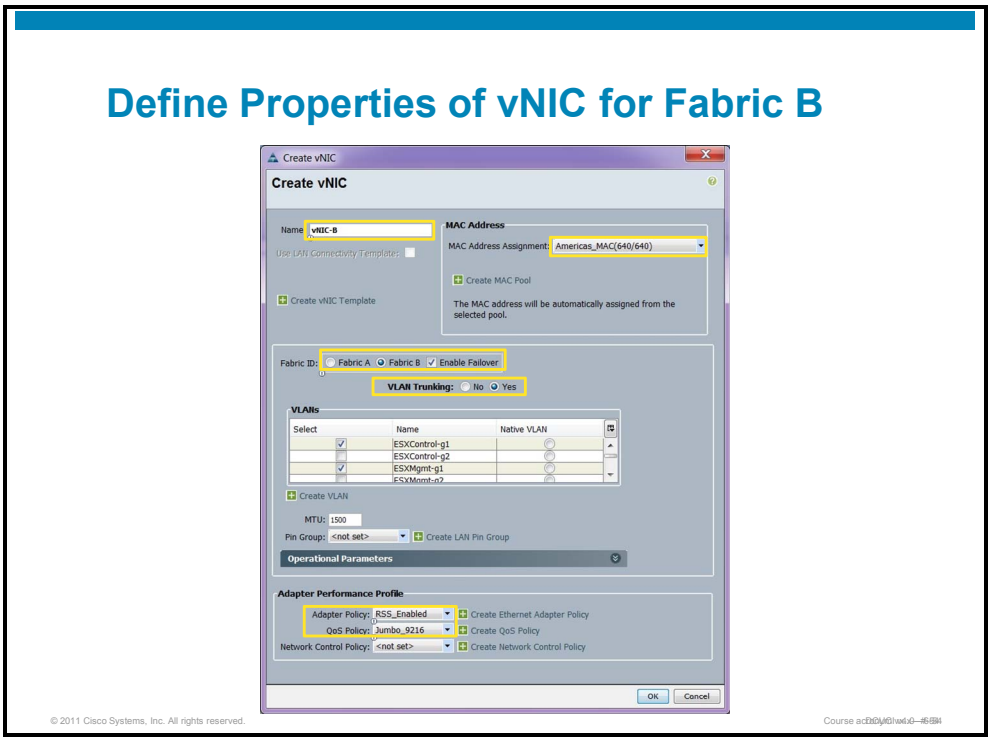

Repeat the process to create the vNIC for fabric B. It requires a unique name and it will be assigned to fabric B. All other parameters will be identical to the vNIC for fabric A.

vNIC Configuration Summary

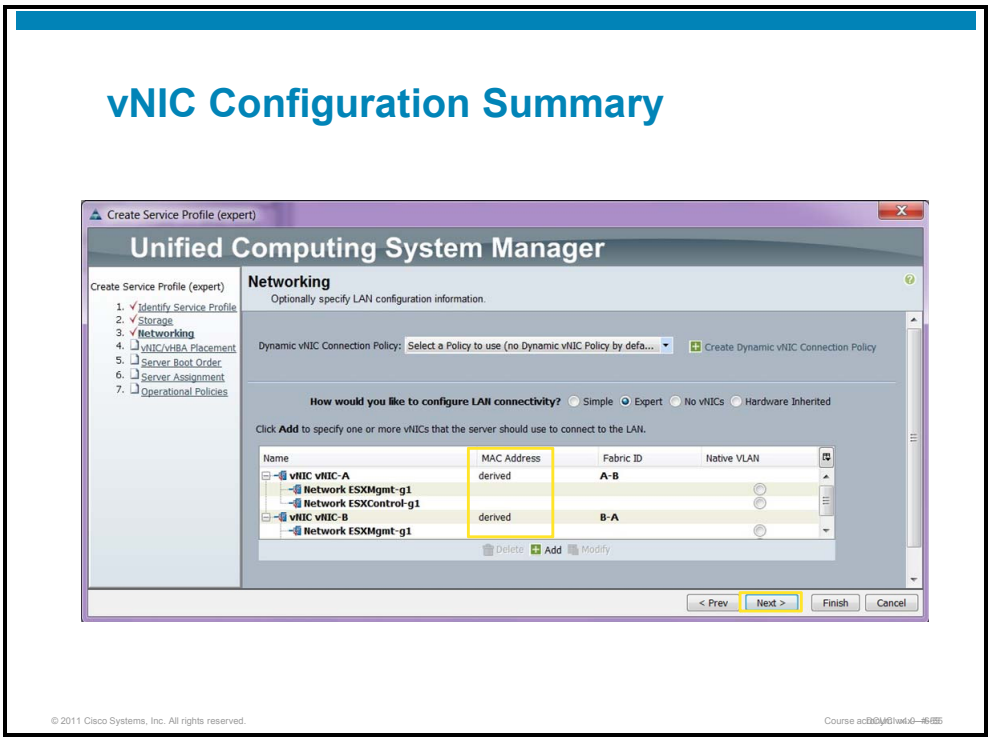

The vNIC summary window is used to validate the configuration of the newly created virtual interface cards. Click **Next** to continue the wizard.

Note	The MAC address assignment indicates “derived” because the actual assignment of the address will not occur until the service profile wizard has completed.

vNIC and vHBA Placement on Full-Slot Blades

This topic discusses vNIC placement on blade servers with dual mezzanine cards.

Dual Mezzanine vNIC Placement

- Full-slot blades have two mezzanine cards.
- vNIC placement can be automatic or manual.

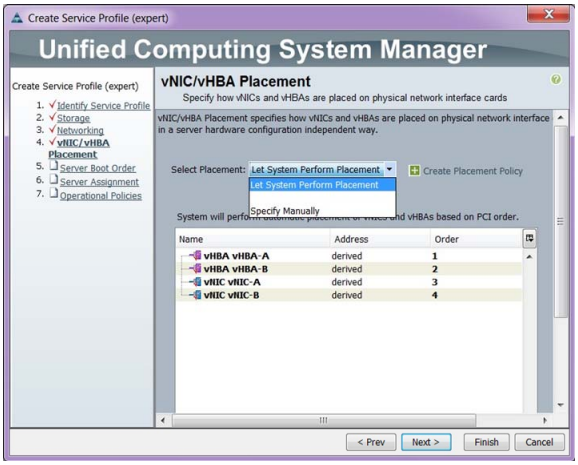

The Cisco UCS B250 and B440 full-slot blade servers include two slots for mezzanine cards. Because a vNIC is a virtual definition of a network interface card, it could be placed on the appropriate fabric on either of the mezzanine cards present in the full-slot server.

In a half-slot blade with a single mezzanine card, simply allow the system to select the only mezzanine card. If manual control is desired, select **Specify Manually** from the Select Placement drop-down list. vCon1 maps vNICs to the left mezzanine slot, and vCon2 maps a vNIC to the right mezzanine slot (as viewed from the front panel of the blade server). Click **Next** to continue in the wizard.

Binding of a vHBA to a Fibre Channel Boot Target

This topic discusses finding boot targets to vHBAs.

Boot Order Creation

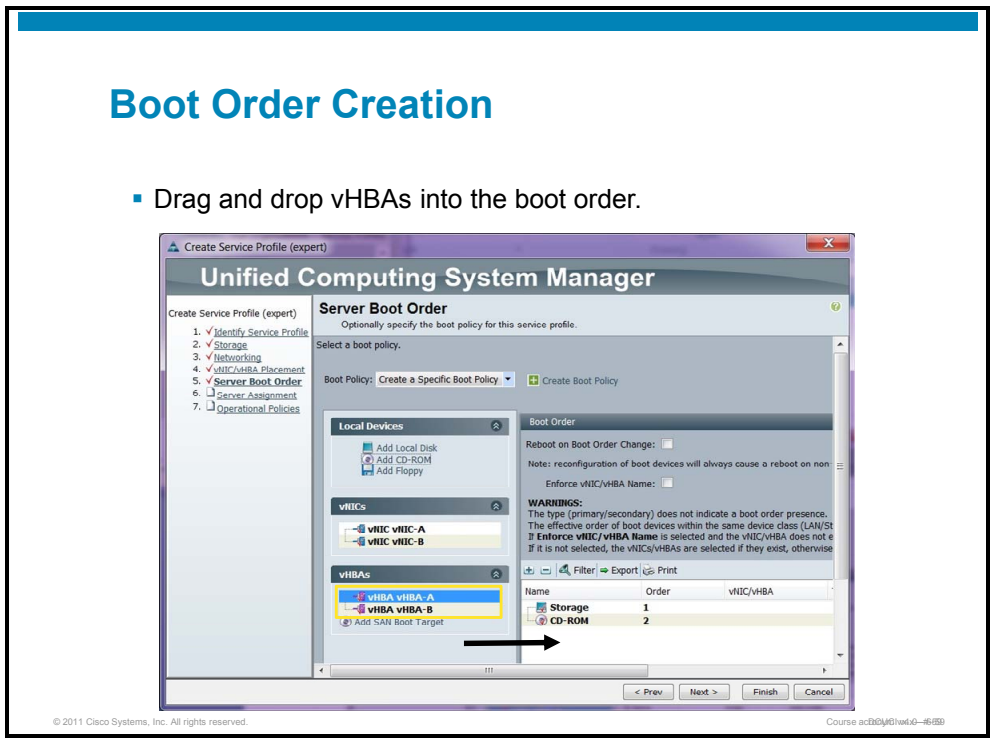

To select vHBAs to boot from a logical unit number (LUN), click and drag the first vHBA into the boot order whitespace.

Drag and Drop Pop-Up for Fabric A

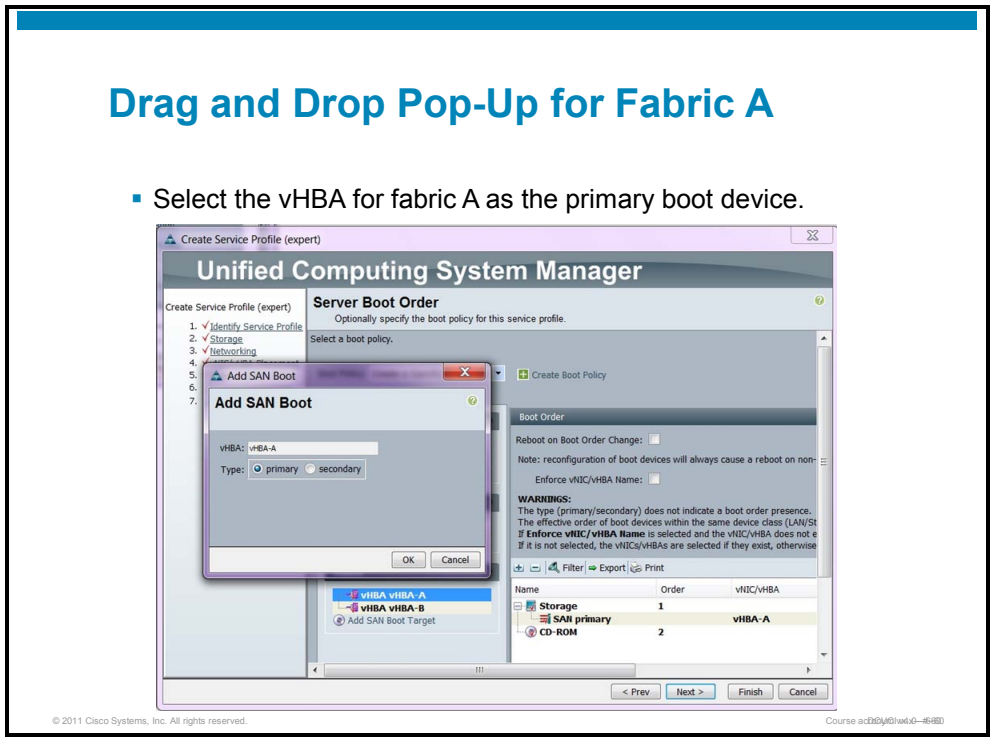

A pop-up window will appear with the name of the vHBA and the choice to make this the primary or secondary boot device. Select **Primary** for the vHBA on fabric A, then click **OK**.

Drag and Drop Pop-Up for Fabric B

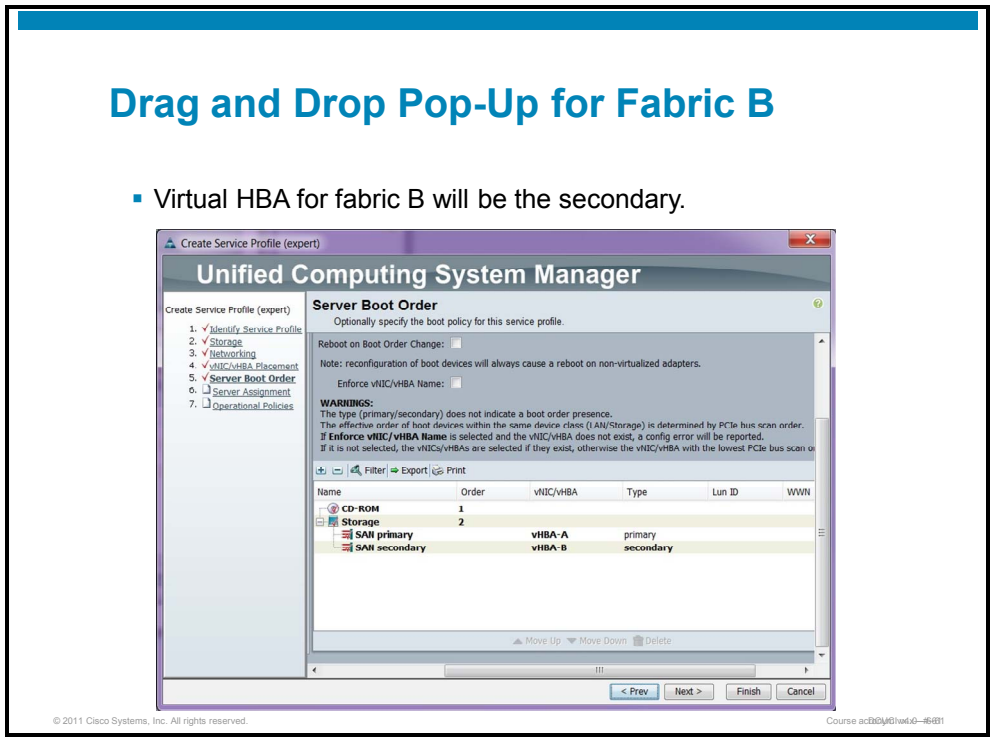

A pop-up window will appear with the name of the vHBA and the choice to make this the primary or secondary boot device. Select **Secondary** for the vHBA on fabric B, then click **OK**.

Set Boot Target Menu

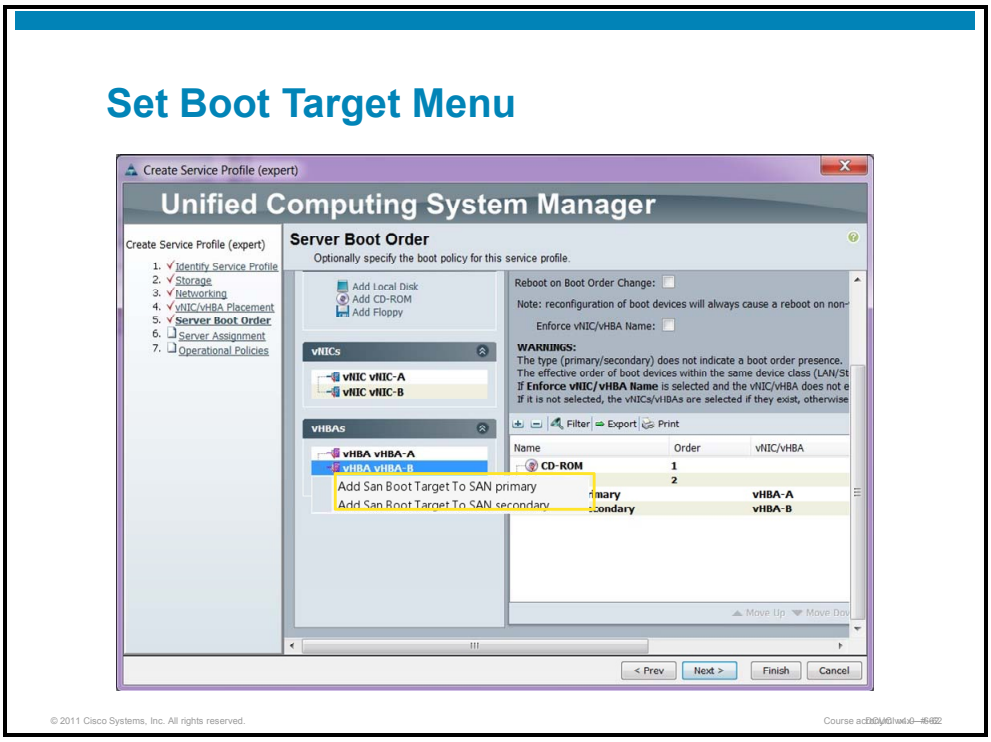

Click the Add SAN Boot Target below the vHBAs and select **Add SAN Boot Target to SAN Primary**.

Set Boot Primary A

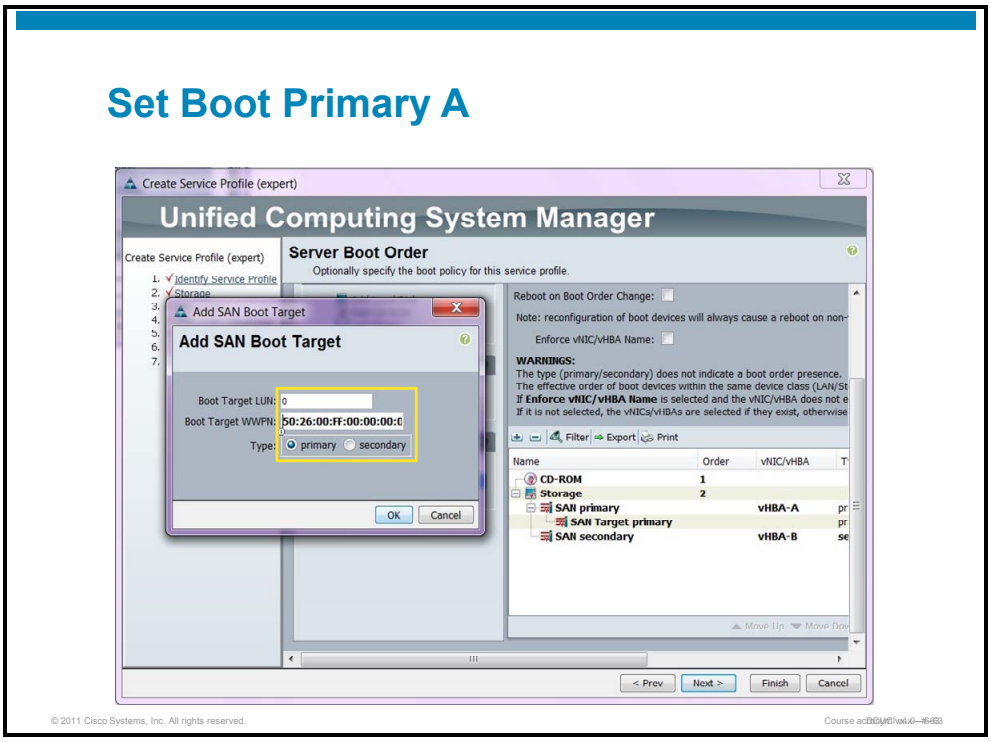

In the pop-up window, enter the boot LUN (always LUN 0 on Cisco UCS systems), the WWPN of the boot target, and set the type to **Primary**.

Set Boot Primary B

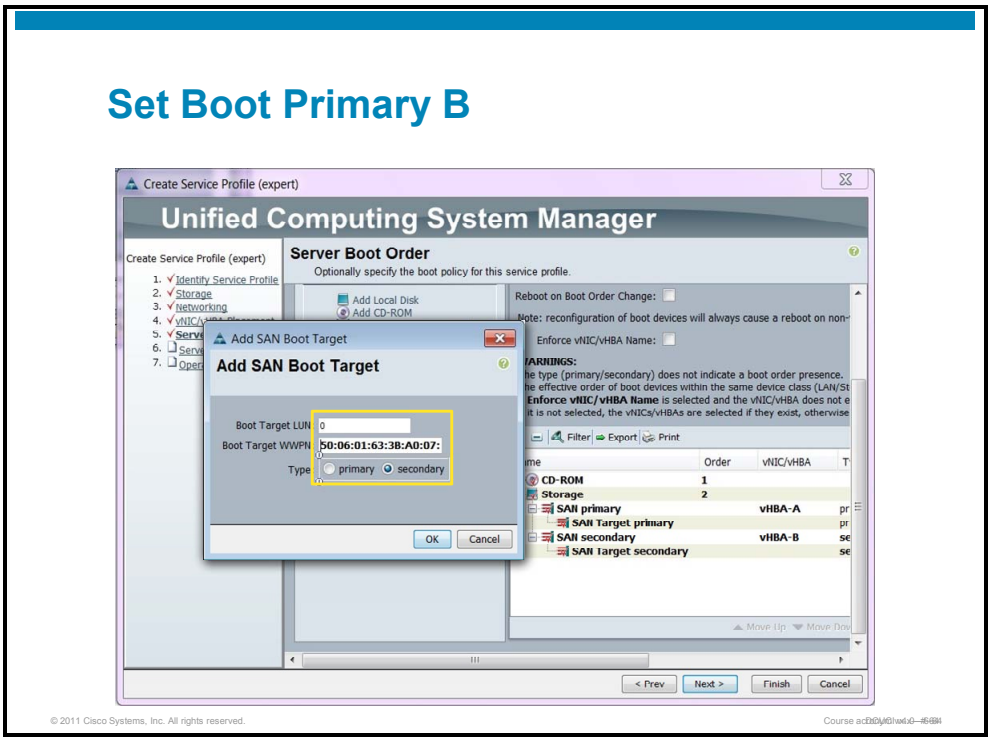

Repeat the steps that are required to set the primary, but set the type to **Secondary**. In the event that the primary boot device fails, the secondary device will attempt to boot the system from the other vHBA.

Boot Order Summary

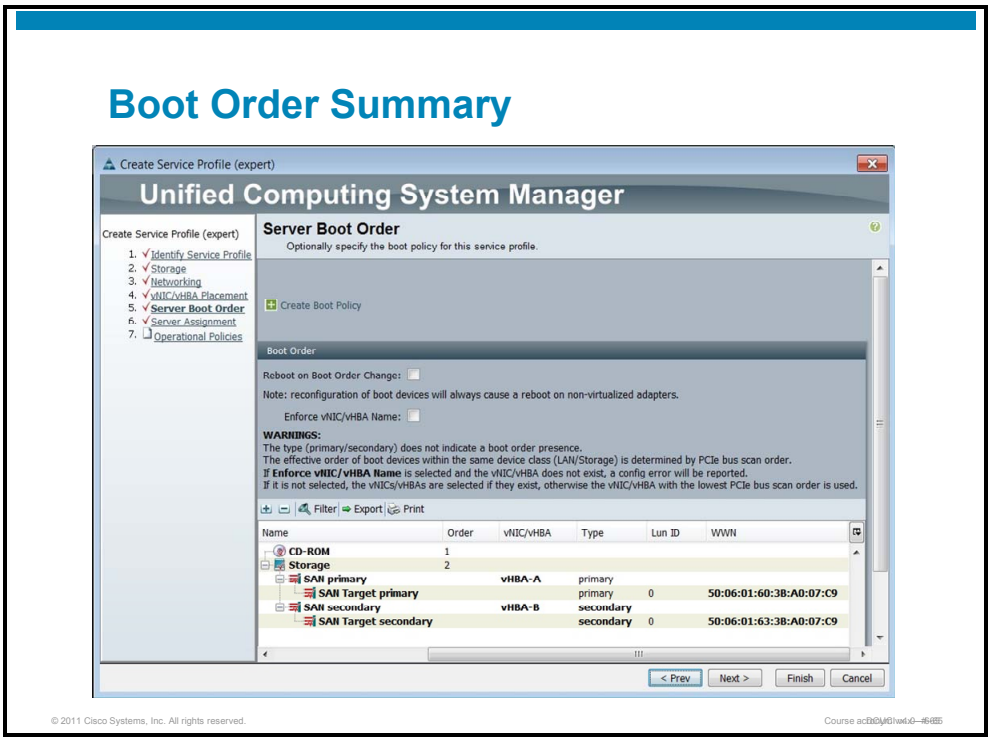

After boot devices are configured, the boot order summary window allows you to verify and make modifications to the boot order before committing the configuration.

There are two checkboxes:

- **Reboot on Boot Order Change:** Requires that the blade associated with the service profile reboot immediately.
- **Enforce vNIC/vHBA Name:** Means that the system uses any vNICs or vHBAs in the order that is shown in the Boot Order table. If not checked, the system uses the priority that is specified in the vNIC or vHBA.

Note	If the configuration of a vHBA is changed (other than the boot order), the system will immediately reboot.

Server Assignment

Server assignment is one of the final configuration decisions within the expert service profile wizard. The Cisco UCS administrator can simply finish the wizard and manually assign the service profile at a later time, or manually select a server from the list of unassociated servers. Service profiles can only be associated with a server that does not have a profile that is actively associated with it.

Assign Service Profile to Server from a Pool

Assign Service Profile to Server from a Pool

- Select a server pool from the drop-down list.

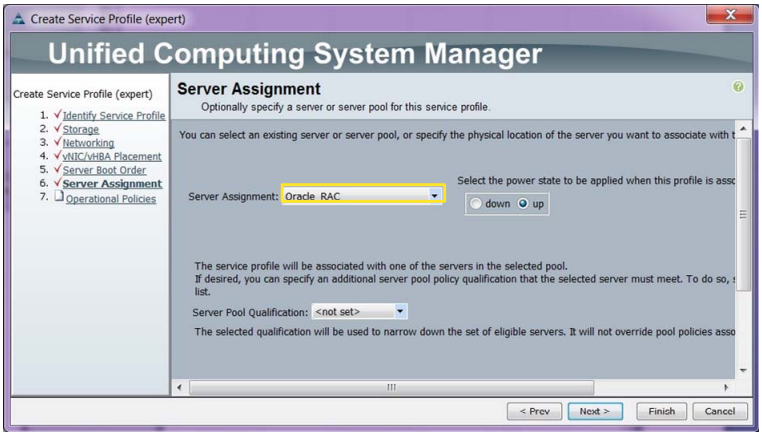

From the Server Assignment drop-down list, select an available pool. Cisco UCS Manager will remove that server from all of the pools of which it is currently a member and associate the service profile with that blade. Click **Next** to continue the wizard.

Assign Service Profile Directly

Assign Service Profile Directly

- Select a server from all unassociated server blades.

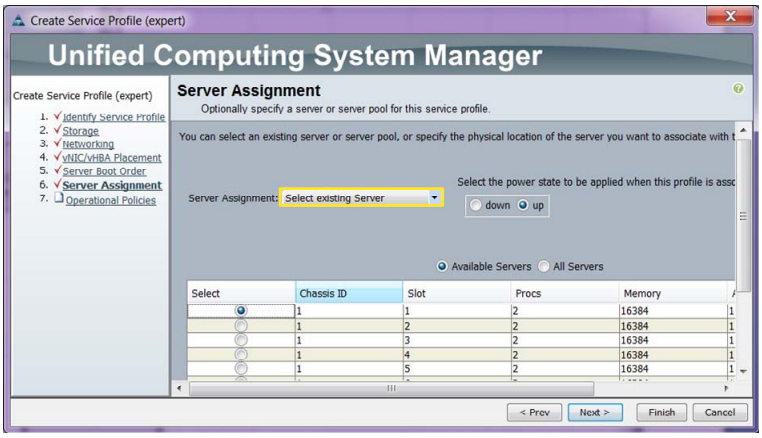

Choose **Select Existing Server** from the Server Assignment drop-down list. Click the radio button of an unassociated server.

The power state radio button allows you to choose the initial power state after the service profile has been successfully associated with the blade server. If the SAN team has not provisioned boot LUN in time for the service profile, you should leave the power state down.

Assign IPMI and SoL Policies to the Service Profile

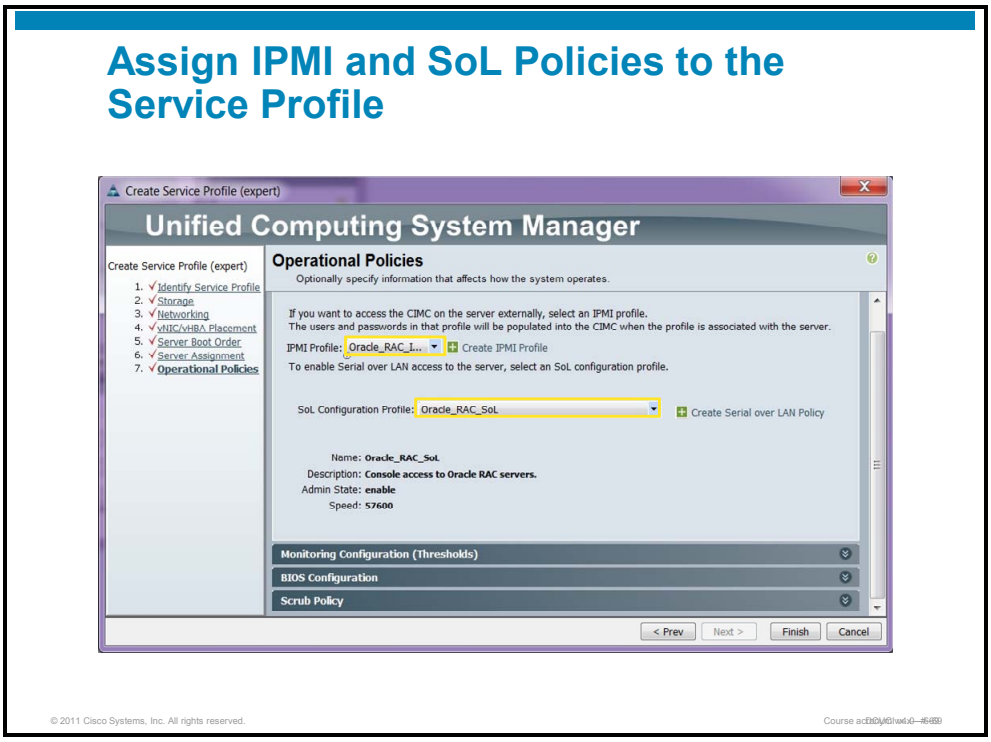

Operational Policies is the last page of the wizard. Use the drop-down lists to select the IPMI and SoL policies.

Assign BIOS and Scrub Policies to the Service Profile

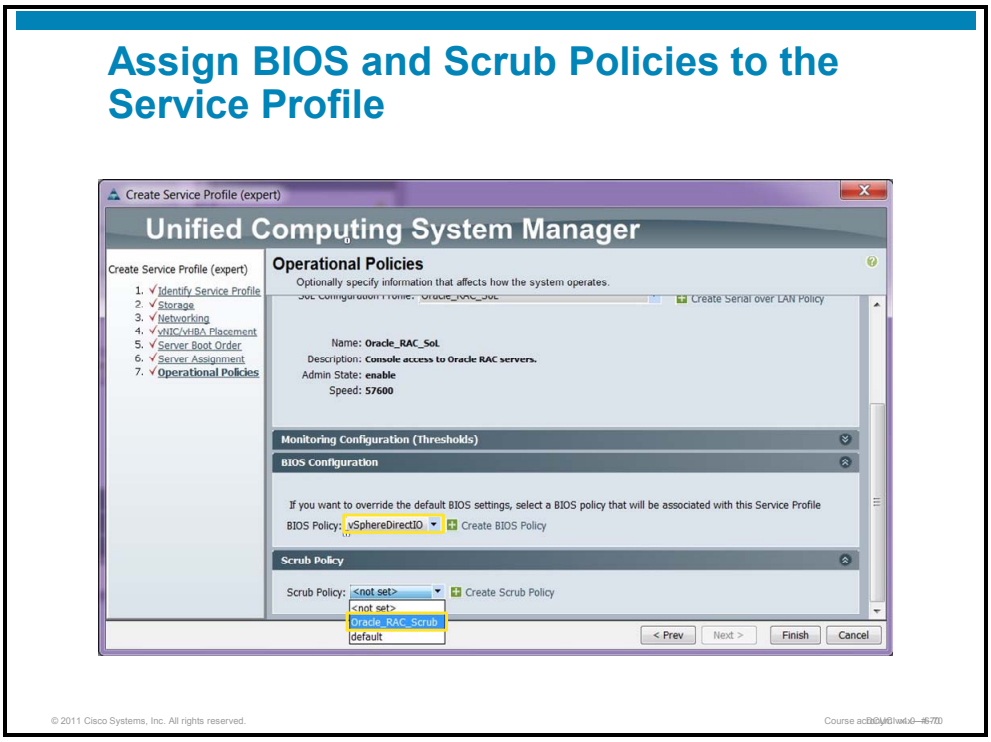

While still on the Operational Policies page of the wizard, expand the BIOS Configuration and Scrub Policy subwindows. Use the drop-down lists to select the BIOS and scrub policies.

Required vs. Optional Components of the Service Profile Definition

This topic differentiates between required elements and optional elements of the service profile.

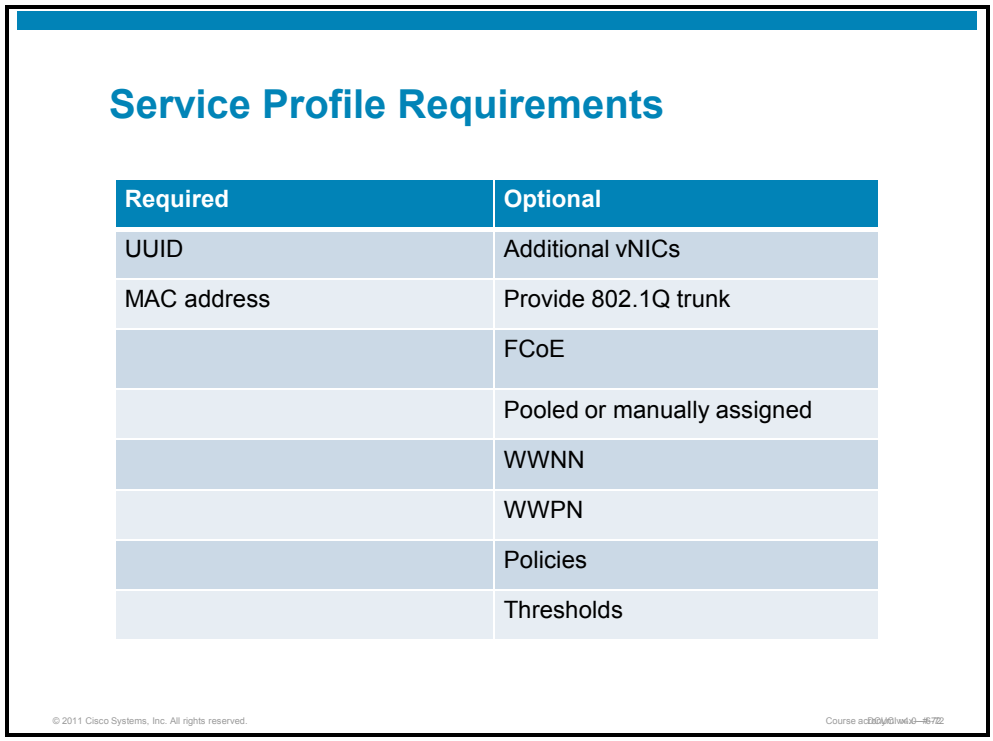

The table summarizes required and optional elements of all service profiles.

Summary

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

Summary

- A blade server requires a service policy to achieve external communication through the mezzanine card.
- Cisco UCS Manager version 1.3 includes manipulating BIOS settings within a BIOS policy in the service profile.
- Adapter policies allow configuration of RSS, checksum offloading, failback timer, and transmit and receive buffers.
- The LAN Uplinks Manager allows the modification of QoS system classes to tune bandwidth priority, lossless fabric, multicast, and MTU.
- IPMI and SoL policies are applied to service profiles to allow external access to the Cisco Integrated Management Controller and serial console.
- Scrub policy for local disks and BIOS can be applied to a service profile that allows local disks and BIOS settings to be erased upon disassociation.
- The expert service profile wizard allows complete control over the assignment of identity, policy, and thresholds.

Summary (Cont.)

- The expert service profile wizard is initiated from the Server tab in the navigation pane.
- UUID can be assigned from a pool, manually assigned, or derived from the server BIOS.
- WWNN and WWPN assignment can be performed from a pool, manually assigned, or (depending on the mezzanine model) derived from hardware.
- MAC address assignment can be performed from a pool, manually assigned, or derived from hardware.
- Full-slot blade servers include two mezzanine slots in the service profile, and offer manual or automatic selection binding vNICs and vHBAs to a slot.
- You must configure the binding of a vHBA to a Fibre Channel boot target.
- Server assignment can be directly selected from a list of unassociated servers, assigned at a later time, or signed from a pool.
- UUID and MAC address assignments are the only required elements in a service profile.

Lesson 3

Creating Service Profile Templates and Cloning Service Profiles

Overview

Service profile templates build on the idea of manually created service profiles. With the potential for a large population of blade servers in a given Cisco Unified Computing System (UCS), manual creation of service profiles would be both slow and error-prone. The use of templates allows the Cisco UCS administrator the ability to create server definitions in a consistent and rapid process.

Unlike manually created service profiles, service profile templates must use pools for identity and server assignment. Derived hardware values are never used.

Objectives

Upon completing this lesson, you will be able to configure service profile templates and automate the creation of service profiles that are based on the template. This ability includes being able to meet these objectives:

- Create a service profile template and describe the need for pooled resources and identities
- Describe the reasons to create differentiated service profile templates to allow variations of policy
- Automate the creation of a server farm using service profile templates
- Describe the hidden pitfalls when using updating templates
- Unbind a service profile from its template
- Clone a service profile and demonstrate understanding of cloning requirements

Service Profile Templates

This topic discusses the configuration of service profile templates and the need for pooled resources and identities.

Creating a Service Profile Template

Creating a Service Profile Template

- Process is similar to service profile creation
- Requires server pool (cannot select individual blade)
- Requires pooled identities (UUID, MAC, WWNN, WWPN)

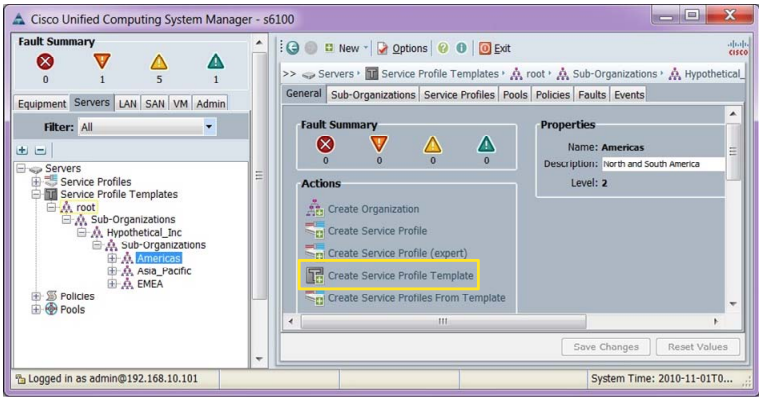

The process of creating a service profile template is nearly identical to creating a service profile manually. The principal difference is that service profile templates cannot be directly applied to a compute node and no hardware elements can use a derived value.

Name the New Template

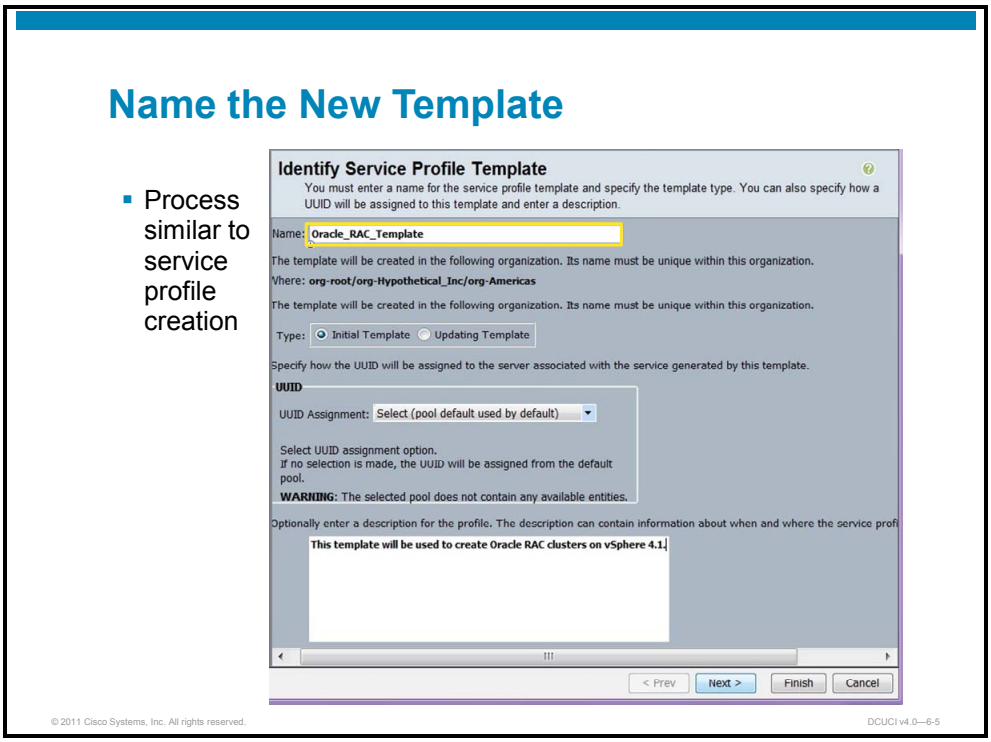

Service profile templates require a name, just as manually created service profiles. Templates can also contain an optional description.

Template Types

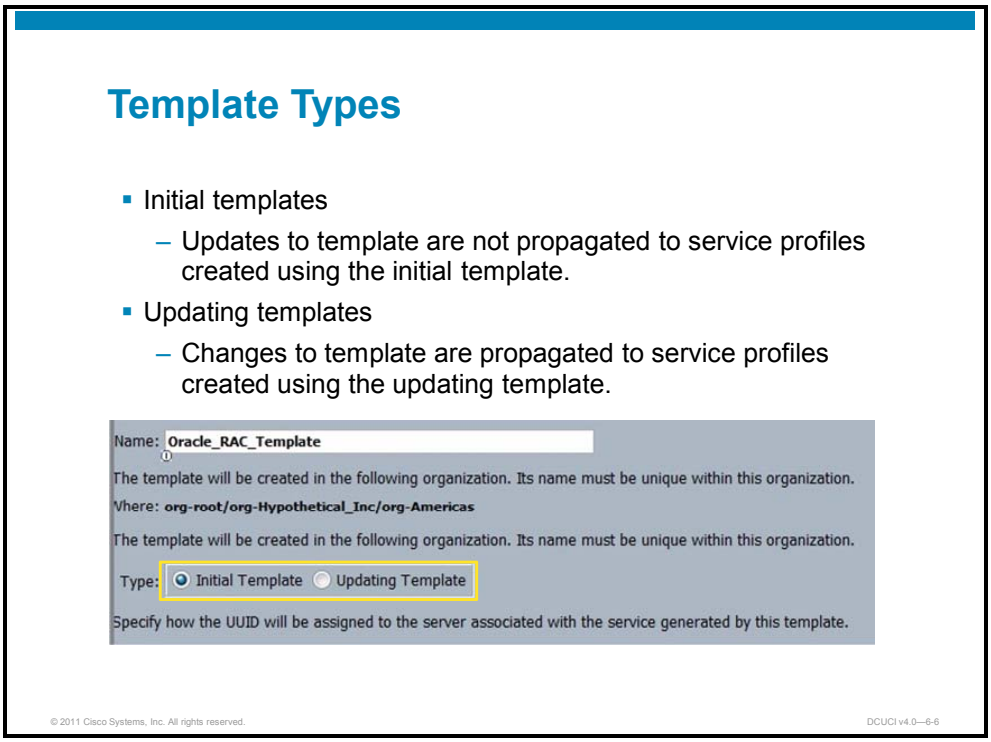

There are two types of templates. For both types, profiles that are created from a template cannot be modified. The ability to bind or unbind a service profile from its template will be discussed later in this lesson.

- **Initial templates:** This type of template maintains no connection to service profiles created from this definition. Changes to the template do not propagate to service profiles created from the template.
- **Updating templates:** This type of template maintains a semi-permanent link to service profiles spawned from this definition. Any changes to an updating template will be immediately propagated to all service profiles created from the template.

Apply UUID Pool

Apply UUID Pool

- All service profiles created from this template will use pooled UUIDs.

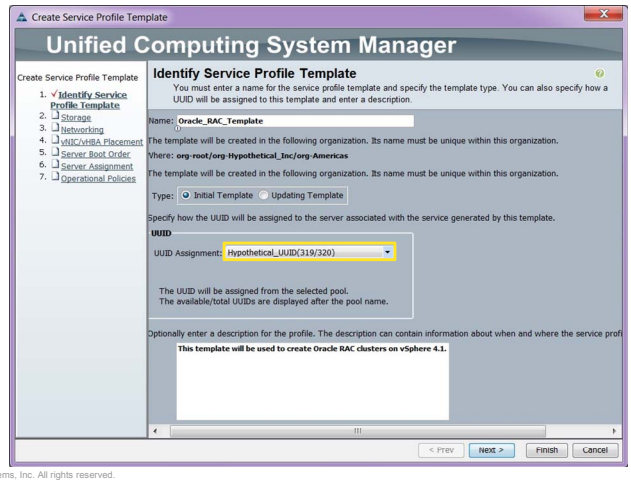

To facilitate stateless computing, the universally unique identifier (UUID) must be assigned from a pool. The UUID is unique in that it is the only identity resource that has the option of using the hardware default in the BIOS. Click **Next** to continue the wizard.

Apply WWNN Pool

Apply WWNN Pool

- All service profiles created from this template will use pooled WWNNs.

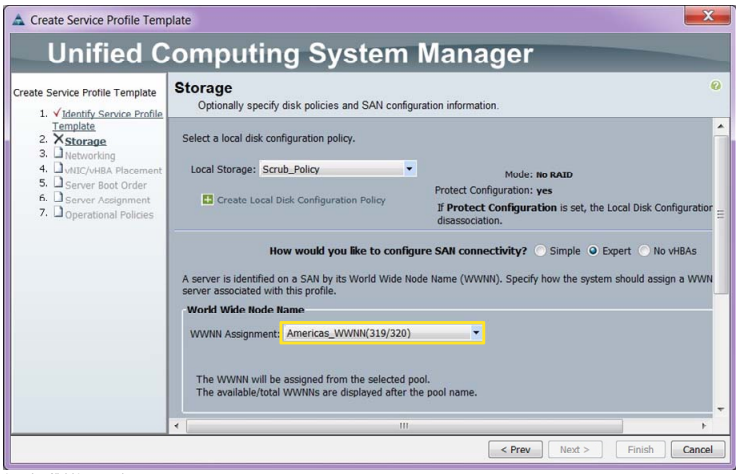

To enable Fibre Channel over Ethernet (FCoE) support to service profiles generated from this template, enter the name of the world wide node name (WWNN) pool.

Create vHBA for Fabric A

Create vHBA for Fabric A

- All service profiles created from this template will use pooled WWPNs.

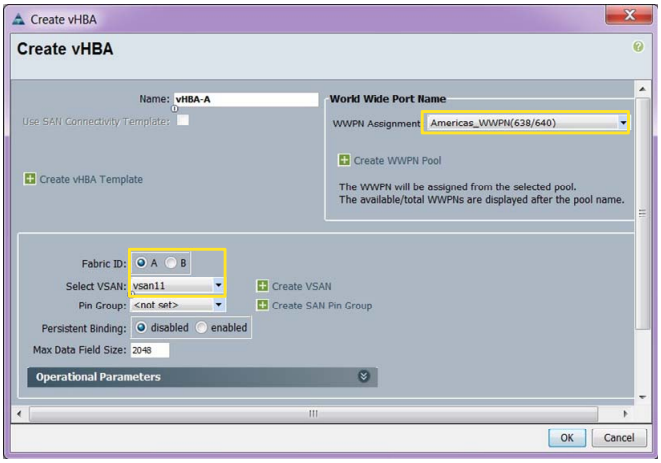

Click the **Expert** radio button to provide complete control over creating vHBAs. Click the **Add (+)** button to create the virtual host bus adapter (vHBA) for fabric A. As in the service profile creation wizard, enter a name, fabric affiliation, VSAN, and world wide port name (WWPN) pool.

Create vHBA for Fabric B

Create vHBA for Fabric B

- All service profiles created from this template will use pooled WWPNs.

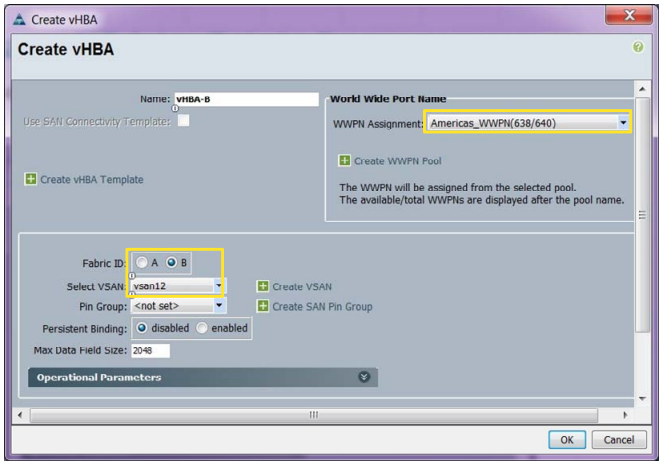

Click the **Expert** radio button to provide complete control over creating vHBAs. Click the **Add (+)** button to create the vHBA for fabric B. As in the service profile creation wizard, enter a name, fabric affiliation, VSAN, and WWPN pool.

vHBA Templates

vHBA Templates

- All the elements of a vHBA can be stored in a template.
- Templates are created under the SAN tab in Policies.

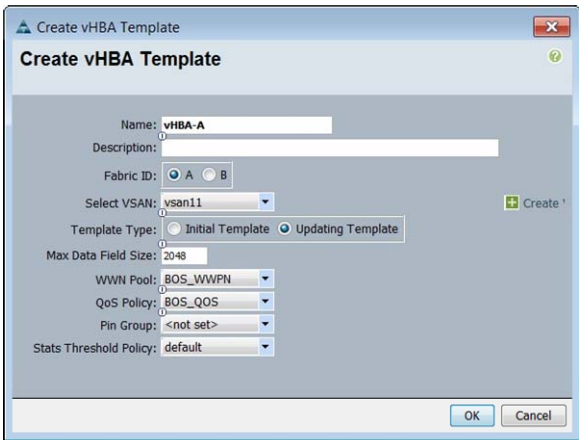

vHBA templates allow Cisco UCS administrators to speed the repetitive task of entering vHBA parameters in many templates or service profiles. Enter the information once in the template and never set it again.

Create vNIC for Fabric A

Create vNIC for Fabric A

- All service profiles created from this template will use pooled MACs.

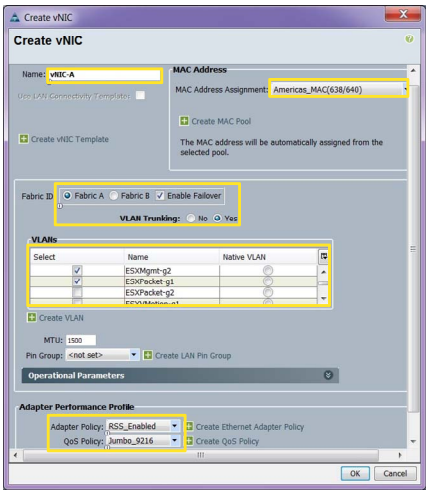

The definition of a virtual network interface card (vNIC) in the service profile template wizard is identical to the way that the definition is created in the manual service profile wizard. Enter a name for the vNIC, MAC address pool, access VLAN or VLANs associated with an 802.1Q trunk, and adapter performance profiles. In this example, the receive-side scaling (RSS) and jumbo frame policies will be bound to every service profile generated from this template.

Create vNIC for Fabric B

Create vNIC for Fabric B

- All service profiles created from this template will use pooled MACs.

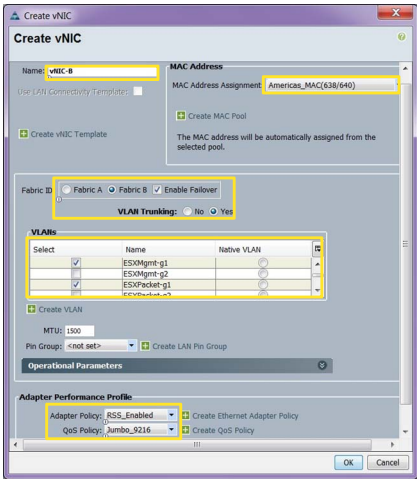

The definition of the vNIC for fabric B is identical to the creation of the vNIC for fabric A, except for the name.

vNIC Templates

vNIC Templates

- All the elements of a vNIC can be stored in a template.
- Templates are created under the LAN tab in Policies.

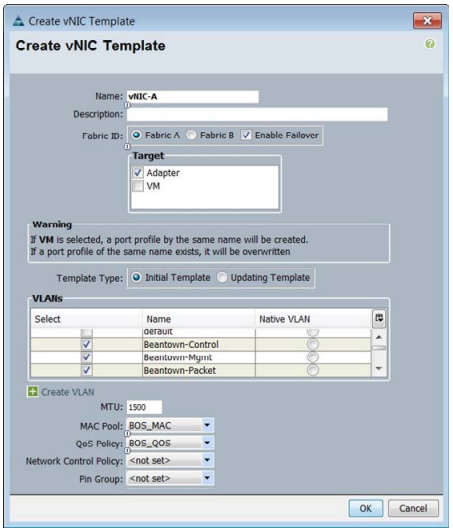

vNIC templates allow Cisco UCS administrators to speed the repetitive task of entering vNIC parameters in many templates or service profiles. It is especially useful when many VLANs must be selected for a trunk interface.

Boot Order and Boot Target

Boot Order and Boot Target

- Create the boot order as if it were a service profile and not a template.

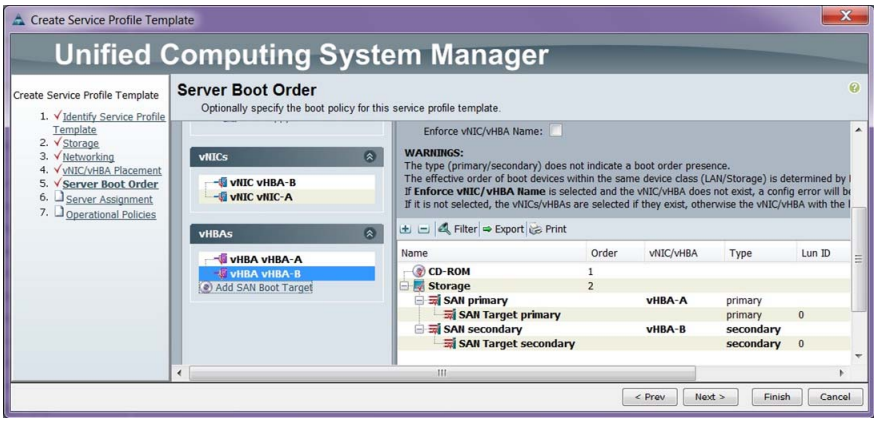

The Cisco UCS administrator has two options regarding the boot order. The example illustrates a SAN boot policy, where every initiator is mapped to the same target logical unit number (LUN). This is only possible if the storage system maps the source WWPN with a unique LUN. This method is very useful if the SAN administrators can provide premapped WWPNs to LUNs.

The second option is to define the vHBAs as bootable, but leave the boot target definition for a later time.

Template Server Assignment

Template Server Assignment

- Templates can only assign servers from a server pool.

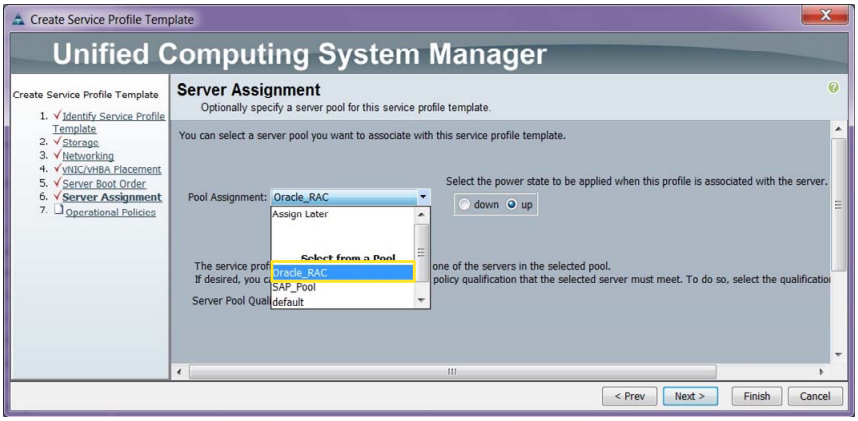

It is clear from the drop-down list for server assignment that manual assignment is not an option. Servers must be assigned from a pool.

IPMI and SoL Policy Assignment

IPMI and SoL Policy Assignment

- Add IPMI and SoL policies created earlier.

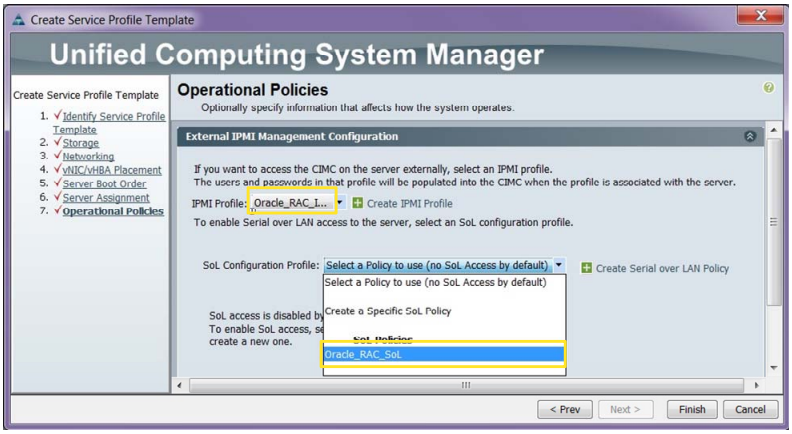

The Intelligent Platform Management Interface (IPMI) and Serial over LAN (SoL) policies that were created earlier can be applied to the template. All service profiles that are generated from this template will inherit both policies.

BIOS and Scrub Policy Assignment

BIOS and Scrub Policy Assignment

- Add BIOS and scrub policies.

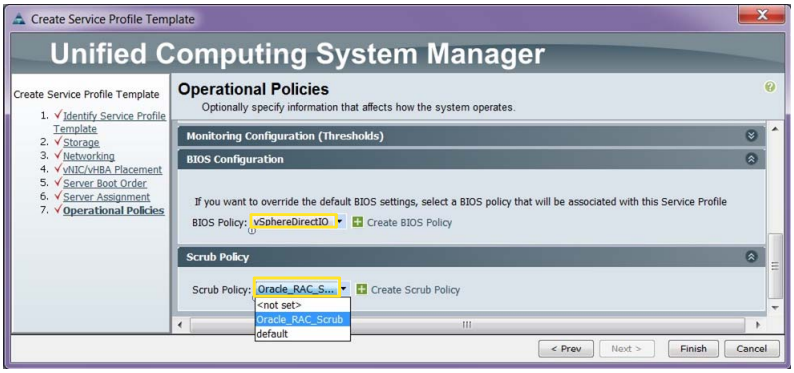

The BIOS and Scrub policies are applied to the template, and will be assigned automatically to every service profile generated by the template.

Modify Template

Modify Template

- Pool assignments and policies can be changed after template creation.

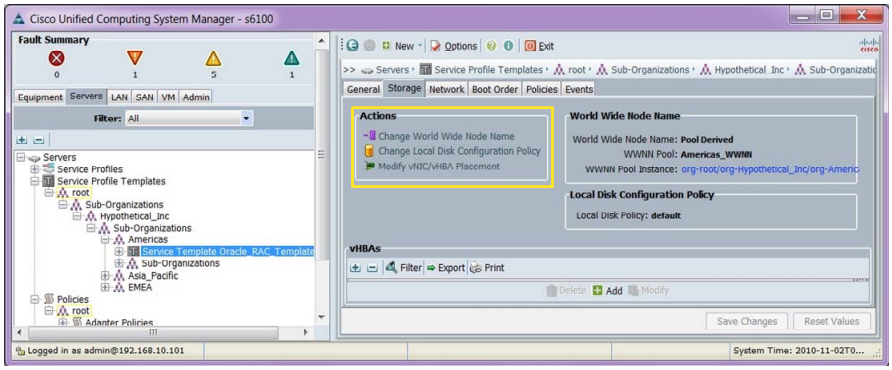

After a service profile template has been generated, it can be modified in a very similar manner to a manually created service profile. Certain changes made to an updating template will be propagated to every service profile generated by the template.

Creating Differentiated Service Profile Templates

This topic discusses the need for differentiated service profile policy in heterogeneous computing environments to allow variations of policy.

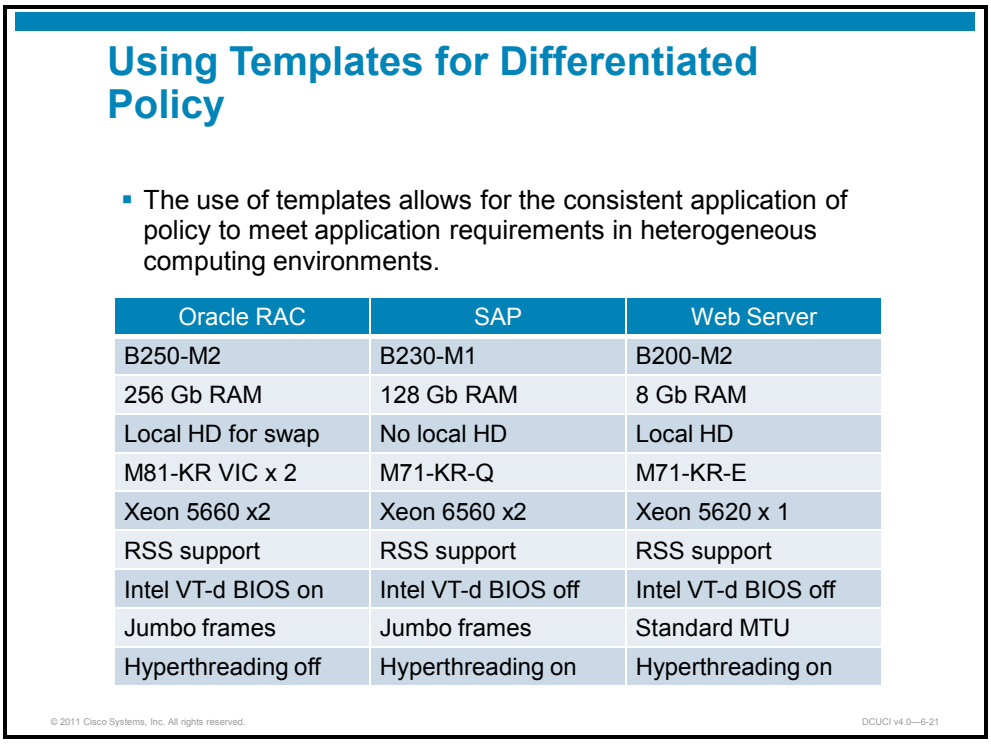

The figure shows the power and flexibility of using templates for differentiated policy. Because groups of applications share similar or identical requirements, service profile templates can be created to seamlessly provide the identity of server resources that are needed to serve the application. One of the important operational benefits to this approach is consistency of policy across the entire class of applications.

Automating Creation of a Server Farm Using Service Profile Templates

This topic discusses automated compute node provisioning from templates.

Creating Service Profiles from Template

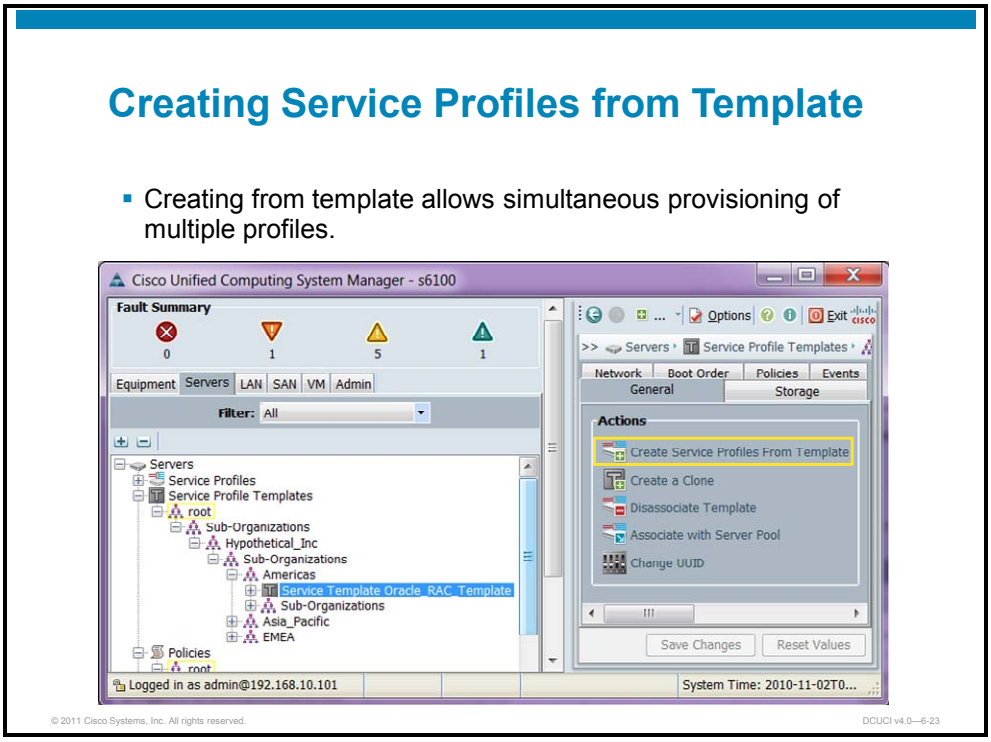

Another benefit of using service profile templates is that the Cisco UCS administrator can automate the provisioning of one to hundreds of compute nodes into simple operations. After a service profile template is built and points to identity and resource pools with sufficient resources, automation can begin.

Select the service profile template and the organization where the new service profiles are to be created in the navigation pane. In the content pane, click the link **Create Service Profiles From Template**.

Select Prefix and the Number of Profiles

Select Prefix and the Number of Profiles

- Provide naming prefix and number of service profiles based on the service profile template.
- In the example, the naming prefix would create service profiles Oracle_RAC1 through Oracle_RAC20.

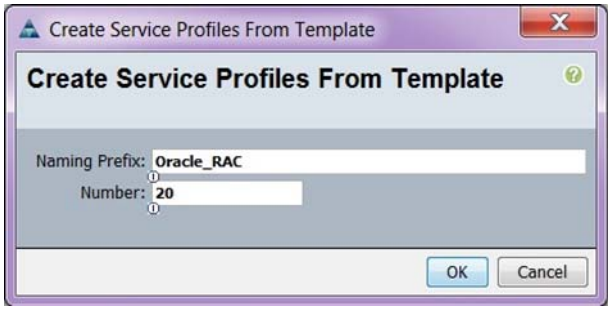

The dialog box prompts you for a naming prefix and the number of service profiles to be generated with that prefix. Immediately after you click **OK**, service profiles will appear under the organization. If the server assignment in the template points to a server pool, a new service profile will immediately begin to associate with the next available server in the pool.

Describe the Hidden Pitfalls When Using Updating Templates

This topic discusses issues related to the use of updating templates that can directly affect production operations.

Updating Template Issues to Consider

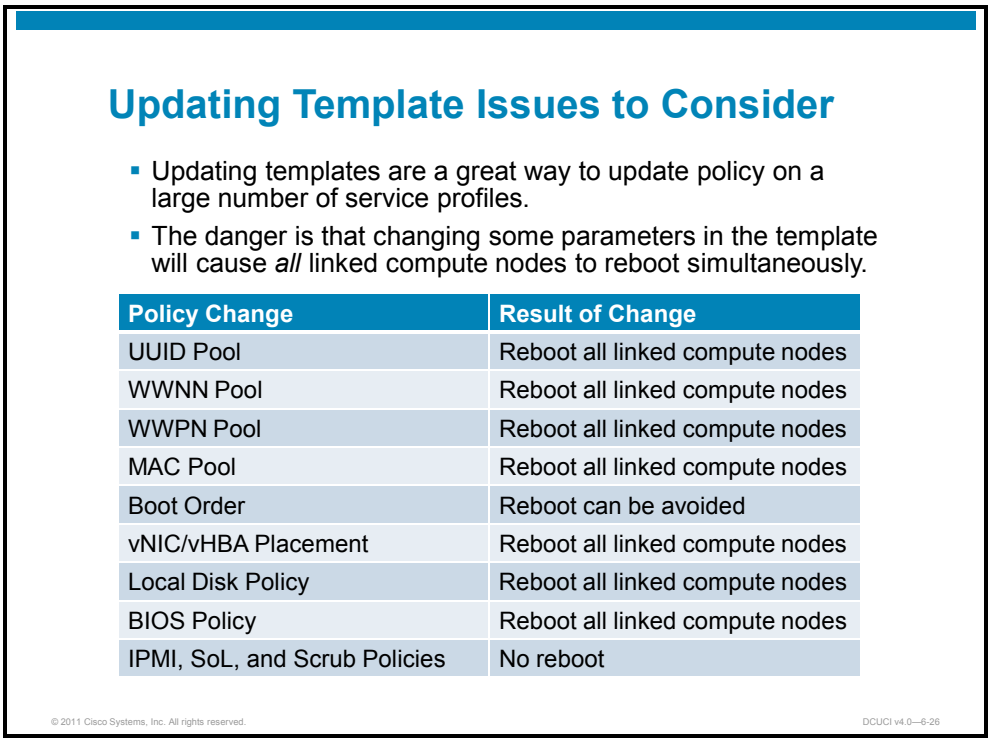

Changes to updating templates are immediately propagated to any service profiles that were generated from that template. If none of the generated service profiles are associated to a compute node, there is no risk to an update. However, if certain changes are made to the updating template, it will cause all linked compute nodes to reboot. A summary of template modifications and their associated reactions are shown in the table.

Updating Template Warning

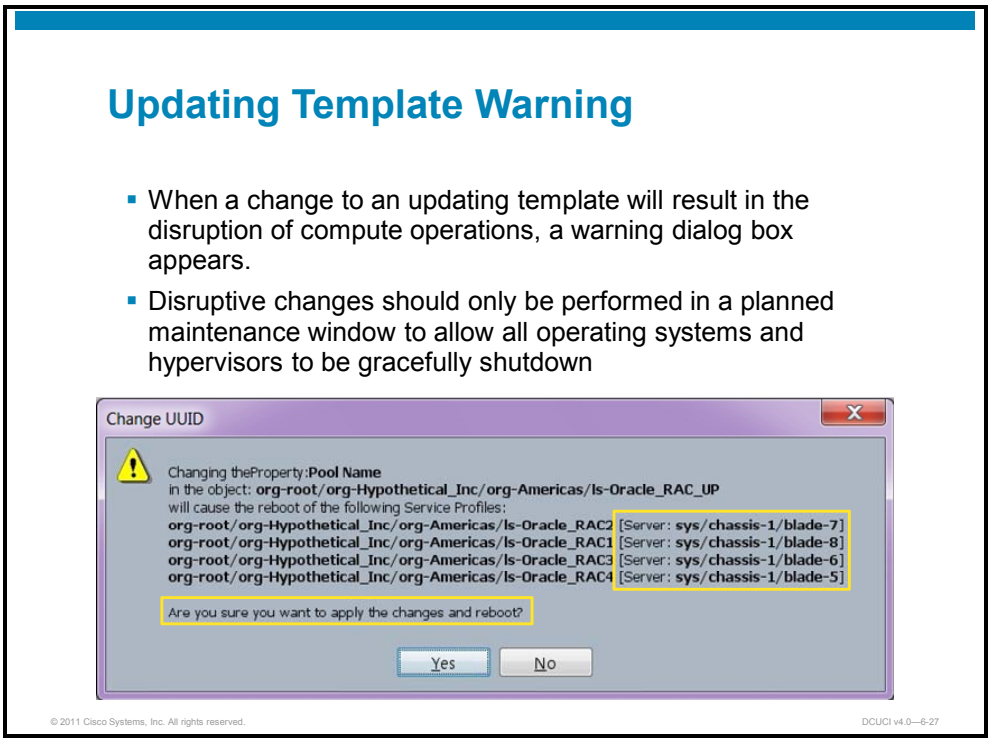

Beginning with Cisco UCS Manager version 1.2, the system warns you that if the modification to the updating template is executed, all impacted compute nodes will reboot immediately.

The best practice in this case is to perform the update in a scheduled and approved maintenance window that provides for the graceful shutdown of all compute nodes that the change will affect.

Unbind a Service Profile from Its Template

This topic discusses how to unbind a service profile from its template.

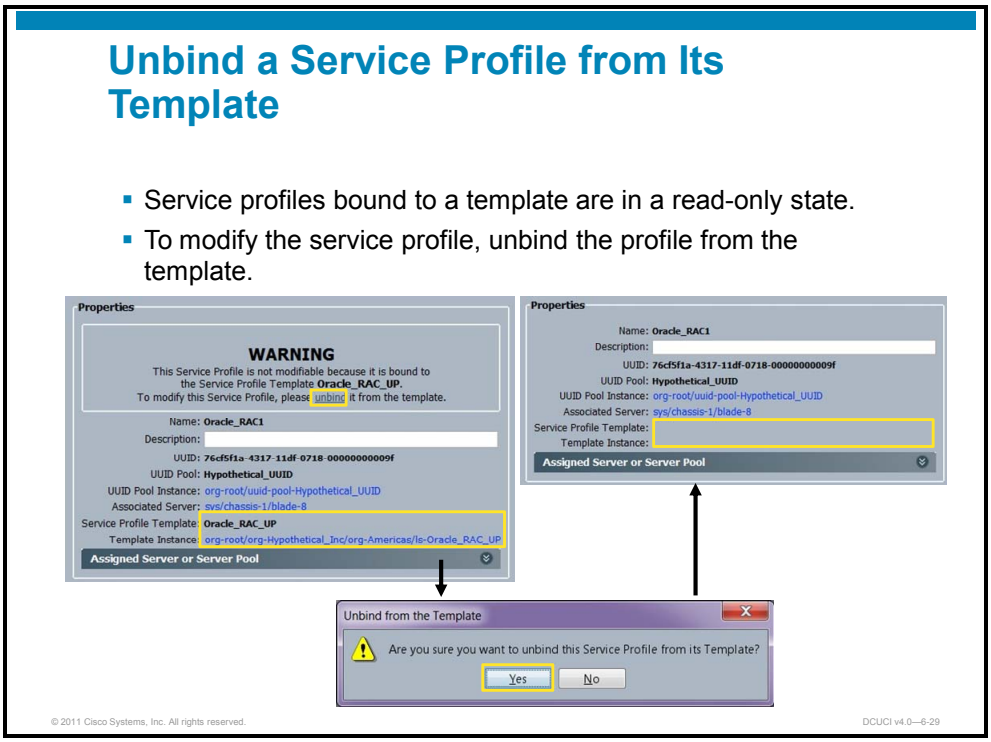

One of the consequences of generating service profiles from a template is that the resulting service profiles are created in a read-only state. The example displays a conspicuous warning message alerting the Cisco UCS administrator that no changes can be made to the service profile unless it is unbound from the template. By clicking the unbind link, a small dialog box appears asking the administrator to confirm the operation. When the operation is confirmed, the service profile no longer displays the warning or its link to its parent template.

Bind a Service Profile to a Template

Bind a Service Profile to a Template

- Stand-alone service profiles can be bound to a template at a later time.

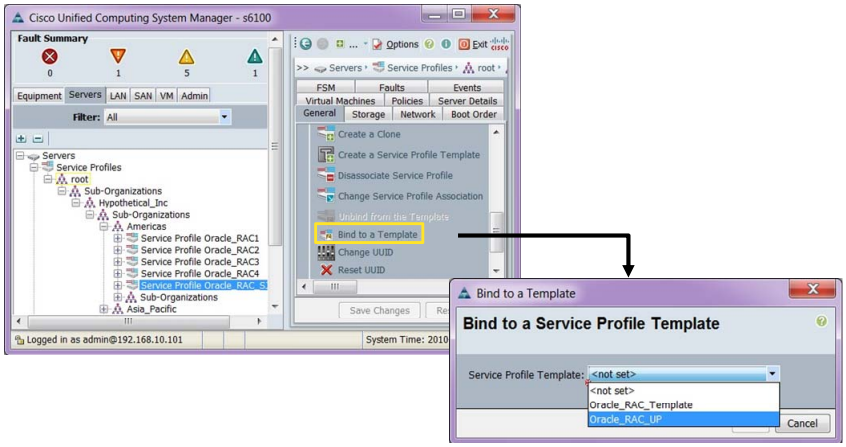

It is also possible to bind a manually created service profile to a template. If the previously created service profile is bound to an initial or updating template, it will retain the identity information for UUID, WWNN, WWPN, and MAC address unless the template uses different pools. If the template uses a different named pool, identity information will be replaced with data pulled from the pools of the template.

Cloning a Service Profile

This topic discusses cloning operations and service profiles in service profile templates.

Service Profile Cloning

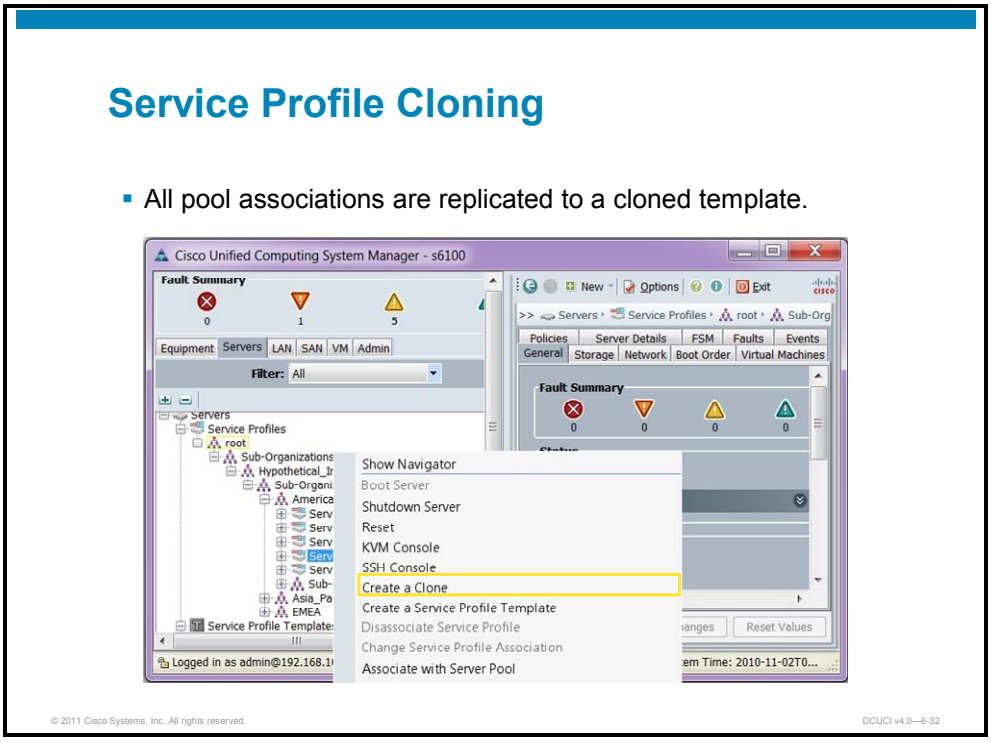

Service profiles and service profile templates can both be cloned. Simply right-click on the name of the service profile or service profile template and select **Create a Clone**. The result of this operation is that all pooled identities in the clone will be refreshed with unique values. The boot order is cloned verbatim.

Clone Destination

Clone Destination

- Unique values for MAC and WWN will be immediately assigned to the profile from the appropriate pool.
- Select the destination organization where the clone should be created.

When you commit to creating a clone, a pop-up dialog box prompts you for the name of the clone and destination organization where the clone will be created.

Note	Remember that once an object is created in a particular organization, it cannot be moved or renamed. If an object is created in the wrong organization, it must be deleted and then re-created in the correct organization.

Summary

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

Summary

- Service profile templates require all identity and resources to be derived from a pool.
- Differentiated service profile templates allow for the consistent application of policy in heterogeneous computing environments.
- Service profile templates can be leveraged to create an arbitrary number of service profiles associated to compute nodes.
- Updating templates are useful for maintaining consistent policy across a large population of compute nodes, but modifying certain parameters will cause all linked compute nodes to reboot.
- Service profiles created from templates cannot be changed unless they are unbound from the template.
- Clones made from existing service profiles maintain the boot order for beta, but select new identities from identity pools.

Lesson 4

Managing Service Profiles

Overview

While a Cisco Unified Computing System (UCS) is operational, periodic maintenance to service profiles may be necessary. Corporate policy and regulatory compliance requirements can dictate changes to policy. Updates to Cisco UCS Manager software regularly introduce support for new features, capabilities, and new blade servers. The pace of innovation within Cisco UCS has been swift and dramatic. Cisco UCS has grown from a product with a single blade in the fabric interconnect to a rich selection of components ready to meet the needs of the most demanding applications in enterprise data centers and cloud computing. Service profile maintenance will be an ongoing periodic process.

Objectives

Upon completing this lesson, you will be able to disassociate and associate service profiles from compute nodes. You will recognize which parameters of a service profile trigger a disruptive change. This ability includes being able to meet these objectives:

- Use Cisco UCS Manager to associate and disassociate a service profile to a server blade
- Describe what changes to a service profile trigger a Cisco UCS utility operating system update, and outage to a server
- Describe the importance of planning the organization where a service profile is created
- Use Cisco UCS Manager to move a service profile to a new server blade in the event of hardware failure

Associating and Disassociating a Service Profile to a Server Blade

This topic discusses associating and disassociating service profiles with compute nodes.

Associate a Service Profile with a Compute Node

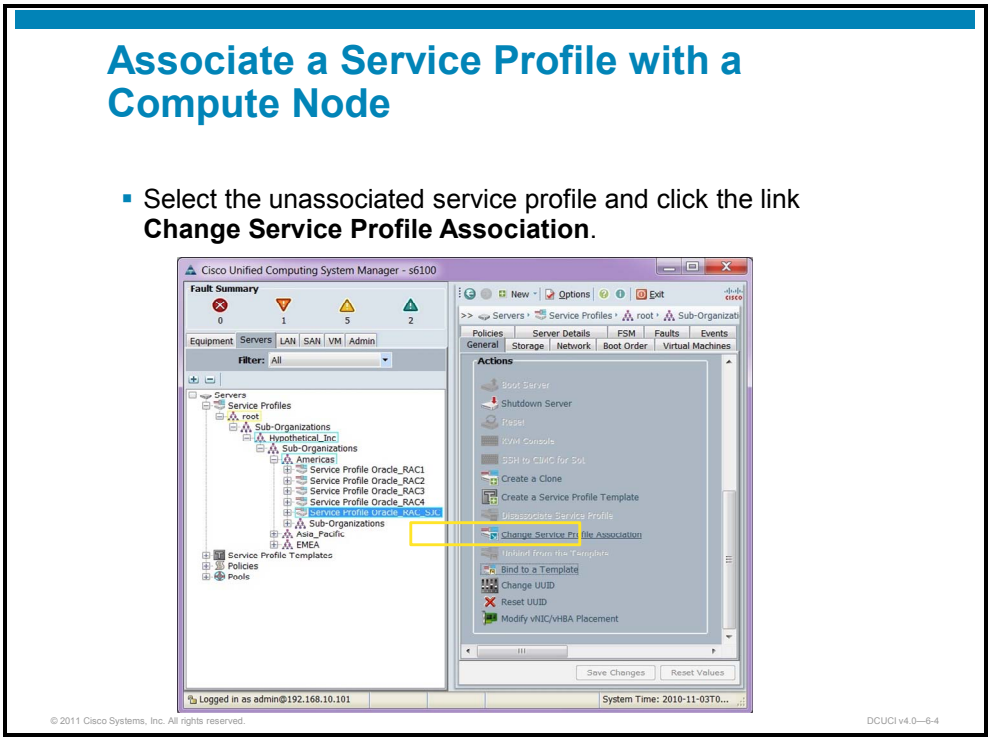

If you decided not to immediately assign a service profile to a compute node, you can select the desired service profile in the navigation pane. In the General tab of the content pane, click the link labeled **Change Service Profile Association**.

Associate a Service Profile with a Server Pool

Associate a Service Profile with a Server Pool

- Select the appropriate server pool from the Server Assignment drop-down list.

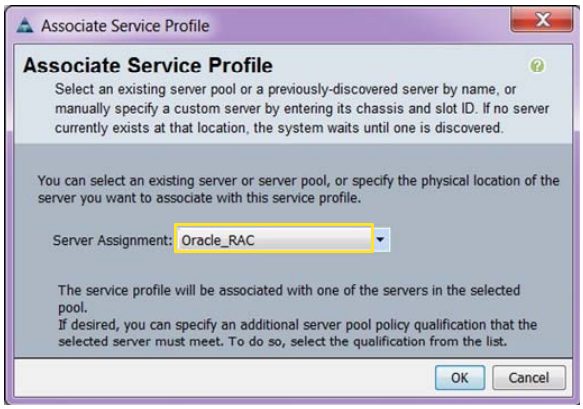

A pop-up dialog box prompts you to select either from an existing unassociated server or server pool. Unlike service profile templates, service profiles that are not bound to a template are not required to select a server from a pool only. Click **OK** to begin the association process.

Observe FSM Status During Service Profile Association

Observe FSM Status During Service Profile Association

- You can observe the association process in the FSM tab in the content pane of the service profile.

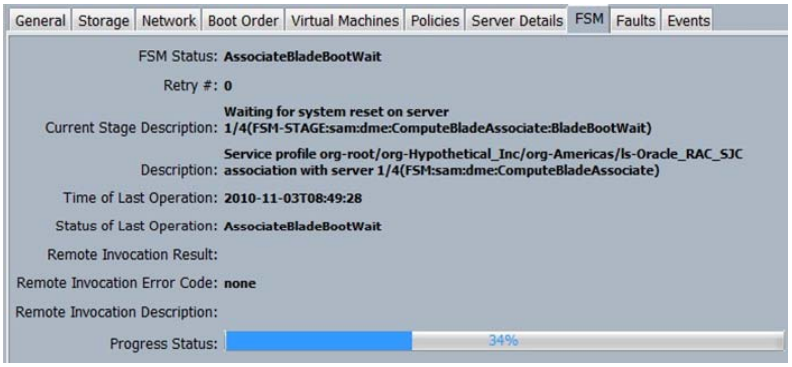

You can follow the complete process of association by clicking the **FSM** tab in the content pane. Recall from "Monitoring System Events" that service profile association and disassociation are complex processes that are assigned to a finite state machine (FSM).

If the service profile is unable to associate with the compute node that is selected, the FSM will provide information on which step of the process a failure occurred. This is very useful for troubleshooting service profile association issues.

Note	Be aware that the FSM status indicator may appear to stop and lock up. Some stages of the association process can take one minute or longer to complete. This is normal.

What Happens During Service Profile Association?

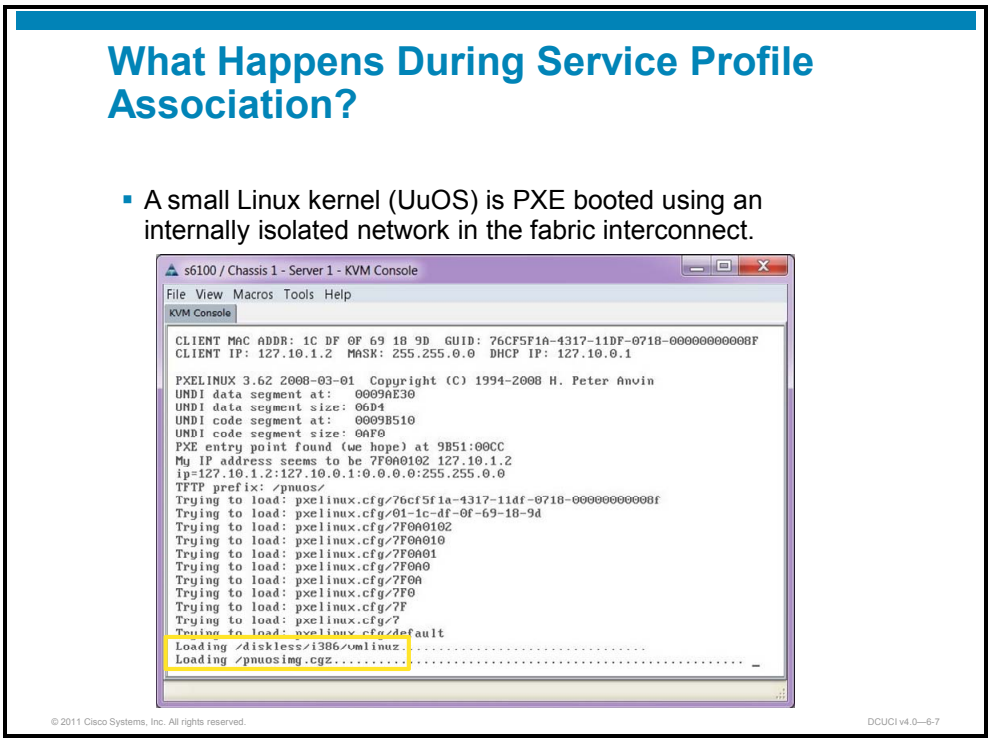

The processes that occur during service profile association and disassociation are very interesting. The first step in associating a service profile to a compute node begins by powering up the server. Next, the server Preboot Execution Environment (PXE) boots a small Linux distribution over a private network connection to the fabric interconnect.

The screenshot in the figure highlights the term “pnuosimg.” Before Cisco UCS was released to the public, this Linux operating system was referred to as Processor Node Utility Operating System, or PNuOS. The official name for this is Cisco UCS Utility Operating System. The old terminology still appears in some contexts.

Note	The black text on white background was reversed from the standard keyboard, video, mouse (KVM) output of white text on black background in a graphics program, for readability. The KVM does not have a choice of text or background colors.

Cisco UCS Utility Operating System

Cisco UCS Utility Operating System

- The Cisco UCS Utility Operating System PXE booted from the fabric interconnect applies all configuration and policy elements of the service profile to a compute node, then exits and reboots the compute node.

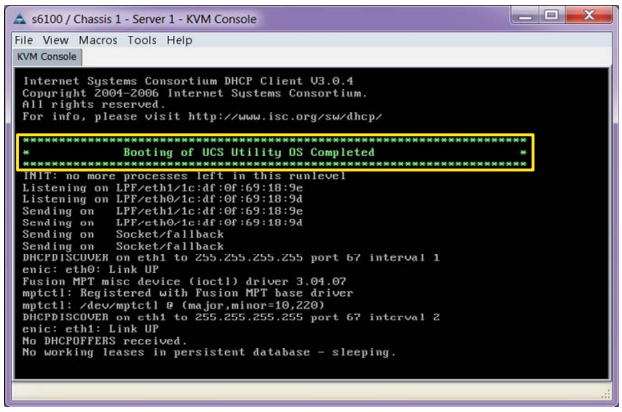

During PXE boot, the server obtains a DHCP address over the private network. This network is completely isolated from the in-band and out-of-band connections that are processed by the fabric interconnect and servers.

The purpose of booting this Linux operating system is to program the compute node. You will see identity information such as universally unique identifier (UUID), MAC address, world wide network node (WWNN), world wide port name (WWPN), BIOS configuration, adapter policies, and so on.

Disassociate a Service Profile from a Compute Node

Disassociate a Service Profile from a Compute Node

- Select the service profile, and click the link **Disassociate Service Profile**.
- Click **Yes** in the warning dialog box to confirm the action.
- Monitor the process in the FSM tab.

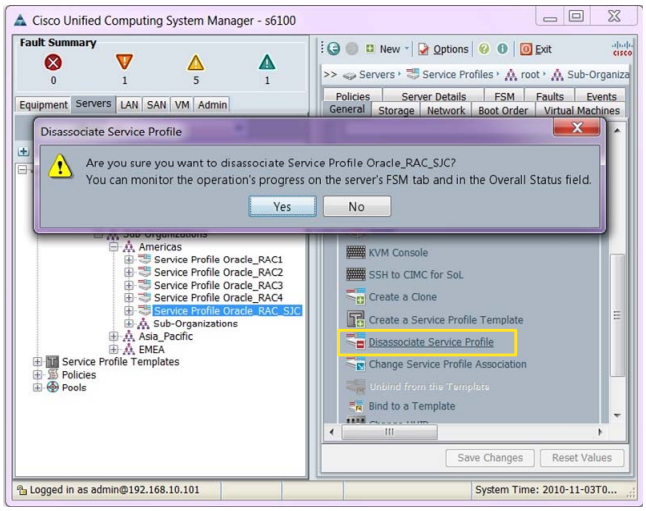

To disassociate a service profile from its compute node, select the service profile in the navigation pane. In the content pane click the link **Disassociate Service Profile**. A pop-up warning dialog asks you to verify the operation. Note also the small comment about observing the process in the FSM tab.

FSM Status During Service Profile Disassociation

FSM Status During Service Profile Disassociation

- You can observe the disassociation process in the FSM tab in the content pane of the service profile.

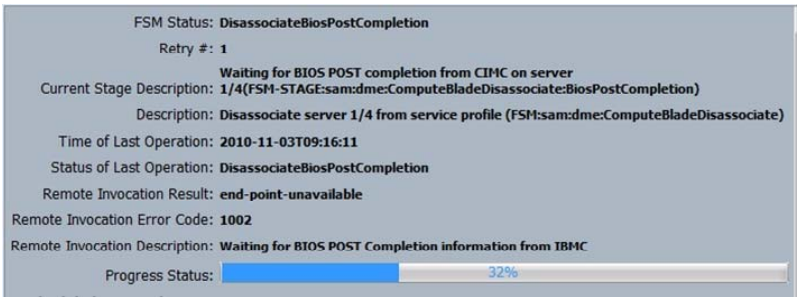

Both association and disassociation processes are monitored by FSM. Click the **FSM** tab in the content pane to observe the process of disassociation.

Changes to a Service Profile that Trigger a Cisco UCS Utility Operating System Update

This topic discusses the type of service profile modifications that automatically trigger a Cisco UCS Utility Operating System update, and outage to a server.

Changes Trigger Cisco UCS Utility Operating System

- When a service profile is associated with a compute node, certain changes made to the service profile trigger a Cisco UCS Utility Operating System update (and instant server reboot).

Configuration Change	Triggers UCSuOS to Run?
Change pooled identity	Yes
Change non-pooled identity	Yes
Change Fibre Channel boot target	Yes
Change boot order	Maybe
Change BIOS config policy	Yes
Change vNIC failover	Yes
Change adapter policy	No
Change VSAN on vHBA	No
Change VLAN on vNIC	No
Change from access VLAN to trunk	No

UCSuOS = UCS Utility Operating System

It is important to understand what types of service profile modifications can be made outside of a change control maintenance window. The table summarizes changes that will trigger a Cisco UCS Utility Operating System to run. As of version 1.2 of Cisco UCS Manager, the system alerts you to the changes that will result in the compute node being immediately rebooted.

Planning the Organization Where a Service Profile Is Created

This topic discusses the importance of creating service profiles and other policy objects in the correct organization.

Creating Service Profiles in the Correct Organization

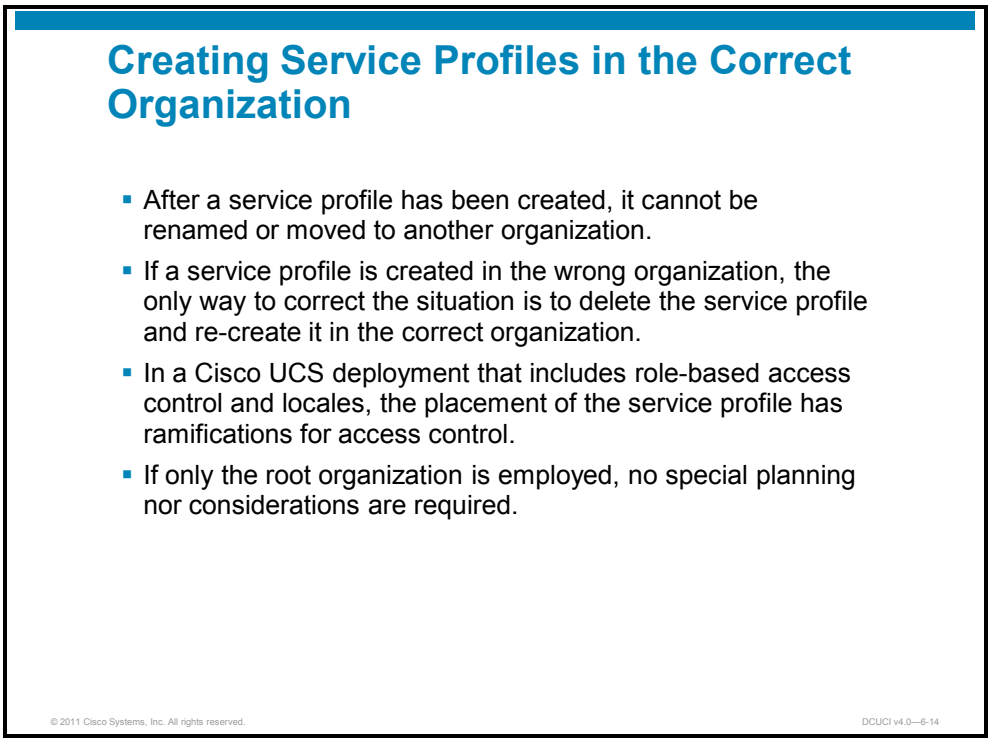

If a given Cisco UCS deployment creates all identity, resource, pool, and policy objects in the root organization, no special planning or considerations are required.

If the Cisco UCS administrators have created a significant number of objects under the root organization and decide to apply administrative hierarchy at a later date, significant work will be required. Every object that needs to move into a nonroot organization will need to be deleted and re-created in the new organization.

When an object has been created in a given organization, it cannot be renamed or moved.

Creating Service Profiles in the Wrong Organization

Creating Service Profiles in the Wrong Organization

- Administrators restricted to the Boston locale cannot manage service profile BOS_Oracle_RAC1 created in the Americas parent organization.

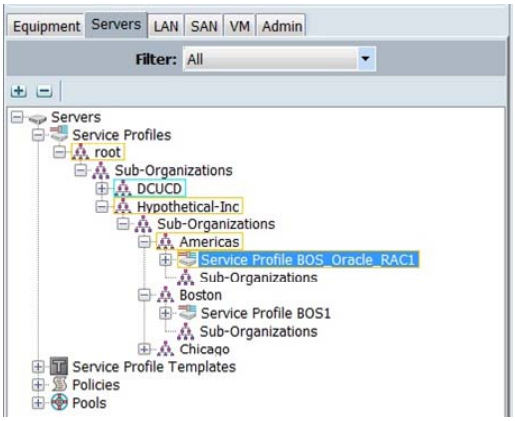

In the example in the figure, a service profile was created in the Americas organization. If this company employs role-based access control (RBAC) and limits administrative scope based on locale, administrators in the Boston organization may have no control over the service profile BOS_Oracle_RAC1. The service profile must be deleted and re-created in the appropriate organization.

Moving a Service Profile to a New Server Blade in the Event of Hardware Failure

This topic discusses using Cisco UCS Manager to move a service profile to a new server blade in case of hardware failure.

A Compute Node Hardware Has Failed

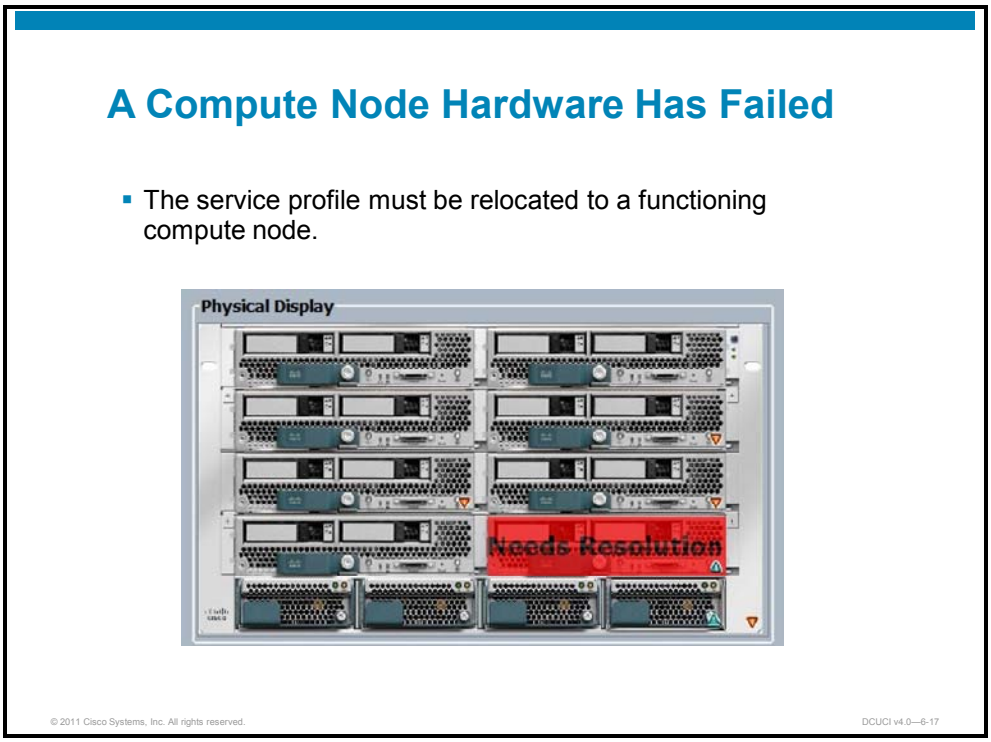

The blade server in slot 8 has experienced a severe failure. An administrator used Cisco UCS Manager to decommission the blade in slot 8. The service profile that is associated with this compute node is automatically de-linked.

Automatic Service Profile Reassociation

Automatic Service Profile Reassociation

- Because the compute node for service profile Oracle_RAC_1 was selected from a server pool, the service profile automatically reassociated to an available server in that pool.

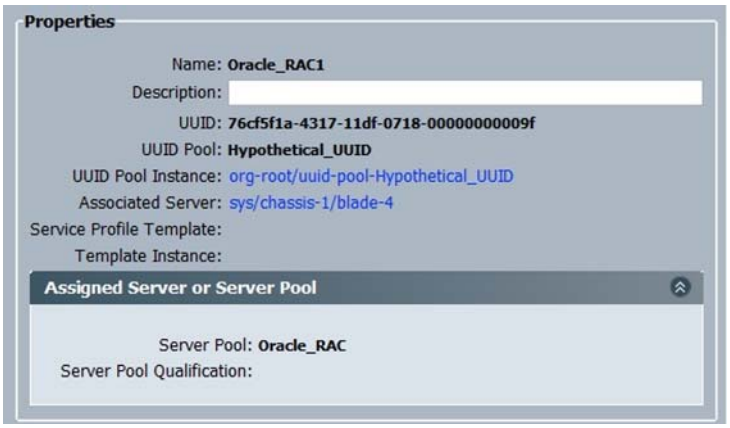

Because the service profile received its server assignment from a server pool, it automatically selected a new server from the same named pool without any further action from the administrator. After the service profile reassociated to the new compute node, the operating system or hypervisor automatically rebooted on the new compute node.

Summary

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

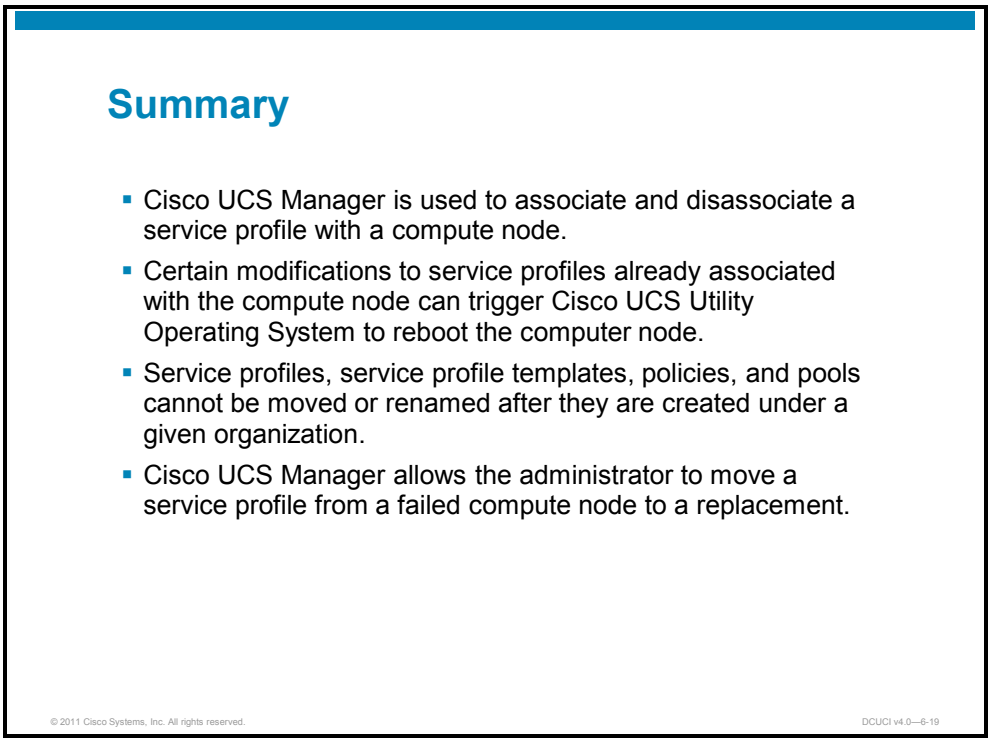

Module Summary

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

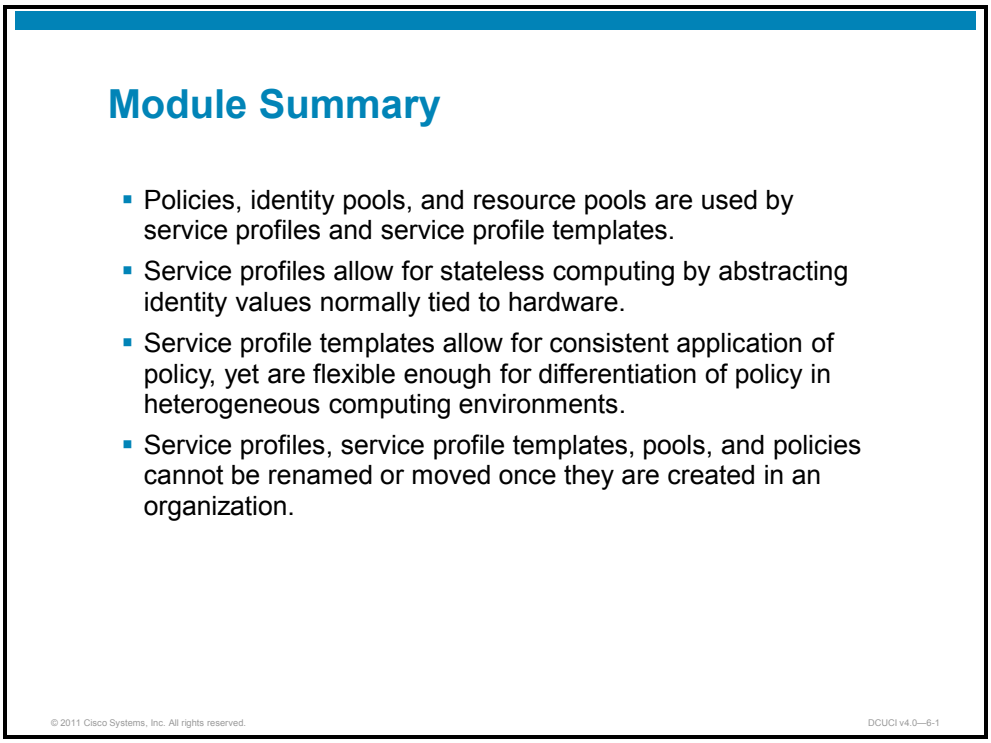

This module examines the building blocks of stateless computing.

In traditional server environments, where a single operating system runs on the underlying hardware, it is very challenging to replicate a consistent configuration across a large population of servers. Adding to this challenge is the fact that unique network and storage addresses are tied to the hardware. In Cisco UCS, these challenges are addressed by a flexible array of reusable policy objects, identity pools, and resource pools. These reusable objects are consumed by a construct that is called a service profile.

Service profiles act as a hardware abstraction layer. Instead of relying on identity locks to the hardware, all elements of a server “personality,” or state, are encapsulated in this policy construct. The benefit of this approach is that if the underlying hardware fails, the service profile is mobile and can be reassociated with another compute node. The operating system or hypervisor boots on the new compute node and is unaware that the hardware has changed. This avoids time-consuming downtime that is associated with re-acclimating an operating system and applications with all of the new hardware identifiers on the replacement server.

Cisco UCS extends the functionality of service profiles with profile templates. Templates offer a way to automate the creation of many service profiles and ensure consistent policy. In heterogeneous computing environments, differentiated policies and service profile templates afford the Cisco UCS administrator a rich set of tools to manage large deployments.

Each compute node requires a unique service profile. There is a one-to-one mapping of service profiles to compute nodes. When the Cisco UCS administrator begins the association process, the Preboot Execution Environment (PXE) boots a special-purpose Linux distribution on a private network connection to the fabric interconnect. The Cisco UCS Utility Operating System is responsible for programming all elements of identity and policy to the compute node. Because some identities and policies are integral to the operation of the compute node, caution should be observed when modifying certain elements of the service profile. Many elements of the service profile will cause the node to reboot to run the Cisco UCS Utility Operating System to apply the request to change.

References

For additional information, refer to these resources:

- Gai, S., Salli, T., et al (2009). Project California: a Data Center Virtualization Server. Raleigh, NC: Lulu.com.
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Configuring Server-Related Pools.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter24.html
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Creating Server Pool Policy Qualifications.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter25.html#task_43C18874A0D245D6987C3530BD4F06C7
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Configuring Service Profiles.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter26.html
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Ethernet and Fibre Channel Adapter Policies.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter23.html#concept_C113E10277D74C5FB66C92A87293B696
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Configuring QoS System Classes.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter18.html#task_2F15A8D4D2B34ED79B7917877D749B47
- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Setting the vNIC/vHBA Placement.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter26.html#task_9E40E74DB2EE43EDA8A74C44AFC9323A

- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Creating a Service Profile Template.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter26.html#task_23DCA7911736413A9D03179A23523A0A

- Cisco Systems, Inc. *Cisco UCS Manager GUI Configuration Guide, Release 1.3(1)—Cloning a Service Profile.*

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/UCSM_GUI_Configuration_Guide_1_3_1_chapter26.html#task_8593C436D54F4C0BA2467C01FFBC0F81

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 statement is valid regarding service profile placement? (Source: "Creating Identity and Resource Pools")

A) Service profiles, policies, templates, and pools can be moved from organization to organization.
B) Service profiles, policies, templates, and pools can be renamed.
C) Service profiles, policies, templates, and pools are the foundation of stateful computing.
D) Service profiles, policies, templates, and pools cannot be moved from organization to organization.
E) Service profiles, policies, templates, and pools are used to enable syslog.

Q2) Which statement is true regarding the creation of identity pools? (Source: "Creating Identity and Resource Pools")

A) Each identity occupies no space in the database.
B) Pools cannot be extended after creation, but additional blocks of identities can be added later.
C) Pools can be extended after creation, but additional blocks of identities cannot be added later.
D) It is a best practice to always create the maximum number of addresses in a pool (9999).

Q3) Which three items are reasons for adopting pooled identity resources? (Choose three.) (Source: "Creating Identity and Resource Pools")

A) ensures global uniqueness of identities in an entire network
B) leads to consistency of policy
C) enables scalable provisioning
D) ensures stateful computing
E) eliminates the need to license hardware-derived identities
F) allows for profile mobility

Q4) How many bits of the UUID are represented in the UUID suffix blocks? (Source: "Creating Identity and Resource Pools")

A) 32
B) 48
C) 64
D) 128
E) 256

Q5) How many bits are represented in the MAC address prefix (OUI)? (Source: "Creating Identity and Resource Pools")

A) 32
B) 48
C) 64
D) 128
E) 256

Q6) Which statement accurately describes RSS? (Source: “Creating Service Profiles”)

A) a method of flow control applied to FCoE
B) a method of load-sharing across multiple interfaces
C) a method of load-sharing across multiple TCP connections
D) a method of distributing TCP sessions across multiple CPU cores
E) a method of load-sharing TCP sessions across multiple FCoE interfaces

Q7) Which item accurately describes where to configure a QoS system class? (Source: “Creating Service Profiles”)

A) LAN Uplinks Manager
B) Internal Fabric Manager
C) LAN tab under QoS
D) SAN tab under QoS

Q8) What are three functions of the Cisco Integrated Management Controller? (Choose three.) (Source: “Creating Service Profiles”)

A) KVM over IP
B) SoL
C) BMC
D) IPMI
E) QoS
F) FCoE

Q9) Which port does SoL use to communicate? (Source: “Creating Service Profiles”)

A) TCP port 22
B) TCP port 23
C) TCP port 623
D) UDP port 22
E) UDP port 23
F) UDP port 623

Q10) What are two functions of a scrub policy? (Choose two.) (Source: “Creating Service Profiles”)

A) remove adapter policies when a service profile is disassociated
B) remove BIOS settings when a service profile is disassociated
C) remove threshold settings when the service profile is disassociated
D) remove disk settings when a service profile is disassociated
E) erase disks when the service profile is disassociated

Q11) Which two of these actions can be configured in the expert service profile wizard but not in the simple service profile wizard? (Choose two.) (Source: “Creating Service Profile Templates and Cloning Service Profiles”)

A) Select access VLAN
B) Configure VSAN trunking
C) Configure VLAN trunking
D) Select MAC address from a pool
E) Use hardware-derived WWNN

Q12) Which part of a cloned service profile is copied verbatim? (Source: "Creating Service Profile Templates and Cloning Service Profiles")

A) UUID
B) MAC address
C) WWNN
D) WWPN
E) boot order

Q13) What must an administrator configure to modify the service profile created from a template? (Source: "Creating Service Profile Templates and Cloning Service Profiles")

A) unbind the service profile from the linked clone
B) unbind the service profile from the linked template
C) unbind the service profile from the linked compute node
D) unbind the service profile from the linked master

Q14) What is a concern when employing updating templates? (Source: "Creating Service Profile Templates and Cloning Service Profiles")

A) There is the potential to consume too many identity resources.
B) Modifying certain template parameters can cause a compute node to reboot.
C) Modifying certain template parameters can cause all linked nodes to reboot.
D) There is the potential to oversubscribe the server pool.

Q15) Which element allows the automated provisioning of service profiles? (Source: "Creating Service Profile Templates and Cloning Service Profiles")

A) MAC pool
B) WWNN pool
C) WWPN pool
D) VLAN
E) template

Q16) Which three configuration changes to a service profile will trigger a Cisco UCS Utility Operating System reboot? (Choose three.) (Source: "Managing Service Profiles")

A) modify MAC pool
B) modify WWNN pool
C) modify boot target
D) modify VLAN
E) change from access VLAN to trunk
F) change adapter policy

Q17) Which three configuration changes to a service profile are safe to perform during production (no automatic reboot)? (Choose three.) (Source: "Managing Service Profiles")

A) modify MAC pool
B) modify WWNN pool
C) modify boot target
D) modify VLAN
E) change from access VLAN to trunk
F) change adapter policy

Q18) Which item correctly describes where Cisco UCS Utility Operating System updates its IP address? (Source: "Managing Service Profiles")

A) dedicated private IP address in each compute node, supplied by the Cisco Integrated Management Controller
B) DHCP server running on the CMC in the IOM
C) DHCP server running on the fabric interconnect
D) external DHCP server
E) BOOTP server running on the CMC in the IOM

Q19) What are two properties of service profiles after they have been created in an organization? (Choose two.) (Source: "Managing Service Profiles")

A) Service profiles can be renamed.
B) Service profiles cannot be renamed.
C) Service profile boot target cannot be changed.
D) Service profiles can be moved to a different organization.
E) Service profiles cannot be moved to a new organization.

Q20) What is the correct method to move a service profile from a failed compute node to a replacement? (Source: "Managing Service Profiles")

A) Unbind the service profile from its template and it will automatically seek a new compute node.
B) Select **Change Service Profile Association** from the Equipment tab.
C) The service profile will automatically reassociate without any administrator intervention.
D) Select **Change Service Profile Association** from the Servers tab.

Module Self-Check Answer Key

Q1) D

Q2) B

Q3) B, C, F

Q4) C

Q5) C

Q6) D

Q7) A

Q8) A, B, D

Q9) F

Q10) B, E

Q11) C, D

Q12) E

Q13) B

Q14) C

Q15) E

Q16) A, B, C

Q17) D, E, F

Q18) B

Q19) B, E

Q20) D

Module 7

Virtual Server Networking

Overview

Virtualization has become a critical requirement of data center networking. This module introduces the Cisco Nexus 1000V Distributed Virtual Switch (DVS), a plug-in designed by Cisco that operates within the VMware vCenter environment. Implementers will be expected to successfully integrate Cisco Nexus 1000V in a VMware environment that runs on a Cisco Unified Computing System (Cisco UCS) infrastructure.

The module contrasts the capabilities, features, and benefits of the Cisco Nexus 1000V DVS with the VMware vNetwork Standard Switch (vSwitch) and vNetwork Distributed Switch (vDS). The module presents the installation method for the Cisco Nexus 1000V Virtual Supervisor Module (VSM), as well as the creation of port profiles and the application of policies to virtual machines (VMs). The module also introduces the Cisco Nexus 1010 Virtual Services Appliance and compares its capabilities to the Cisco Nexus 1000V VSM. This module also covers the configuration of Cisco Virtual Network Link (VN-Link) in hardware with VMware Pass-Through Switching (PTS).

Module Objectives

Upon completing this module, you will be able to describe the capabilities, features, and benefits of the Cisco Nexus 1000V switch; the installation method and capabilities of the Cisco Nexus 1000V VSM; and the configuration of VN-Link in hardware with VMware PTS. This ability includes being able to meet these objectives:

- Describe the Cisco Nexus 1000V switch and its role in a virtual server networking environment
- Describe the VMware networking environment
- Describe the Cisco Nexus 1000V network architecture
- Install and configure a Cisco Nexus 1000V switch
- Configure basic Cisco Nexus 1000V networking
- Configure Cisco UCS Manager for VMware PTS

Lesson 1

Evaluating the Cisco Nexus 1000V Switch

Overview

Implementers must be able to compare the components that comprise the Cisco Nexus 1000V Distributed Virtual Switch (DVS). This lesson describes VMware virtual networking options, including VMware vNetwork Standard Switch (vSwitch) and vNetwork Distributed Switch (vDS). The lesson compares these options to the Cisco Nexus 1000V DVS and presents its capabilities and switching components.

Objectives

Upon completing this lesson, you will be able to compare VMware virtual networking options and the Cisco Nexus 1000V DVS. This ability includes being able to meet these objectives:

- Describe the Cisco virtual switching solution for the VMware vDS
- Describe the switching components of the Cisco Nexus 1000V DVS

Cisco Virtual Switching Overview

This topic provides an overview of Cisco virtual switching.

Evolution of Virtual Networking—Before Virtualization

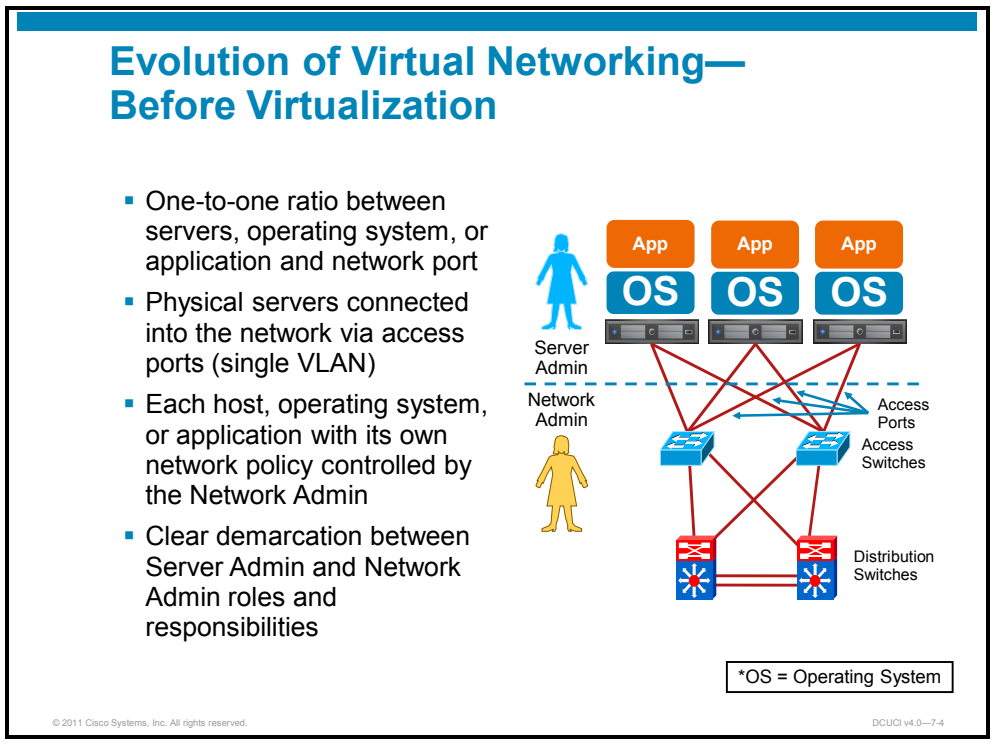

Before virtualization, each server ran its own operating system, usually with a single application running in addition to the operating system. The network interface cards (NICs) were connected to access layer switches to provide redundancy. Network security, quality of service (QoS), and management policies were created on these access layer switches and applied to the access ports that corresponded to the appropriate server.

If a server needed maintenance or service, it was disconnected from the network. During that time, any critical applications needed to be manually offloaded to another physical server. Connectivity and policy enforcement were very static and seldom required any modifications.

Server virtualization has made networking, connectivity, and policy enforcement much more challenging. By using VMware vMotion, the devices that run the applications can move from one physical host to another, which leads to several challenges:

- Providing network visibility from the virtual machine (VM) virtual NIC (vNIC) to the physical access switch
- Creating and then applying policies to the vNICs
- Providing consistent mobility of the policies that are applied to the VMs during a vMotion event

Further complications exist because objects have been virtualized and can move around.

Evolution of Virtual Networking—Virtual Switches

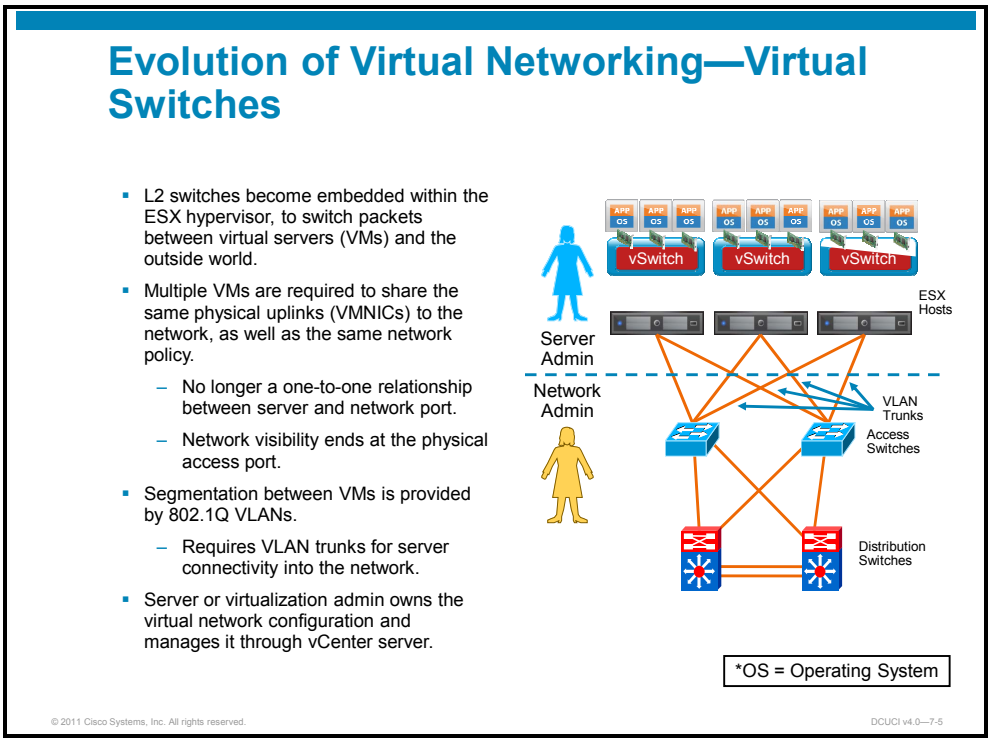

The VMware server-virtualization solution extends the access layer into the VMware ESX server by using the VM networking layer. Several components are used to implement server-virtualization networking:

- **Physical networks:** Physical devices connect ESX hosts for resource sharing. Physical Ethernet switches are used to manage traffic between ESX hosts, like in a regular LAN environment.
- **Virtual networks:** Virtual devices run on the same system for resource sharing.
- **Virtual Ethernet switch (vSwitch):** Like a physical switch, the vSwitch maintains a table of connected devices. This table is used for frame forwarding. The vSwitch can be connected, via uplink, to a physical switch by using a physical VM NIC (VMNIC). The vSwitch does not provide the advanced features of a physical switch.
- **Port group:** A port group is a subset of ports on a vSwitch for VM connectivity.
- **Physical NIC (VMNIC):** The VMNIC is used to uplink the ESX host to the external network.

vNetwork Distributed Switch

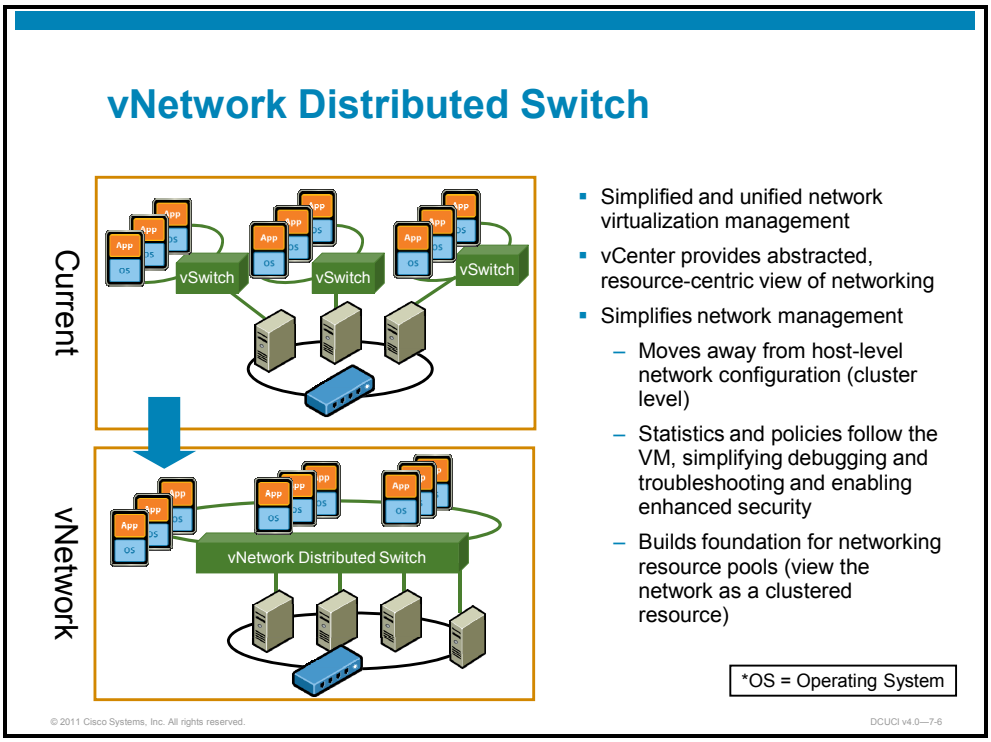

VMware vSphere 4 introduces the vDS—a DVS. With vDS, multiple vSwitches within an ESX cluster can be configured from a central point. The vDS automatically applies changes to the individual vSwitches on each ESX host.

The feature is licensed and relies on VMware vCenter server. The vDS cannot be used for individually managed hosts.

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

VMware vNetwork Evolution

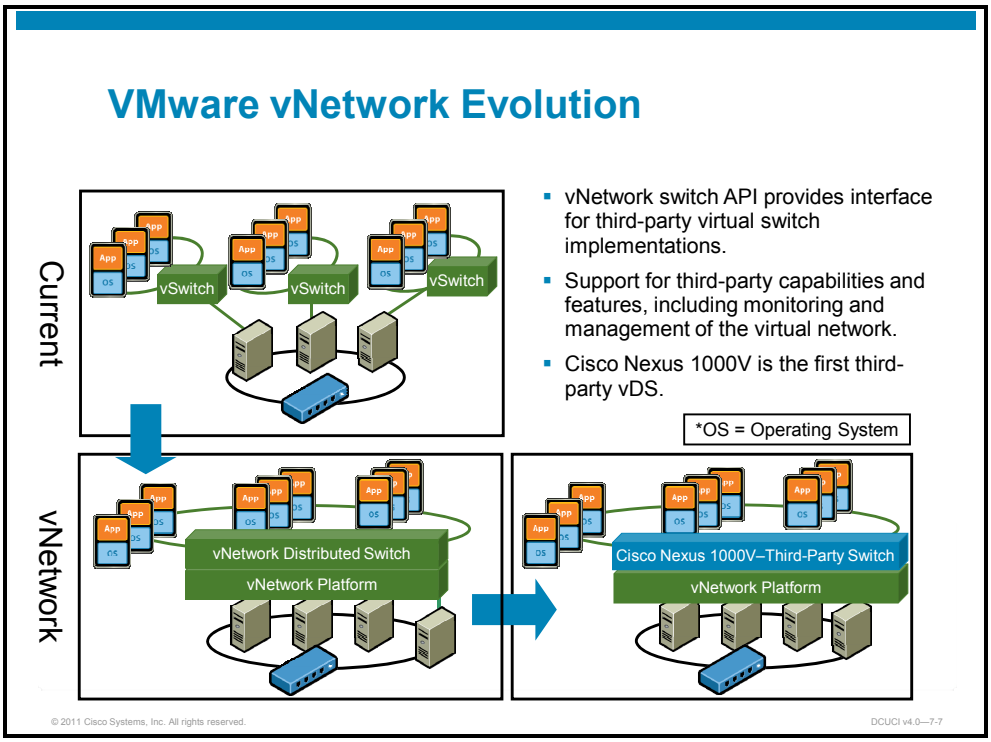

The Cisco server virtualization solution uses technology that was jointly developed by Cisco and VMware. The network access layer is moved into the virtual environment to provide enhanced network functionality at the VM level.

This solution can be deployed as a hardware- or software-based solution, depending on the data center design and demands. Both deployment scenarios offer VM visibility, policy-based VM connectivity, policy mobility, and a nondisruptive operational model.

Cisco Nexus 1000V

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

VN-Link

Cisco Virtual Network Link (VN-Link) technology was jointly developed by Cisco and VMware and has been proposed to the IEEE for standardization. The technology is designed to move the network access layer into the virtual environment to provide enhanced network functionality at the VM level.

Cisco UCS 6100

With the Cisco Unified Computing System (UCS) 6100 Series Fabric Interconnects, VN-Link can be deployed as a hardware-based solution that offers VM visibility, policy-based VM connectivity, policy mobility, and a nondisruptive operational model.

Distributed Virtual Networking

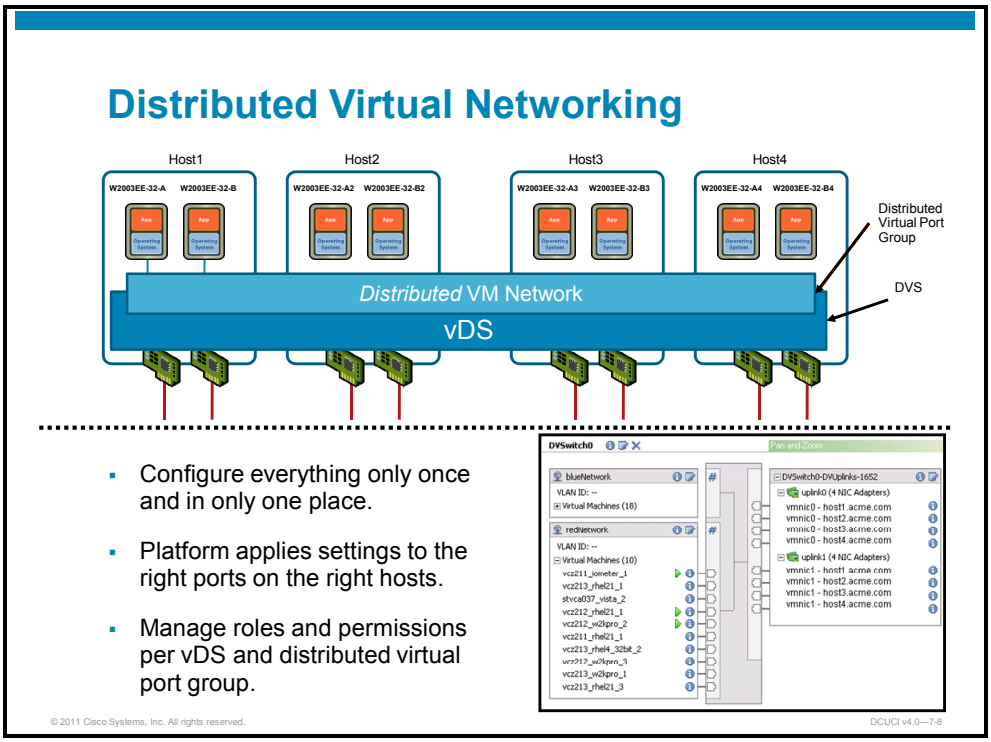

The vDS adds functionality and simplified management to the VMware network. The vDS adds the ability to use private VLANs (PVLANs), perform inbound rate limiting, and track VM port state with migrations. Additionally, the vDS is a single point of network management for VMware networks. The vDS is a requirement for the Cisco Nexus 1000V DVS.

Virtual Switch Options with vSphere 4

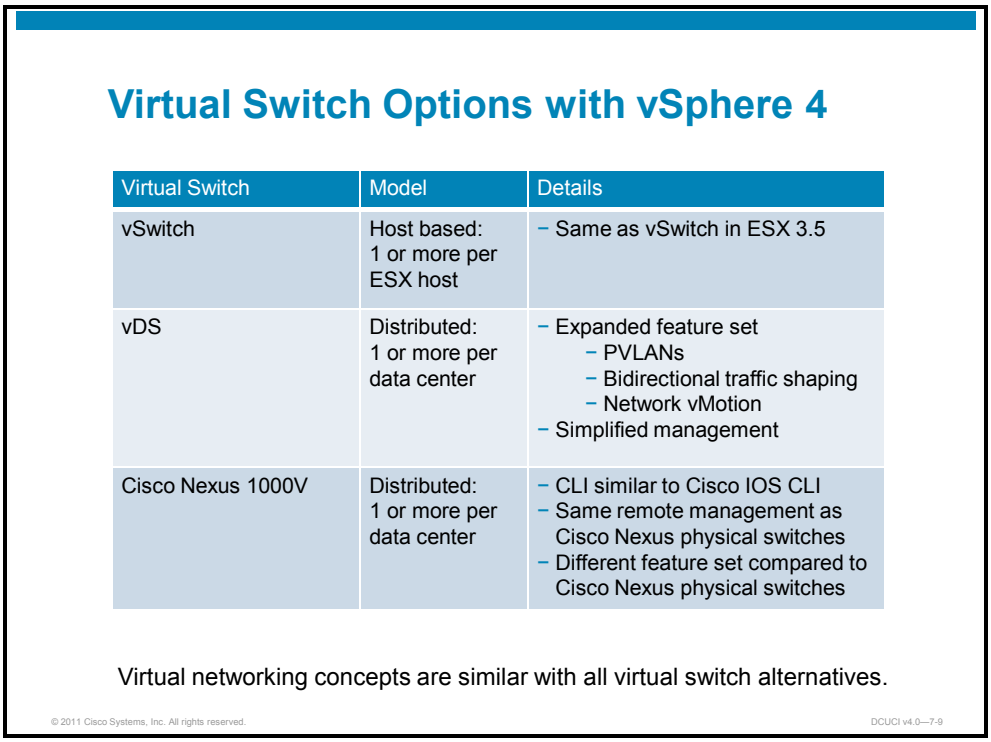

With the introduction of vSphere 4, VMware customers can enjoy the benefits of three virtual networking solutions: vSwitch, vDS, and the Cisco Nexus 1000V.

The Cisco Nexus 1000V bypasses the VMware vSwitch with a Cisco software switch. This model provides a single point of configuration for the networking environment of multiple ESX hosts. Additional functionality includes policy-based connectivity for the VMs, network security mobility, and a nondisruptive software model.

VM connection policies are defined in the network and applied to individual VMs from within vCenter. These policies are linked to the Universally Unique ID (UUID) of the VM and are not based on physical or virtual ports.

Cisco Nexus 1000V Virtual Switching Feature Overview

This topic provides an overview of the virtual switching feature of the Cisco Nexus 1000V solution.

Cisco Nexus 1000V Features

Cisco Nexus 1000V has the same Cisco IOS command-line interface (CLI) and remote management capabilities as Cisco Catalyst and Cisco Nexus physical switches, but it provides a feature set that has been optimized for VMware virtualized environments. The Cisco Nexus 1000V solution maintains compatibility with VMware advanced services such as vMotion, Distributed Resource Scheduler (DRS), Fault Tolerance (FT), and High Availability (HA). The solution extends VM visibility to the physical access switch, while maintaining the rich, advanced feature set offered by physical switches.

From an operations perspective, the network administrator may manage the Cisco Nexus 1000V solution from a console connection or remotely by using Telnet or Secure Shell (SSH). The solution preserves the network administrator function and enables administrators to effectively create advanced security, QoS, and VLAN policies. Because the Cisco Nexus 1000V switch communicates directly with vCenter, configuration changes are reflected directly within the vCenter inventory. Therefore, the server administrator can consume the policies by applying them to either virtual uplinks or VM vNICs.

VM View of Resources

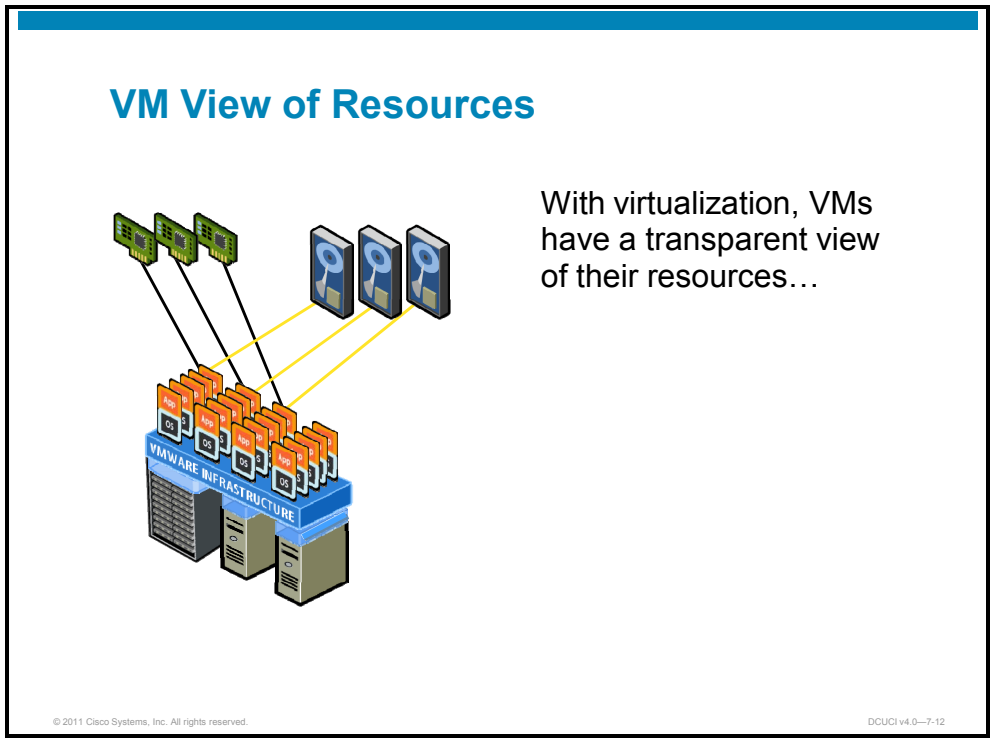

Server virtualization has created many networking challenges, especially at the VM level. VMs have a transparent view of their CPU, memory, storage, and networking resources because the hypervisor services all resource requests that the VMs make.

As a result, creating and assigning advanced networking policies (such as security, QoS, port channels, and so on) to virtual networks at the VM level has been difficult.

VM View of Resources (Cont.)

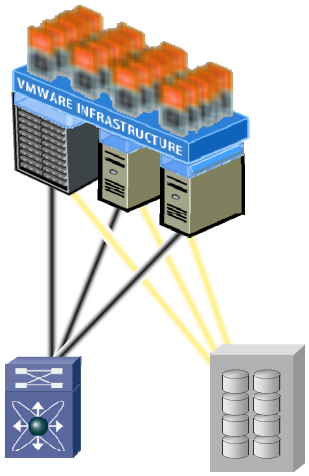

…but correlating network and storage back to VMs is difficult.

VMware vSphere enables CPU and memory to be reserved for VMs that run critical applications. However, this resource-reservation mechanism applies to neither storage nor networking resources.

Because a single VMNIC uplink may be shared, it becomes difficult to identify multiple virtual I/O streams from a single, physical server uplink.

VM Transparency

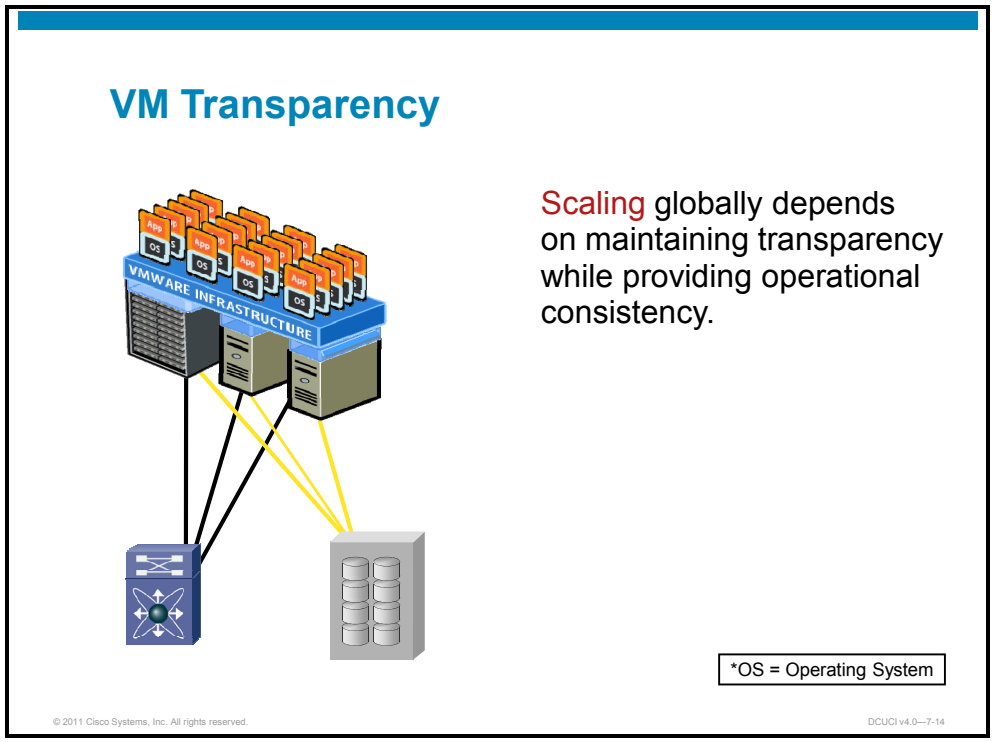

VMs are likely to move from one assigned ESX host to another. The ability to identify unique, individual VMs and treat them differently from a security and network-performance perspective would also be valuable.

To deliver operational consistency, three important parameters must be supported:

- Providing VM visibility to the physical access switch
- Having the ability to maintain the rich, advanced feature set that is offered by high-performance access switches
- Enabling the application of the advanced feature policies at the VM level and assuring that they move with the VM in a vMotion event

Scaling Server Virtualization

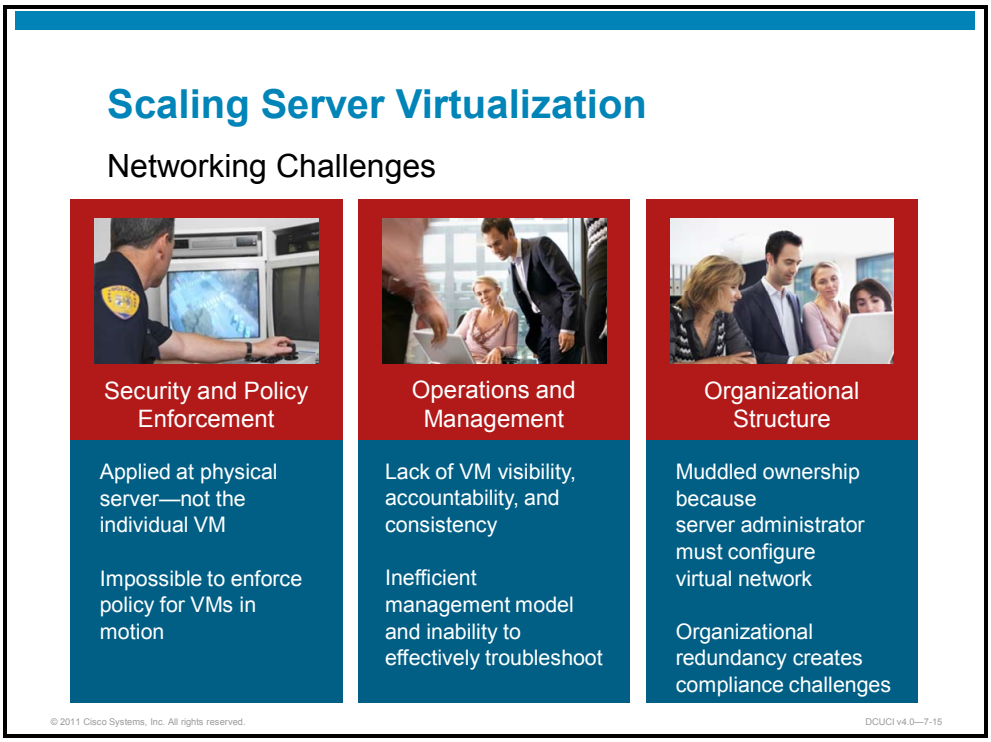

Modern virtualized environments introduce problems for the network. As in a physical environment, the server connects to the network switch via one or more NICs. However, in a virtual environment, each NIC carries traffic for multiple VMs. Therefore, the network switch has no way of knowing from which VM the traffic is being sent.

In addition, from an operations and management perspective, server administrators prefer to perform server tasks, whereas network administrators prefer to create, apply, and manage policies that govern connectivity and network performance.

So, harmony between the skill sets of server and network administrators requires a clear delineation of responsibility.

Server virtualized networking must be able to extend and mirror the robust, advanced feature sets that are offered by high-performance access layer switches.

VN-Link Brings VM-Level Granularity

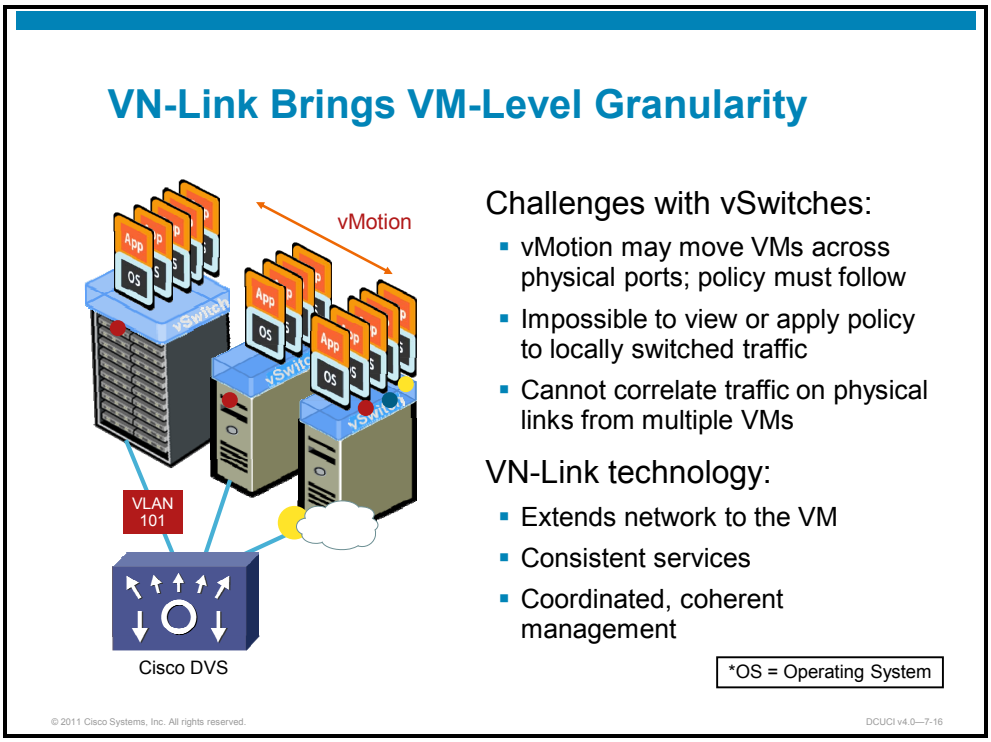

VMware vMotion is a feature that is used within ESX environments. This feature permits you to move VMs, either automatically or manually (via notifications), from one physical machine to another, as resources become overutilized in response to a physical server failure or a VMware fault-tolerance event. This process can involve several issues:

- The policy that is associated with the VM must be able to follow the VM to the new physical location.
- Viewing or applying policies to locally switched traffic when using the vSwitch that is internal to the ESX host can be difficult.
- Differentiating traffic on physical lines from multiple VMs can be difficult.
- Determining with which VLAN the VMs should be associated can be difficult.

Cisco VN-Link uses the vDS framework to deliver a portfolio of networking solutions that can operate directly within the distributed hypervisor layer, offering a feature set and operational model that is familiar and consistent with other Cisco networking products. Cisco VN-Link specifically enables individual VMs to be identified, configured, monitored, migrated, and diagnosed in a way that is consistent with current network operational models.

VN-Link indicates the creation of a logical link between a vNIC on a VM and a Cisco switch that is enabled for VN-Link. This logical creation is the equivalent of using a cable to connect a NIC to a port on an access layer switch.

A switch that is enabled for VN-Link uses the concept of virtual Ethernet (vEthernet) interfaces, which are dynamically provisioned based on network policies that are stored on the switch. These policies are the result of VM provisioning operations by the hypervisor management layer (vCenter). The vEthernet interface maintains network configuration attributes, security, and statistics for a given virtual interface across mobility events.

VN-Link with the Cisco Nexus 1000V

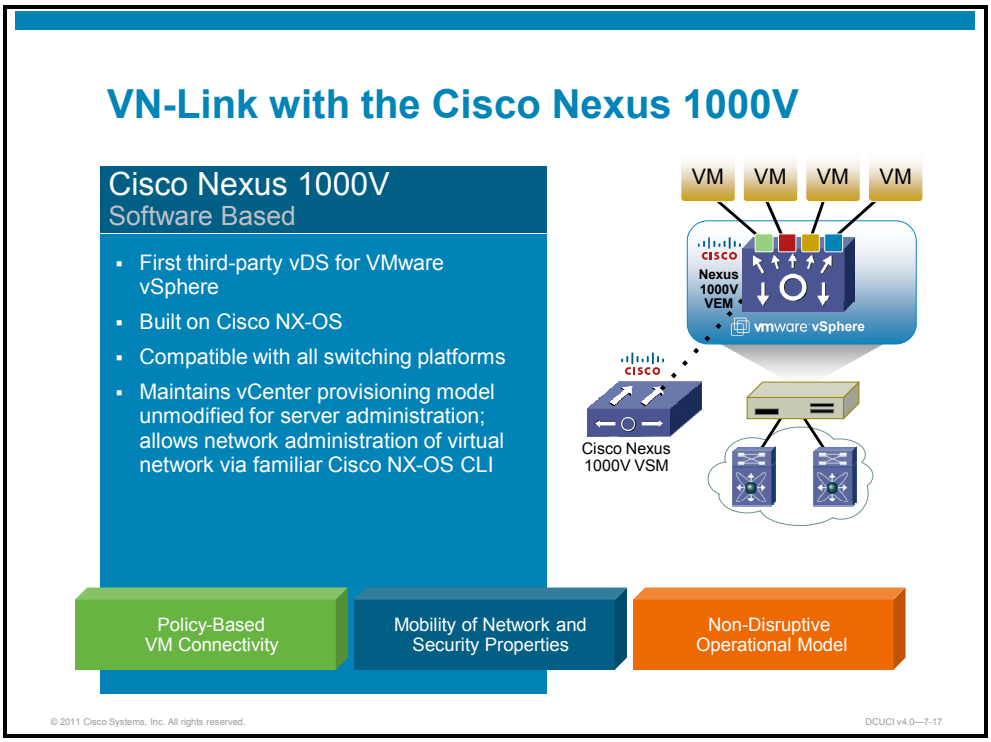

With the introduction of the vDS framework, VMware permits third-party networking vendors to provide their own implementation of distributed virtual switches. When deploying the Cisco Nexus 1000V solution, the vSwitch and port group configuration is offloaded to the network administrator. This process helps to ensure a consistent network policy throughout the data center.

In the Cisco Nexus 1000V, traffic between VMs is switched locally at each instance of a VEM. Each VEM is responsible for interconnecting the local VMs with the rest of the network, through upstream access layer network switches. The Cisco Nexus 1000V Virtual Supervisor Module (VSM) manages the control plane protocols and configuration of the VEMs. The VSM never takes part in the actual forwarding of packets.

Summary

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

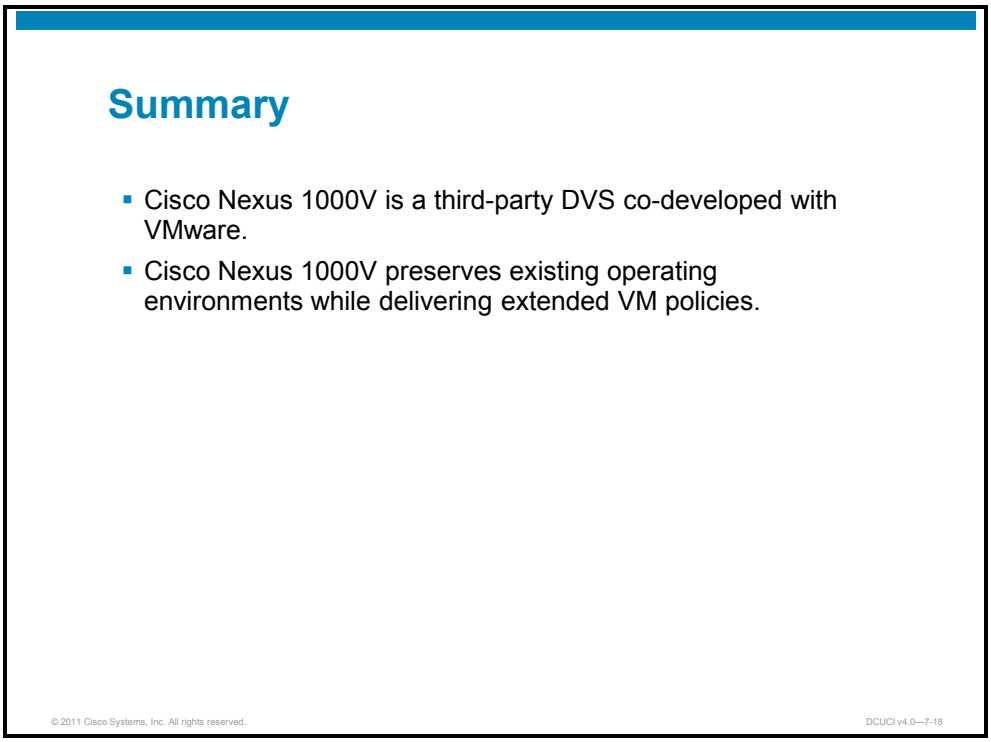

Lesson 2

Working with VMware Ethernet Networking

Overview

Implementers must be able to compare the VMware virtual switching offerings (vSwitch and Distributed Virtual Switch [DVS]) and articulate how deployment of Cisco Nexus 1000V DVS offers greater security, scalability, manageability, and configuration consistency.

This lesson presents the integration of the Cisco Nexus 1000V solution within the VMware distributed switching environment. The lesson also contrasts the unique features, benefits, and capabilities of the Cisco Nexus 1000V solution with the Cisco Unified Computing System virtualization adapter and discusses the differences in how these solutions implement Cisco Virtual Network Link (VN-Link).

Objectives

Upon completing this lesson, you will be able to describe the unique features of the Cisco Nexus 1000V solution and compare it to the Cisco Unified Computing System virtualization adapter. This ability includes being able to meet these objectives:

- Describe how the Cisco Nexus 1000V integrates into VMware vDS
- Describe the Cisco DVS solution
- Describe the unique features that the Cisco Nexus 1000V brings to VMware vDS
- Differentiate between the capabilities of the Cisco Unified Computing System virtualization adapter and the Cisco Nexus 1000V

VMware vDS

This topic discusses the use of the distributed virtual switch (DVS) in VMware environments.

vDS Configuration

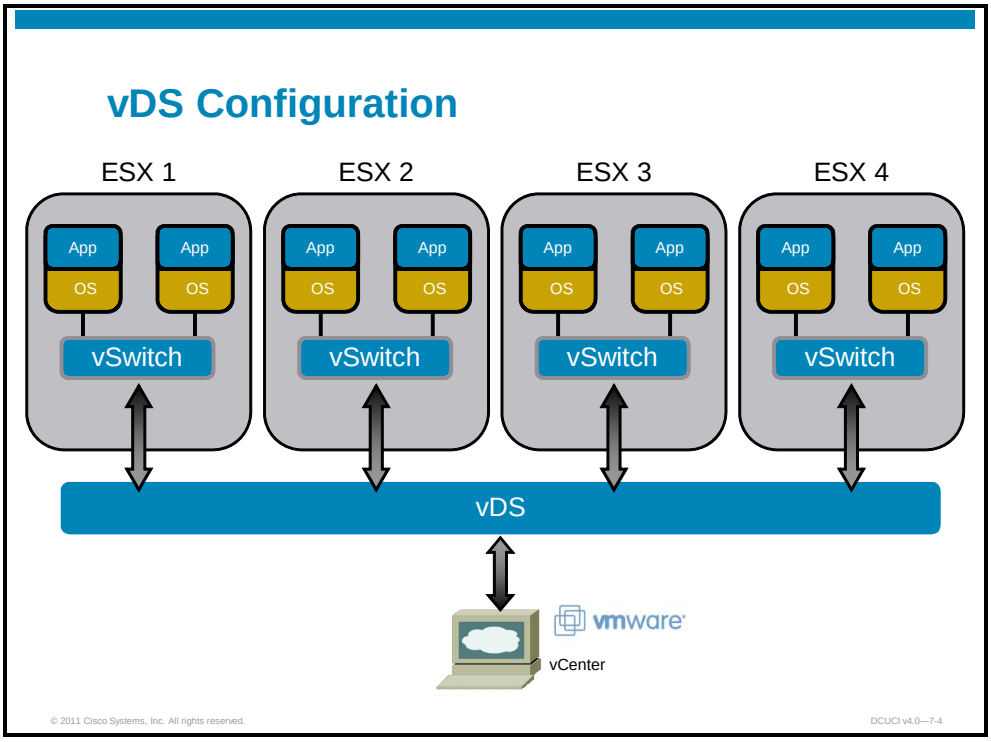

The VMware vNetwork Distributed Switch (vDS) extends the features and capabilities of virtual networks. At the same time, the vDS simplifies provisioning and the ongoing process of configuration, monitoring, and management.

With VMware ESX 3.5 and prior releases, virtual networks were constructed by using VMware vNetwork Standard Switches (vSwitches). Each ESX host would use one or more vSwitches to connect the virtual machines (VMs) with the server network interface cards (NICs) and the outside physical network.

Distributed Virtual Switching

Distributed Virtual Switching

- Single point of control
- Simplified network management
- Removes the requirement for host-level network configuration
- Statistics and policies follow the VM
- Enables network resource pools (view the network as a clustered resource)

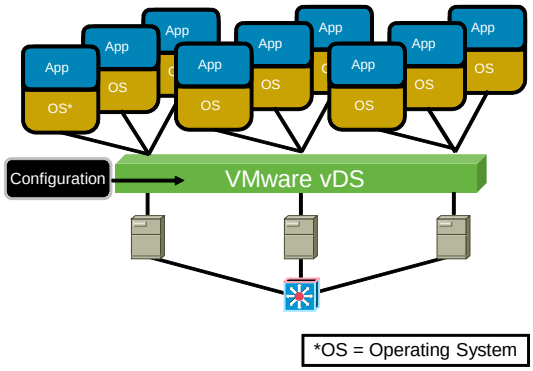

Available with the release of vSphere 4.

In addition to continuing support for the vSwitch, VMware vSphere introduces an additional choice for VMware virtual networking: the vDS. The vDS eases the management burden of per-host virtual-switch configuration management by treating the network as an aggregated resource. Individual, host-level virtual switches are abstracted into one large vDS that spans multiple hosts at the datacenter level. Port groups become distributed virtual port groups that span multiple hosts and ensure configuration consistency for the VMs and virtual ports that are necessary for such functions as VMware vMotion.

Virtual Network Configuration

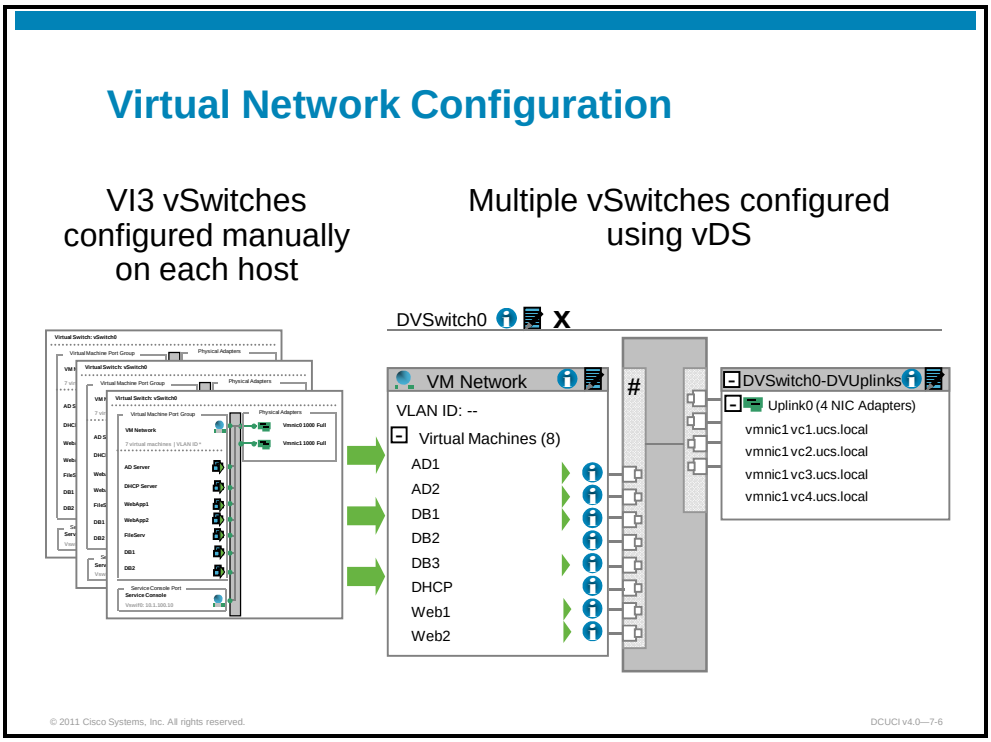

The figure shows the conceptual difference in management of a standard vSwitch environment versus a vDS environment. The vSwitch requires a separate configuration from a separate management panel. The vDS requires just one management panel for the single switch that spans multiple hosts.

vDS Enhancements

vDS Enhancements

- The VMware vDS offers several enhancements to VMware switching:
  - Port state migration (statistics and port state follow VM)
  - Rx rate limiting
  - PVLANs

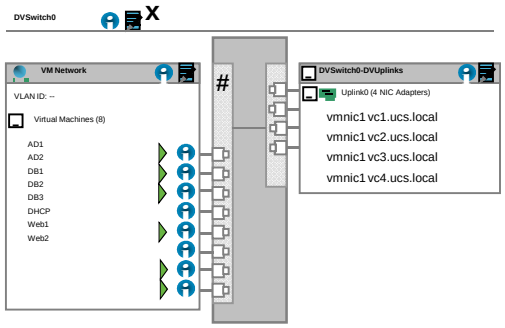

Private VLAN (PVLAN) support enables broader compatibility with existing networking environments that use PVLAN technology. PVLANs enable users to restrict communication between VMs on the same VLAN or network segment. This ability significantly reduces the number of subnets that are needed for certain network configurations.

PVLANs are configured on a DVS, with allocations made to the promiscuous PVLAN, the community PVLAN, and the isolated PVLAN. Distributed virtual port groups can then use one of these PVLANs, and VMs are assigned to a distributed virtual port group. Within the subnet, VMs on the promiscuous PVLAN can communicate with all VMs. VMs on the community PVLAN can communicate among themselves and with VMs on the promiscuous PVLAN. VMs on the isolated PVLAN can communicate only with VMs on the promiscuous PVLAN.

Note	Adjacent physical switches must support PVLANs and must be configured to support the PVLANs that are allocated on the DVS.

Network vMotion is the tracking of VM networking state (for example, counters or port statistics) as the VM moves from host to host on a vDS. This tracking provides a consistent view of a virtual network interface, regardless of the VM location or vMotion migration history. This view greatly simplifies network monitoring and troubleshooting activities when vMotion is used to migrate VMs between hosts.

DVS expands upon the egress-only traffic-shaping feature of standard switches by providing bidirectional traffic-shaping capabilities. Egress (from VM to network) and ingress, or receive (Rx) rate limiting, (from network into VM) traffic-shaping policies can now be applied on distributed virtual port group definitions.

Traffic shaping is useful when you want to limit the traffic to or from a VM or group of VMs, to protect a VM or other traffic in an oversubscribed network. Policies are defined by three characteristics: average bandwidth, peak bandwidth, and burst size.

vSwitch and vDS

VMware vDS and vSwitch are not mutually exclusive.

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

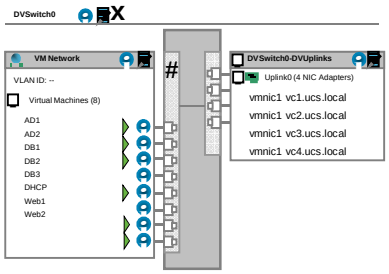

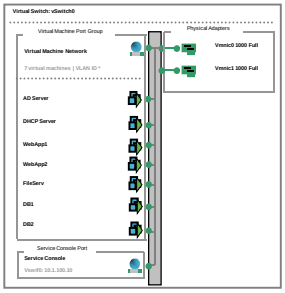

The VMware vSwitch and vDS are not mutually exclusive and can coexist within the same VMware vCenter management environment. Physical VM NICs (VMNICs) may be assigned to either the vSwitch or the vDS on the same VMware ESX or ESXi host.

You can also migrate the ESX service console and VMware VMkernel ports from the vSwitch, where they are assigned by default during ESX installation, to the vDS. This migration facilitates a single point of management for all virtual networking within the vCenter data-center object.

Cisco Nexus 1000V DVS

This topic discusses Cisco Nexus 1000V DVS.

Cisco Nexus 1000V Components

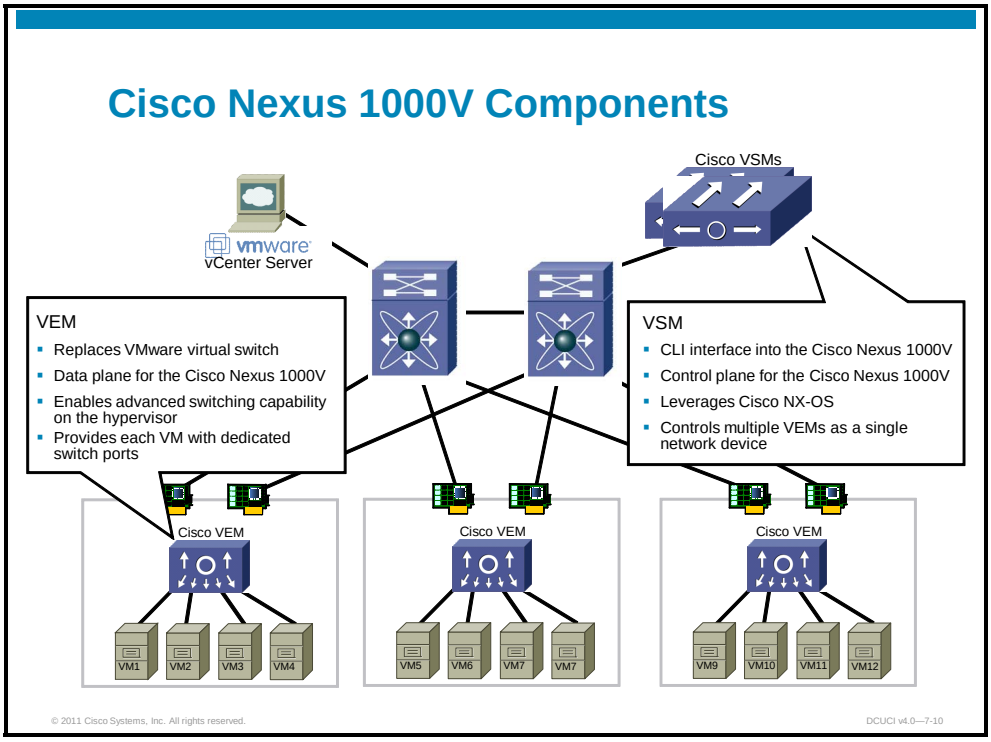

The Cisco Nexus 1000V provides Layer 2 switching functions in a virtualized server environment. Cisco Nexus 1000V DVS replaces virtual switches within the ESX servers. This replacement allows users to configure and monitor the virtual switch by using the Cisco Nexus Operating System (NX-OS) command-line interface (CLI). The Cisco Nexus 1000V also provides visibility into the networking components of the ESX servers and access to the virtual switches within the network.

The vCenter server defines the data center that the Cisco Nexus 1000V will manage. Each server is represented as a line card and is managed as if it were a line card in a physical Cisco switch.

Two components are part of the Cisco Nexus 1000V implementation:

- Virtual Supervisor Module (VSM): The Cisco Nexus 1000V VSM is the control software of the Cisco Nexus 1000V DVS. The VSM runs either on a VM or as an appliance and is based on Cisco NX-OS.
- Virtual Ethernet Module (VEM): The Cisco Nexus 1000V VEM actually switches the data traffic and runs on a VMware ESX 4.0 host. VSM can control several VEMs, with the VEMs forming a switch domain that should be in the same virtual data center that is defined by VMware vCenter.

Cisco Nexus 1000V—Single Chassis Management

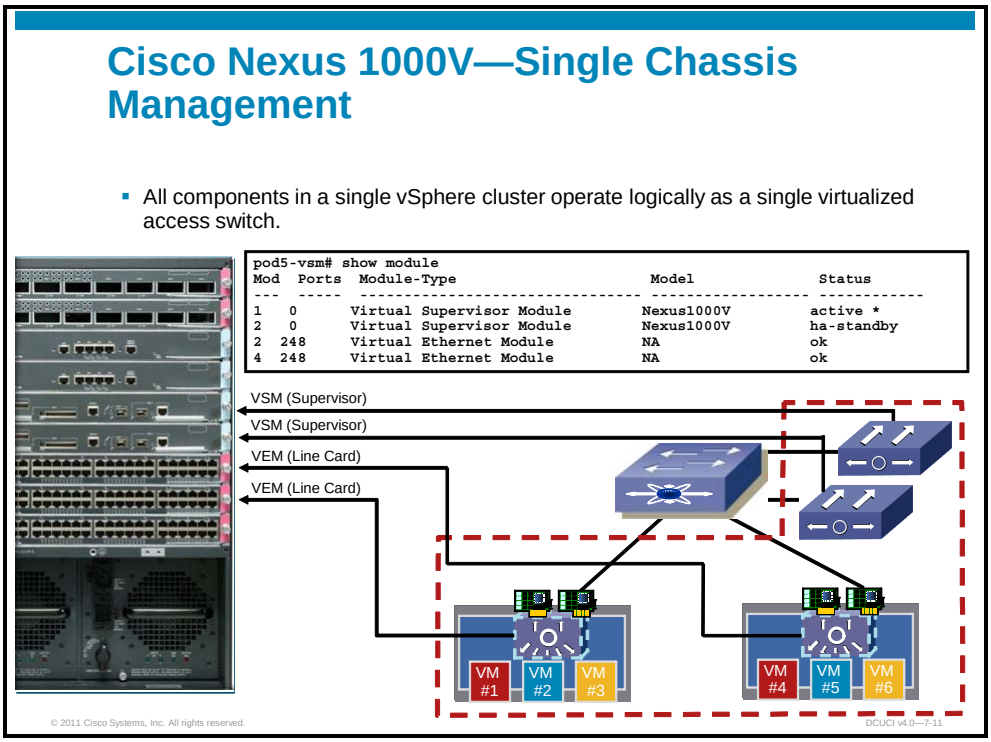

Cisco Nexus 1000V is effectively a virtual chassis: It is modular, and ports can be physical or virtual. The servers are modules on the switch, with each physical NIC port on a module being a physical Ethernet port. Modules 1 and 2 are reserved for the VSM, and the first server or host is assigned automatically to the next available module number. The ports to which the vNIC interfaces connect are virtual ports on the Cisco Nexus 1000V and are assigned a global number.

Cisco Nexus 1000V—Single Chassis Management (Cont.)

```
Upstream-Switch#show cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID            Local Intrfce         Holdtme    Capability    Platform      Port ID

N1KV-Rack10          Eth 1/8               136             S       Nexus 1000V   Eth2/2
N1KV-Rack10          Eth 2/10              136             S       Nexus 1000V   Eth3/2
```

VSMs and VEMs appear as one switch from a management and control plane perspective.

- Protocols such as Cisco Discovery Protocol and SNMP operate as a single switch.
- These control protocols are run on the VSM and carried over the packet VLAN.

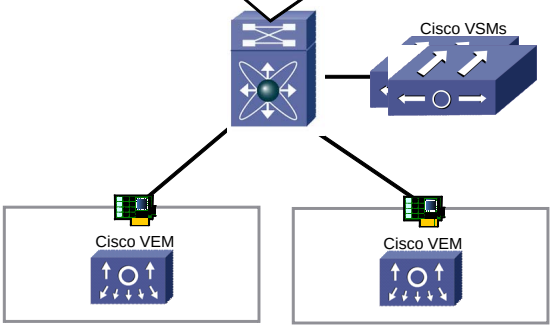

To the upstream switches, the Cisco Nexus 1000V appears as a single switch from the control and management plane. Protocols such as Cisco Discovery Protocol and Simple Network Management Protocol (SNMP) operate as a single switch.

The **show cdp neighbor** command that is shown in the figure indicates that two VEMs are associated with the virtual switch. Therefore, two ports connect the Cisco Nexus 1000V to the upstream switch.

Cisco Nexus 1000V—VSM Deployment Options

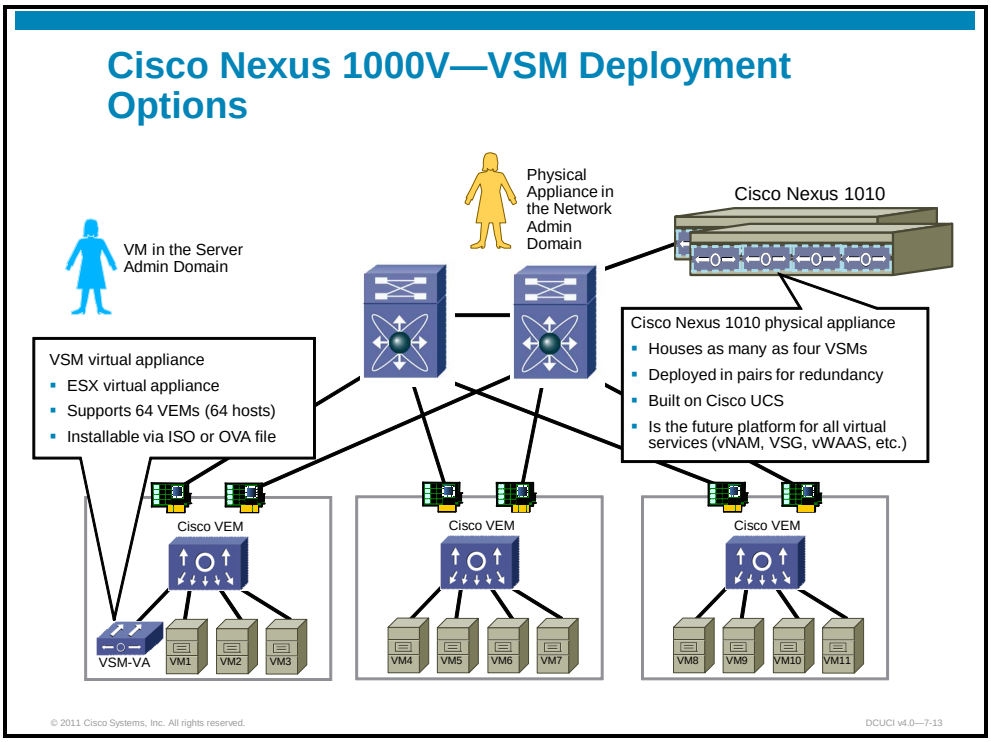

There are two VSM options.

- **VSM virtual appliance:** The VSM runs on an ESX host as an ESX virtual appliance, with support for 64 VEMs. Installation of the VSM virtual appliance is provided though an ISO or Open Virtual Appliance (OVA) file.
- **VSM physical appliance:** A Cisco Nexus 1010 Physical Server can host four VSM virtual appliances. The VSM physical appliance typically is deployed in pairs, for redundancy purposes.

Cisco Nexus 1000V—VSM High-Availability Options

Cisco Nexus 1000V—VSM High-Availability Options

VSMs should always be deployed redundantly for high availability.

- Virtual VSM high availability
  - VSM VMs should not reside on the same physical host. This can be accomplished by using VMware anti-affinity rule.
- Cisco Nexus 1010 high availability
  - Two Cisco Nexus 1010 systems must be deployed for high availability.
  - High availability pair is formed between two Cisco Nexus 1010 systems, based on control VLAN and domain ID information.
  - Cisco Nexus 1010 software takes control of load balancing active/standby VSMs between the high-availability pair.

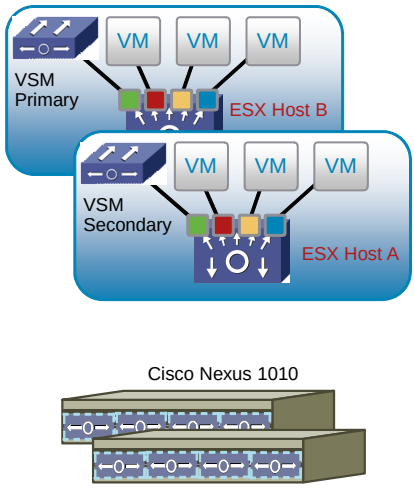

The Cisco Nexus 1000V provides high availability of control plane functionality by using active (primary) and standby (secondary) VSMs.

Both the primary and secondary VSMs install as VMs on ESX or ESXi hosts. To provide maximum redundancy and high availability, each VSM should be installed on separate hosts.

The Cisco Nexus 1010 architecture requires the deployment of two devices in a high-availability configuration.

Cisco Nexus 1000V Communication—Extending the Backplane

Cisco Nexus 1000V Communication—Extending the Backplane

- Because the VSM and VEM are not physically connected, the VSM (supervisor) must program the VEM (line cards) over a network.
- The VSM-to-VEM communication uses the same backplane protocol (AIPC) found in the Cisco Nexus 7000 and MDS platforms.
- There are two ways to extend communication between VSM and VEM:
  - Over Layer 2 cloud, using control and packet VLANs
  - Over Layer 3 cloud, using Layer 3 control capability

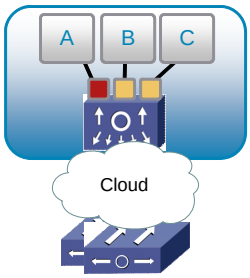

Cisco Nexus 1000V VSM

Maintaining communication between the VSM (control plane) and the individual VEMs (data forwarding) is of paramount importance for successful operation. The control protocol that is used by the VSM to communicate with the VEM is borrowed from the Cisco Nexus data center switch and Cisco Multilayer Director Switch (MDS) 9000 Series Fibre Channel switches.

The VSM-to-VEM communication may be implemented by using a Layer 2 model that uses control and packet VLANs or by using a Layer 3 cloud control capability.

VSM and VEM Communication—Layer 2 Connectivity

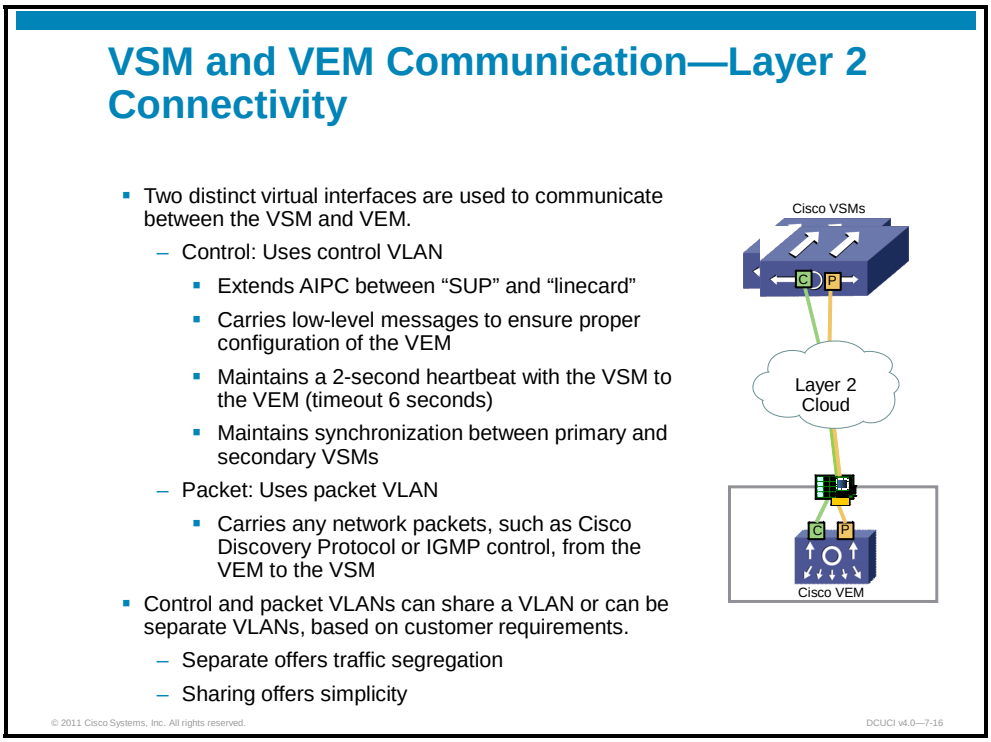

Communication between the VSM and VEM is provided through two distinct virtual interfaces: the control and packet interfaces.

The control interface carries low-level messages to each VEM, to ensure proper configuration of the VEM. A 2-second heartbeat is sent between the VSM and the VEM, with a 6-second timeout. The control interface maintains synchronization between primary and secondary VSMs. The control interface is like the Ethernet out-of-band channel (EOBC) in switches such as Cisco Nexus 7000 Series Switches.

The packet interface carries network packets, such as Cisco Discovery Protocol or Internet Group Management Protocol (IGMP) control messages from the VEM to the VSM.

Customers may choose to implement separate VLANs for the control, management, and packet VLANs, or they can share the same VLAN.

Being VLAN interfaces, the control and packet interfaces require Layer 2 connectivity.

VSM and VEM Communication—Layer 3 Connectivity

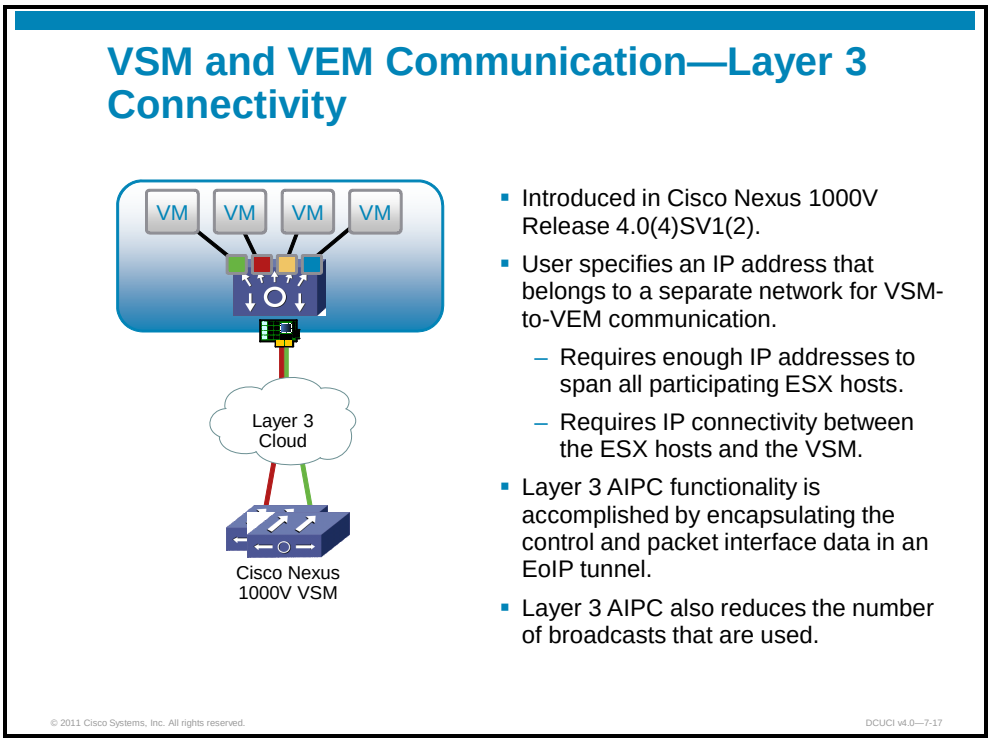

With the release of the Cisco Nexus 1000V Release 4.0(4)SV1(2) image, you can now manage VSM-to-VEM communication by using a Layer 3 network. The administrator specifies a separate IP subnet and sufficient IP addresses to span all participating ESX or ESXi hosts, which requires IP connectivity between the VSM and the ESX or ESXi hosts.

The Cisco Nexus 1000V uses Advanced Interprocess Communications (AIPC), a backplane protocol that the Cisco Nexus 7000 Series Switches also uses. AIPC encapsulates the control and packet interface data in an Ethernet over IP (EoIP) tunnel, which reduces the number of broadcasts that are used.

VSM and VEM Communication—Important Considerations for Layer 3 Control

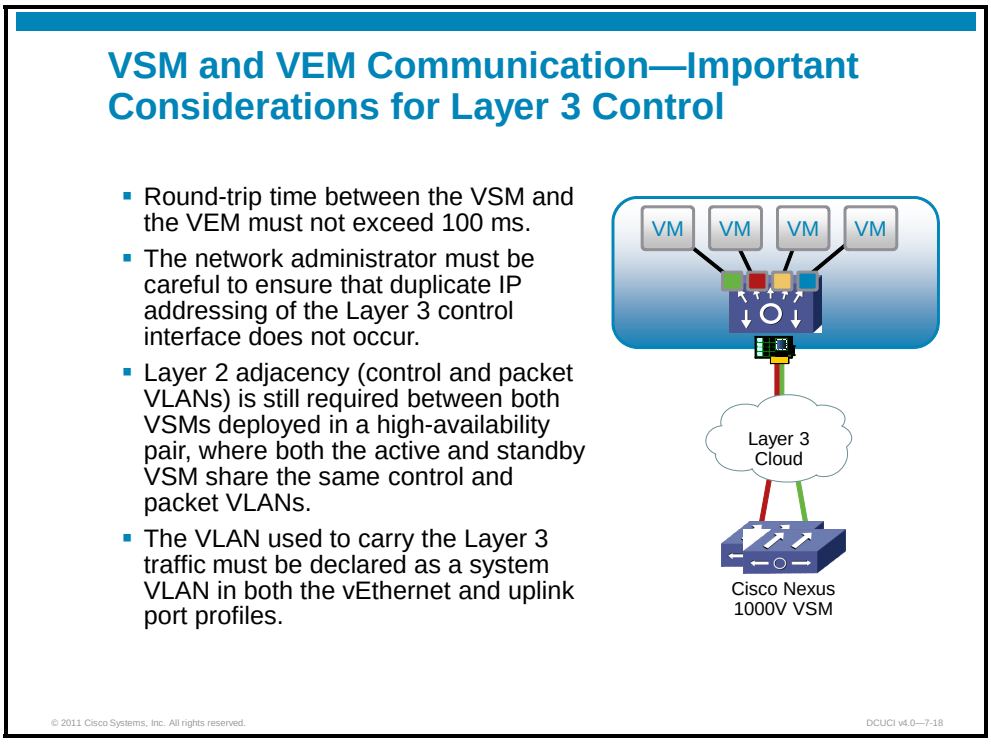

Effective Layer 3 control requires important considerations and requirements. The round-trip time between the VSM and VEM must not exceed 100 ms. The network administrator must also assure that duplicate IP addresses do not appear.

Layer 2 adjacency must be maintained between active and standby VSMs that share the same control and packet VLANs. The VLAN that is configured to carry the Layer 3 traffic must be specified as the system VLAN in both the virtual Ethernet (vEthernet) and uplink port profiles.

Cisco Nexus 1000V Component—vCenter Communication

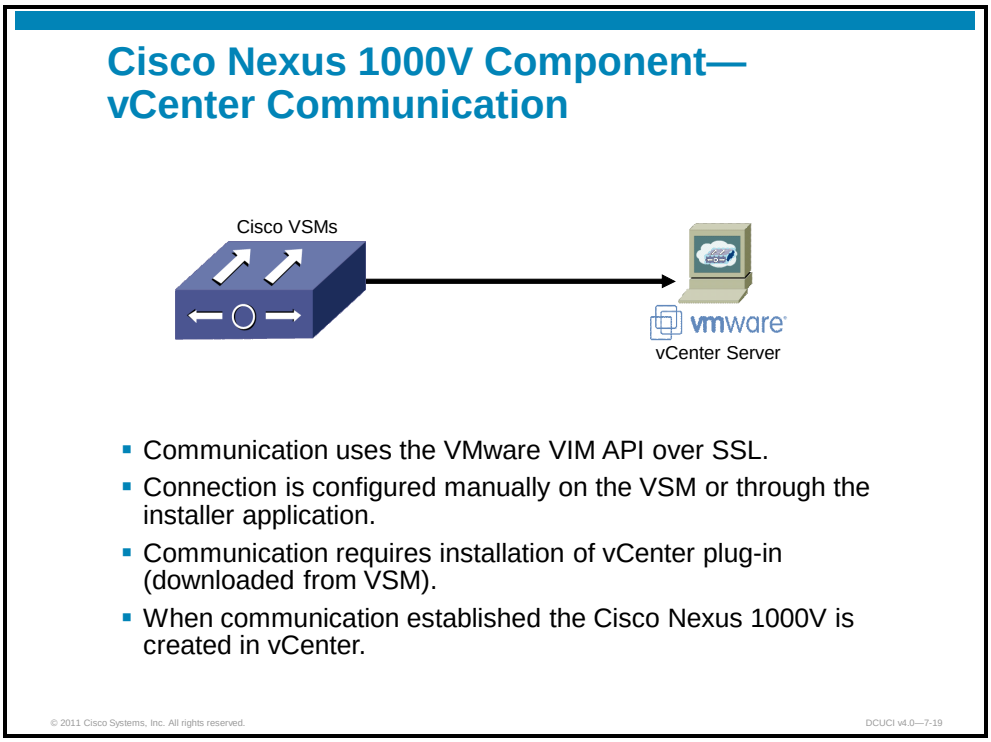

Communication between the VSM and vCenter is provided through the VMware VIM application programming interface (API) over Secure Sockets Layer (SSL). The connection is set up on the VSM and requires installation of a vCenter plug-in, which is downloaded from the VSM.

After communication between the two devices is established, the Cisco Nexus 1000V is created in vCenter.

This interface is known as the out-of-band (OOB) management interface and should be in the same VLAN as your vCenter and host-management VLAN, although that is not a requirement.

Cisco Nexus 1000V—Domain ID

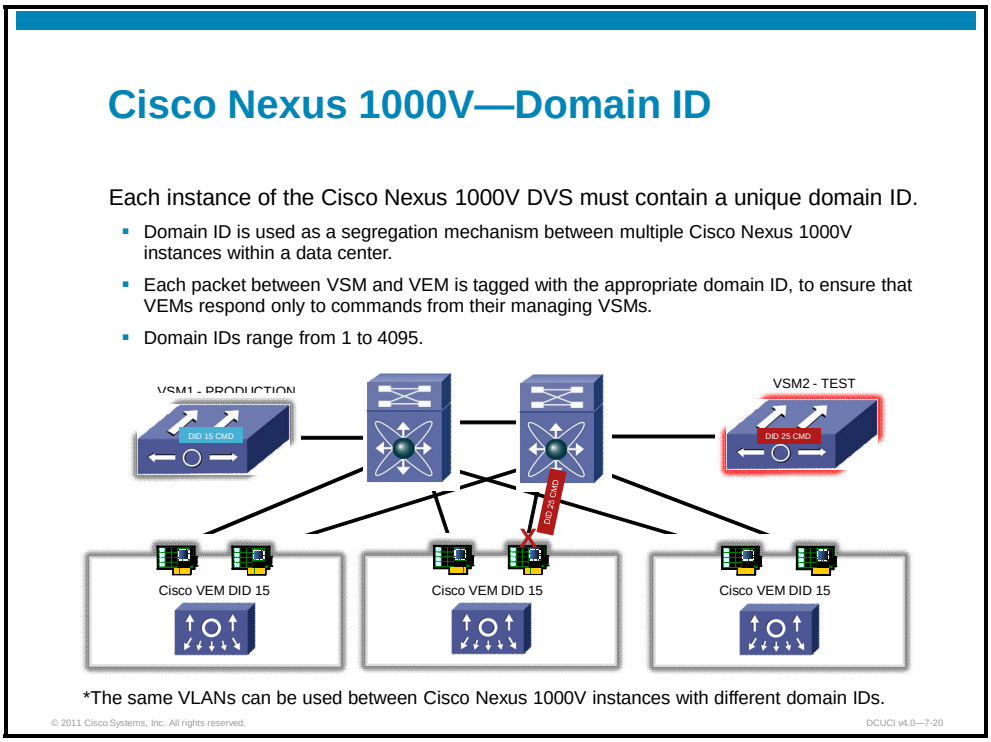

The domain ID configuration setting within the Cisco Nexus 1000V CLI provides a means to uniquely identify and separate multiple instances of the Cisco Nexus 1000V DVS within a vCenter management environment.

The domain ID may be any number between 1 and 4095.

Cisco Nexus 1000V—Opaque Data

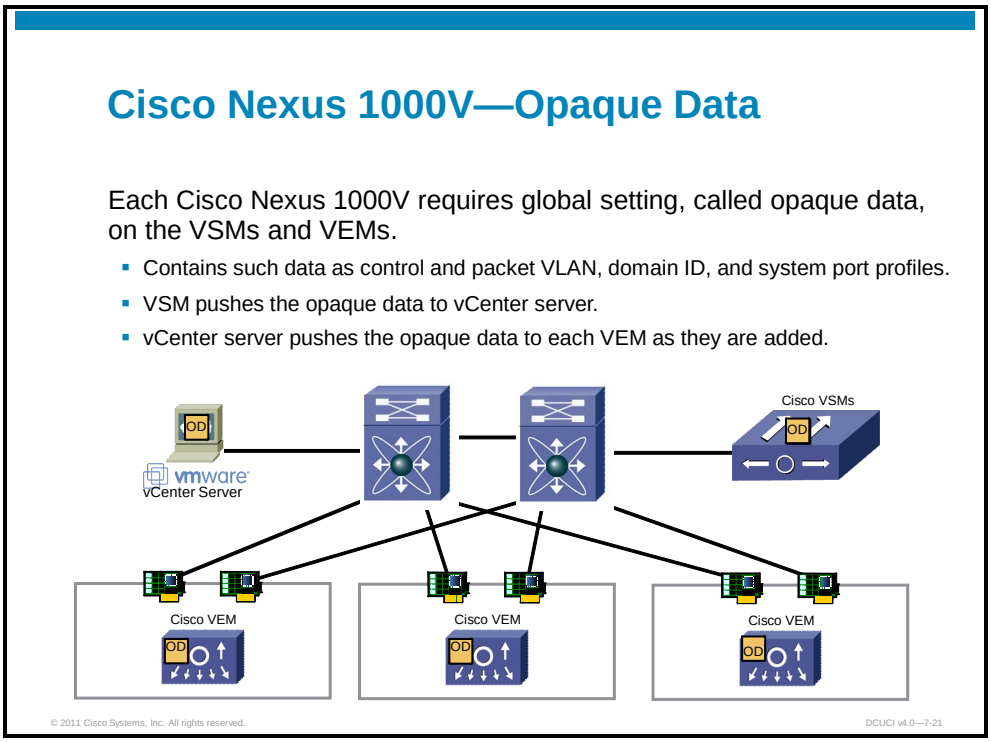

The domain ID is a parameter of the Cisco Nexus 1000V and is used to identify a VSM and VEM that relate to one another. The domain ID of the Cisco Nexus 1000V is defined when the VSM is first installed and becomes part of the opaque data that is transmitted to vCenter. Each command that the VSM sends to any associated VEM is tagged with this domain ID. When a VSM and VEM share a domain ID, the VEM accepts and responds to requests and commands from the VSM. If the VEM receives a command or configuration request that is not tagged with the proper domain ID, then the VEM ignores that request. Similarly, if the VSM receives a packet that is tagged with the wrong domain ID from a VEM, the VSM ignores that packet.

Cisco Nexus 1000V Administrator Roles

This topic discusses administrator roles of the Cisco Nexus 1000V.

The Cisco Nexus 1000V also provides a single, redundant point for managing multiple ESX or ESXi hosts. This management includes the ability to automate policy migration and mobility in virtualized environments by using a feature called port profiles.

Network policies that are enforced by a port profile follow the VM throughout its lifecycle, whether the VM is being migrated from one server to another or is being suspended or restarted. In addition to migrating the policy, the Cisco Nexus1000V moves the network state of the VM, such as the port counters and flow statistics. VMs that participate in traffic-monitoring activities, such as Cisco NetFlow or Encapsulated Remote Switched Port Analyzer (ERSPAN), can continue these activities, uninterrupted by vMotion operations.

Standard VMware Administrator Roles

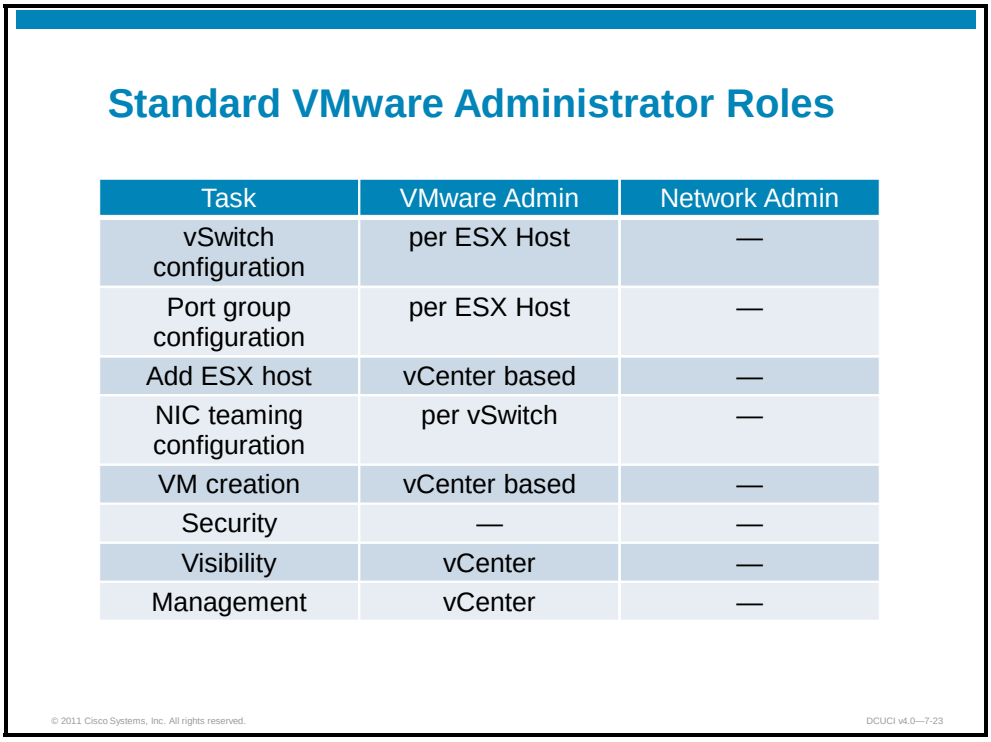

The figure shows the challenges that server administrators face when using the vSwitch or vDS to manage network connectivity and policies. The network administrator has little or no participation, despite the fact that these issues are part of their area of expertise and skill set.

Shifting the burden of network connectivity, policy creation, and configuration from the server to the network administrator vastly improves operations, management, and continuity. Server administrators are tasked only with consuming or assigning the port profile (port groups within the vCenter), which is well within their comfort zone.

Cisco Nexus 1000V Administrator Roles

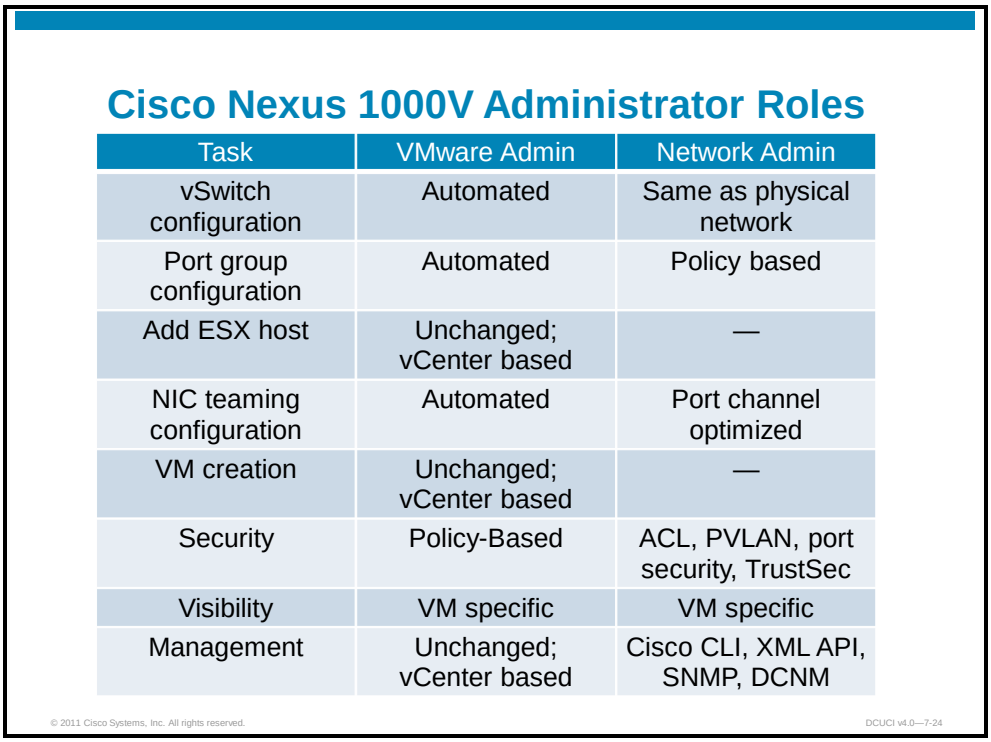

When a new VM is provisioned, the server administrator selects the appropriate port profile. The Cisco Nexus 1000V creates a new switch port that is based on the policies that are defined by the port profile. The server administrator can reuse the port profile to provision similar VMs, as needed. Port profiles are also used to configure the physical NICs in a server. These port profiles, which are known as uplink port profiles, are assigned to the physical NICs as part of the installation of the VEM on an ESX host.

Comparing VN-Link in Software and Hardware

This topic compares the Cisco Nexus 1000V and the Cisco M81KR/P81E Virtual Interface Card (VIC).

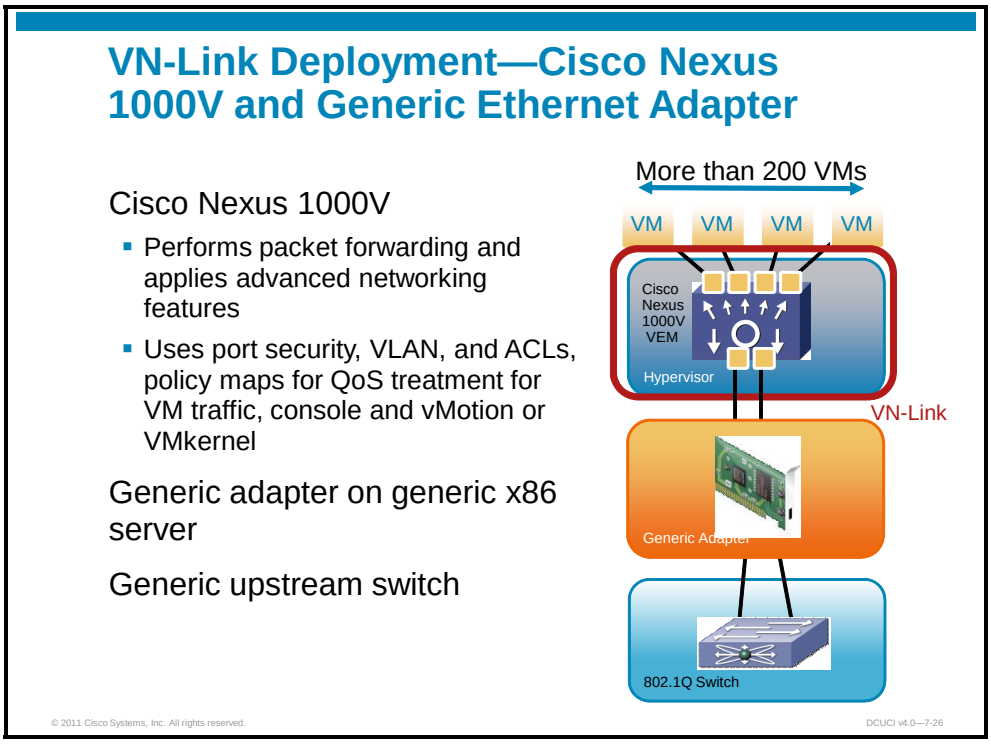

The Cisco Nexus 1000V was developed with VMware to deliver transparency to various server hardware platforms. The Cisco Nexus 1000V may be used with generic NICs on generic x86-based servers. In addition, the upstream physical access layer switch may also be generic.

This generic support enables the Cisco Nexus 1000V to be installed and configured within existing architectures, minimizing disruption and maximizing functionality.

VN-Link Packet Flow—Cisco Nexus 1000V and a Generic Adapter

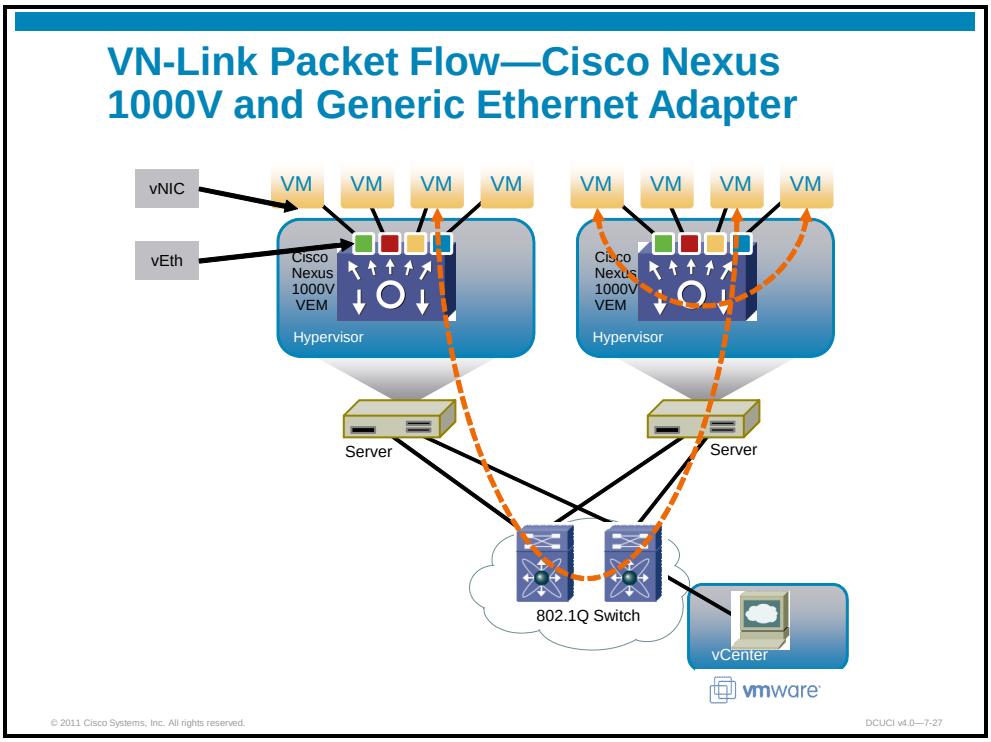

The Cisco Nexus 1000V is similar to physical Ethernet switches. For packet forwarding, the Cisco Nexus 1000V uses the same techniques that other Ethernet switches apply, keeping a MAC address-to-port mapping table that is used to determine where packets should be forwarded. The Cisco Nexus 1000V maintains forwarding tables in a slightly different manner than other modular switches. Unlike physical switches with a centralized forwarding engine, each VEM maintains a separate forwarding table. No synchronization exists between forwarding tables on different VEMs. In addition, there is no concept of forwarding from a port on one VEM to a port on another VEM. Packets that are destined for a device that is not local to a VEM are forwarded to the external network, which in turn may forward the packets to a different VEM.

VN-Link Products—Cisco UCS 6100 and VIC

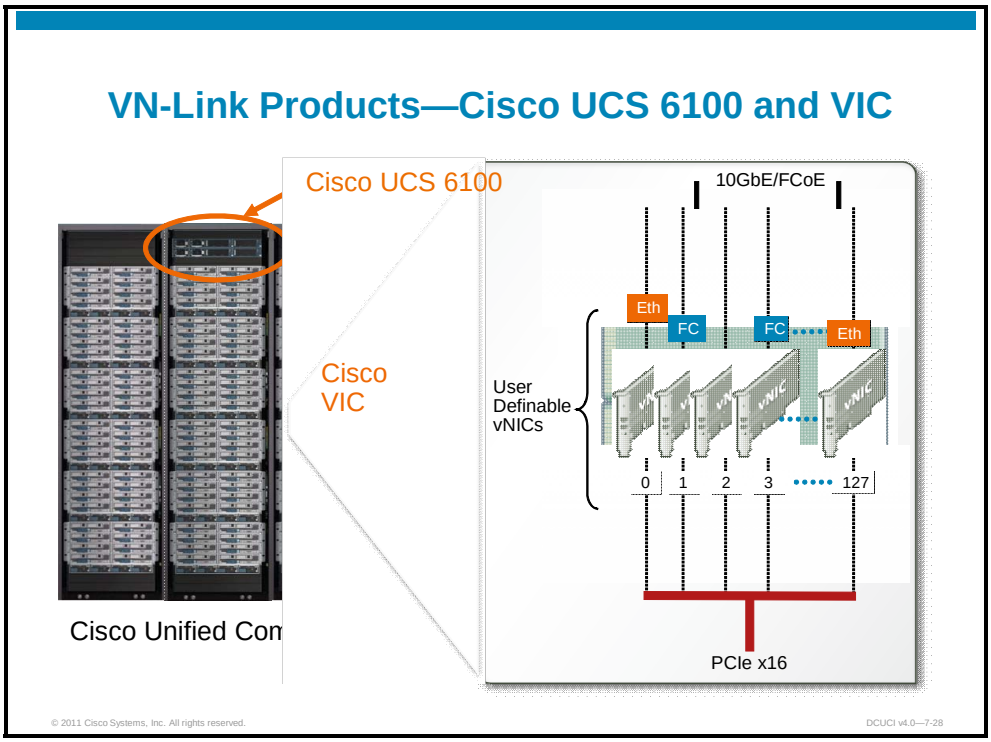

The Cisco M81KR/P81E is a VIC that operates within the Cisco Unified Computing System.

The VIC adapter is a hardware-based VN-Link solution, which uses a virtual network tag (VNTag) to deliver VM visibility at the physical access switch. The VIC adapter assigns a VNTag to each vNIC that is created as part of the service profile within Cisco UCS Manager. The assigned VNTag is locally significant and has visibility between the VIC and Cisco UCS 6100 Fabric Interconnects.

The VNTag provides traffic separation between VMs that share the same VIC adapter and is a reliable method for virtualizing each VM vNIC onto the physical access layer switchport.

VN-Link Deployment—VIC and Cisco UCS 6100 Series with VMware VMDirectPath

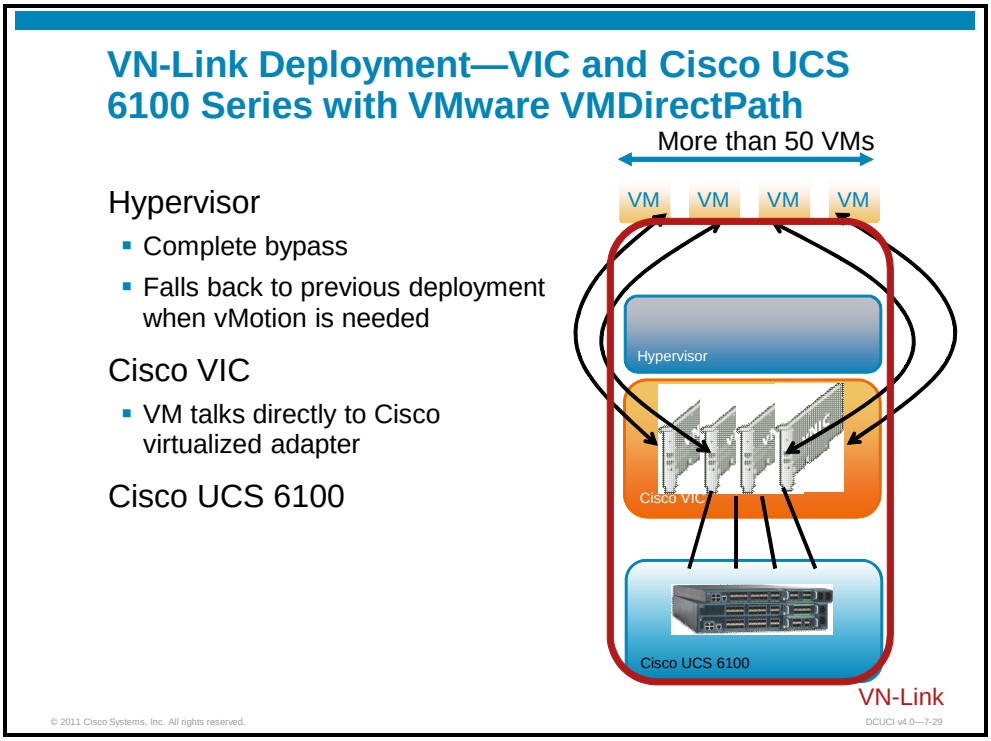

The VIC adapter offers the Pass-Through Switch (PTS) feature. When configured, this feature enables I/O processing to be offloaded from and bypass the local hypervisor to the hardware VIC. Therefore, the VM vNIC communicates directly with the VIC adapter. The M81KR VIC mezzanine adapter option works within the Cisco UCS B-Series platform, and the P81E Peripheral Component Interconnect (PCI) adapter works within the Cisco UCS C-Series platform.

VN-Link Packet Flow—VIC Cisco UCS 6100 Series

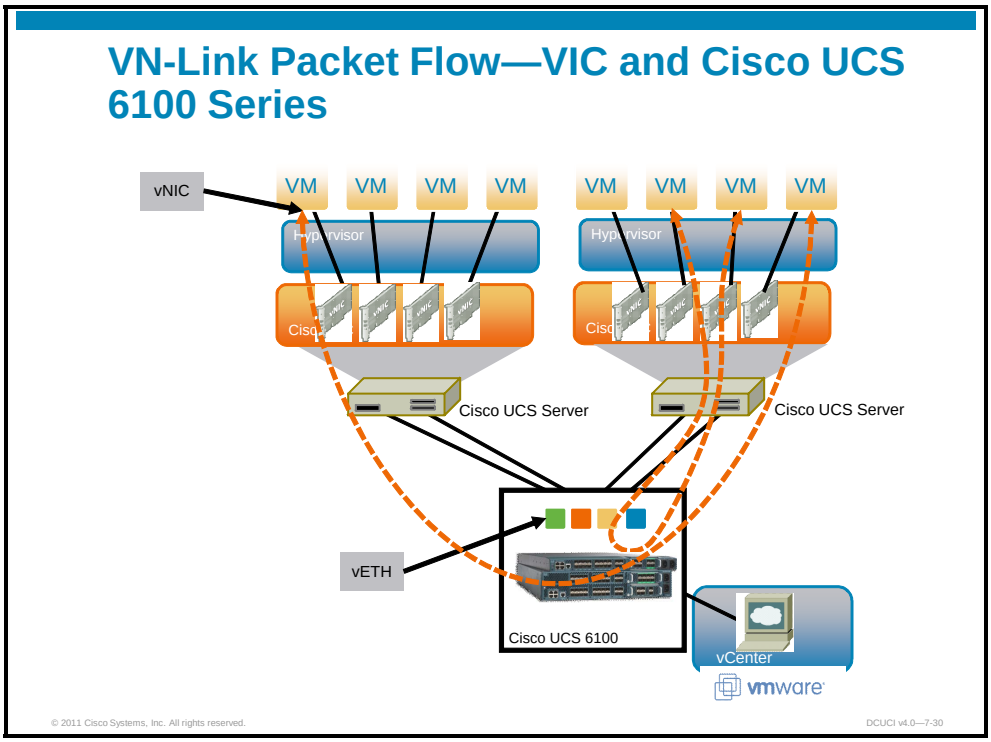

Packet flow within a Cisco Unified Computing System cluster that includes the VIC begins with the operating system, constructing frames in a traditional fashion. The VIC adds the VNTag, which consists of a new, 6-byte file that is inserted directly behind the source MAC address within the Ethernet frame format. The VNTag also has a new EtherType that is assigned for this service.

The virtual interface switch provides both ingress and egress processing, and the VIC forwarding is based upon the VNTag value. VNTag assignment and removal remain local to the Cisco Unified Computing System cluster. Finally, the operating system stack receives the frame transparently, as though the frame is directly connected to the physical access layer switch.

Summary of All VN-Link Offerings

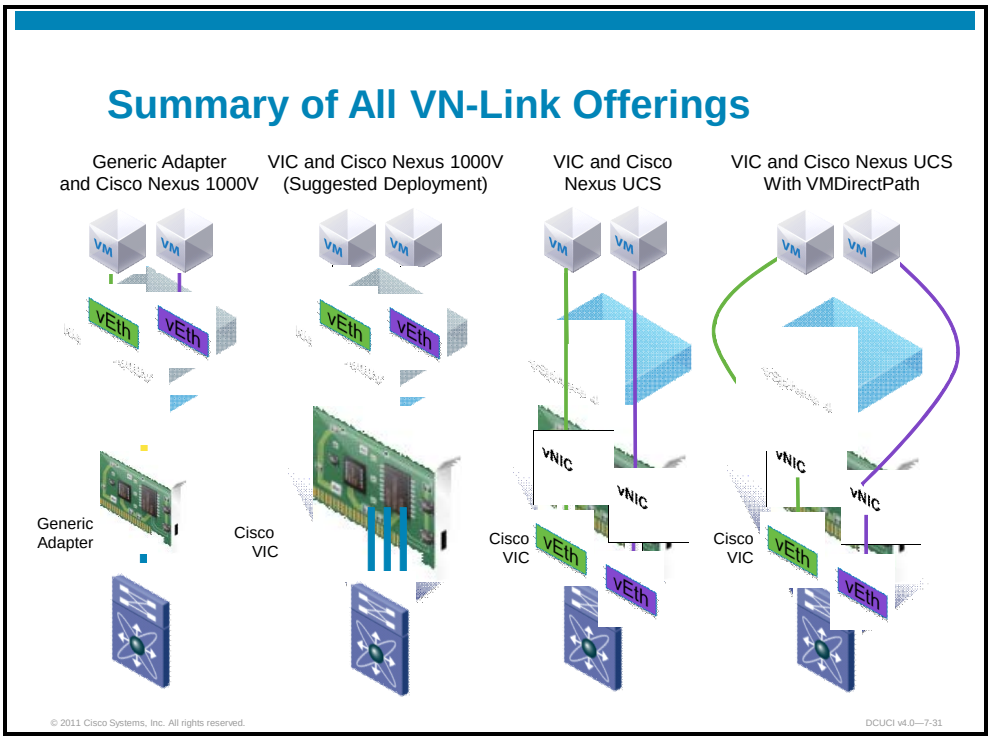

The main solution that VN-Link offers is the concept of a VM adapter with a virtual patch cord that is connected to a physical access layer switch. This solution essentially translates to providing VM vNIC adapter visibility and virtualizing this interface onto the physical switch port.

This solution may be implemented by using the Cisco Nexus 1000V with a generic Ethernet adapter, in which the VEM assigns the vEthernet interface, which survives vMotion events. The Cisco Nexus 1000V may also be combined with the VIC adapter within a Cisco Unified Computing System cluster. In this configuration example, the VEM would assign the vEthernet interface to deliver the VN-Link implementation.

Within a Cisco Unified Computing System cluster, the VIC adapter offers the hardware-based implementation of VN-Link, by using VNTag assignment on the VIC. This option offers two unique choices that are based on performance: the PTS feature or the use of VMDirectPath.

Why All the Different Models?

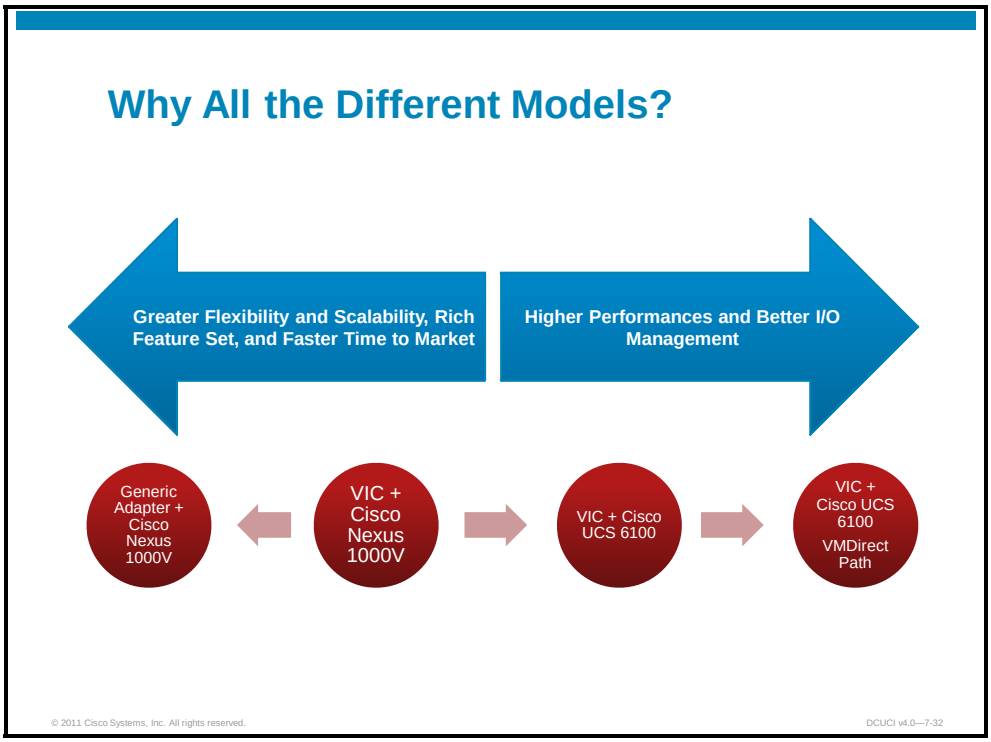

Cisco has created a flexible set of networking options to support various existing and new virtual environments. For large, diverse environments that require advanced networking feature sets and flexible deployments, the use of the Cisco Nexus 1000V and Cisco Nexus 1010, combined with either generic or VIC adapters, is a solid choice.

For environments that may require higher performance and more predictable I/O management, the use of a Cisco UCS B- or C-Series server and the VIC adapter, with or without the configured PTS feature, is another deployment option.

The Cisco VN-Link solution caters to a wide continuum of networking requirements within virtualized server environments.

Summary

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

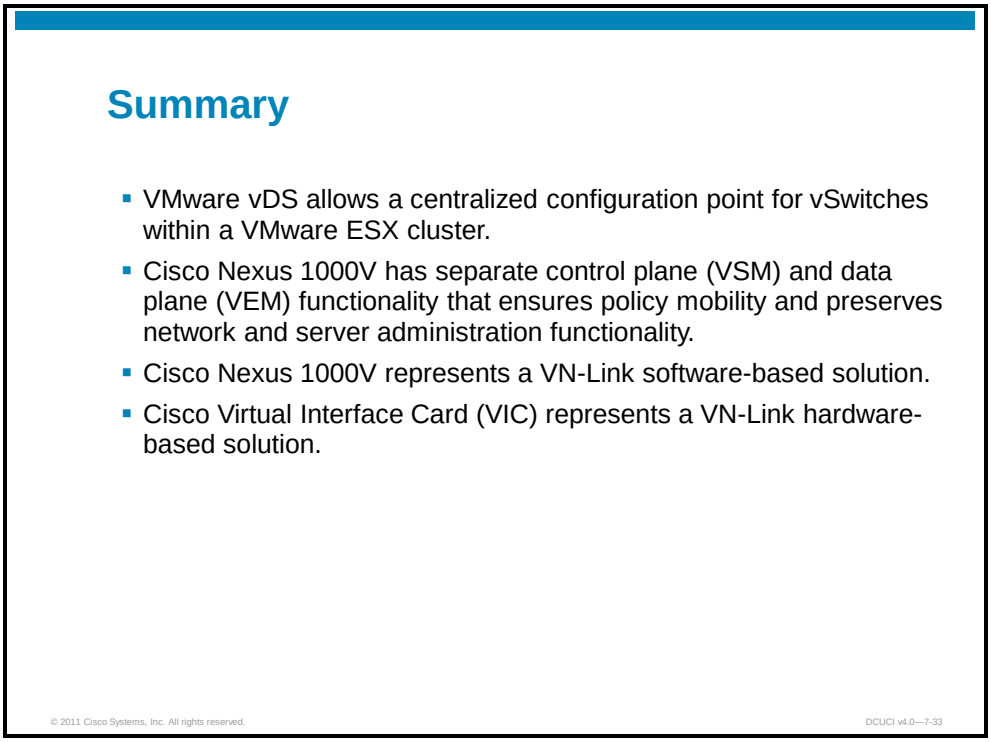

Lesson 3

Characterizing Cisco Nexus 1000V Architecture

Overview

Implementers must be able to identify component features of the Cisco Nexus 1000V architecture and articulate the architectural solution at a high level. This lesson presents the architecture of the Cisco Nexus 1000V, as well as its capabilities for delivering a software-based Cisco Virtual Network Link (VN-Link) solution.

Objectives

Upon completing this lesson, you will be able to characterize the architecture of the Cisco Nexus 1000V and its capabilities in delivering VN-Link. This ability includes being able to meet these objectives:

- Describe the Cisco Nexus 1000V solution
- Describe the Cisco Nexus 1000V architecture

Cisco Nexus 1000V Overview

This topic provides an overview of Cisco Nexus 1000V.

Cisco Nexus 1000V Series DVS

Cisco Nexus 1000V Series DVS

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

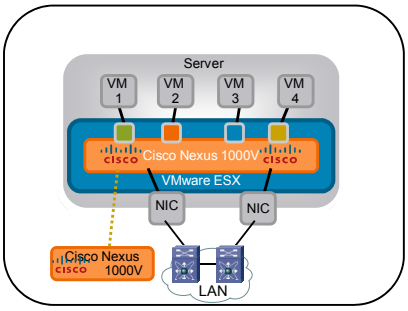

The Cisco Nexus 1000V Series Switch is a software solution that is used in place of the VMware Standard Switch (vSwitch) to provide advanced functionality. The Cisco Nexus 1000V allows all physical network best practices to be applied at the virtual machine (VM) level. Achieving this level of service with the VMware vSwitch is not possible.

Cisco Nexus 1000V

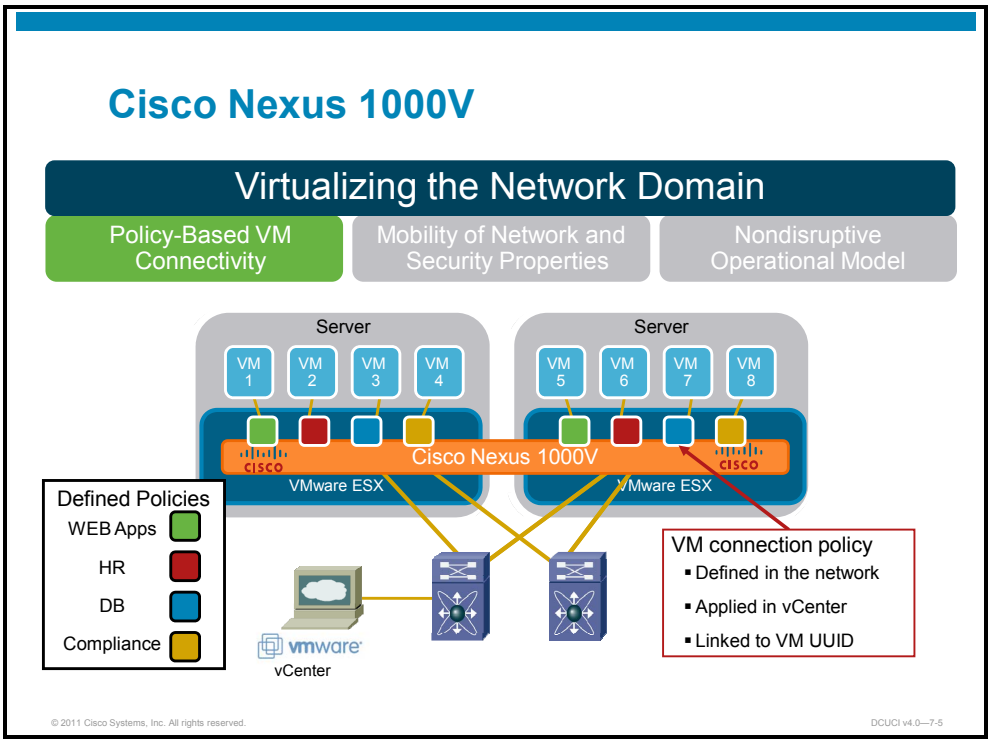

The distributed architecture of the Cisco Nexus 1000V deployment allows a single Cisco Nexus Operating System (Cisco NX-OS) software-based supervisor module to manage the switching capabilities of as many as 64 VMware ESX servers.

Each 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 can run as a VM that runs on a VMware ESX or ESXi host, or as a physical appliance that is known as a VMware vCenter virtual appliance.

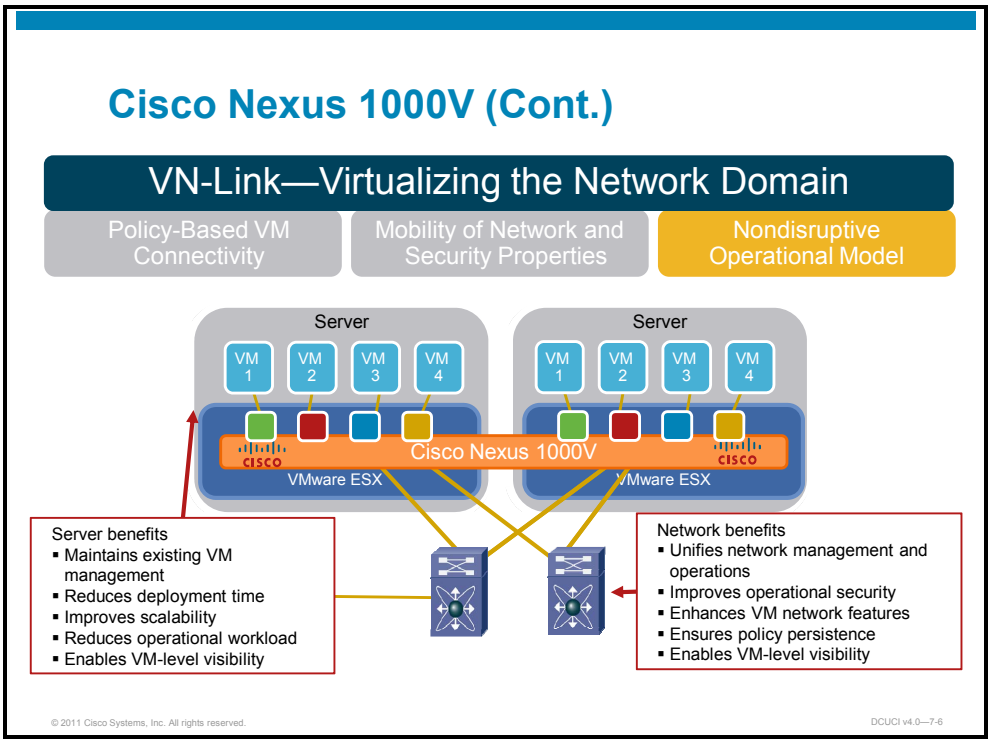

A new type of port has been created to identify the individual VMs. This port type is a virtual Ethernet (vEthernet) port. This port type represents a virtual network interface card (vNIC), regardless of which physical server it is on. vEthernet ports are not tied to a specific server or physical VM NIC (VMNIC), but instead represent the vNIC of a virtual server. The vEthernet port for a VM will remain the same even if the VM migrates. The vEthernet port remains the same, so if a VM obtains vEthernet 1 on host 1 and is moved to host 2, then the VM retains vEthernet 1 as its port. Keeping the same vEthernet port allows for network configuration and policy mobility.

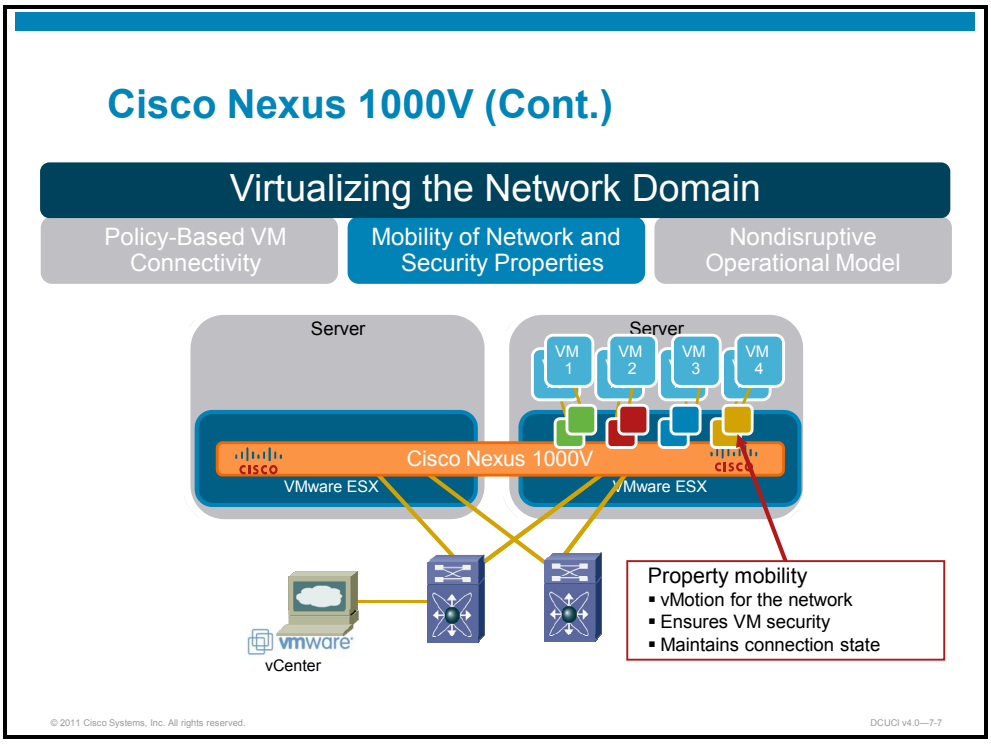

This figure demonstrates the policy mobility that the Cisco Nexus 1000V Distributed Virtual Switch (DVS) offers. A VMware vMotion event (either manual or automatic) has occurred, causing the VMs on the first ESX server to be moved to the second ESX server, which is now running all eight VMs.

The policies have been assigned to the VM vEthernet interfaces. Because these interfaces remain with the VM, then by inheritance, the policies that are assigned to the vEthernet interfaces remain with the VM.

Managing Network Policies with the Cisco Nexus 1000V

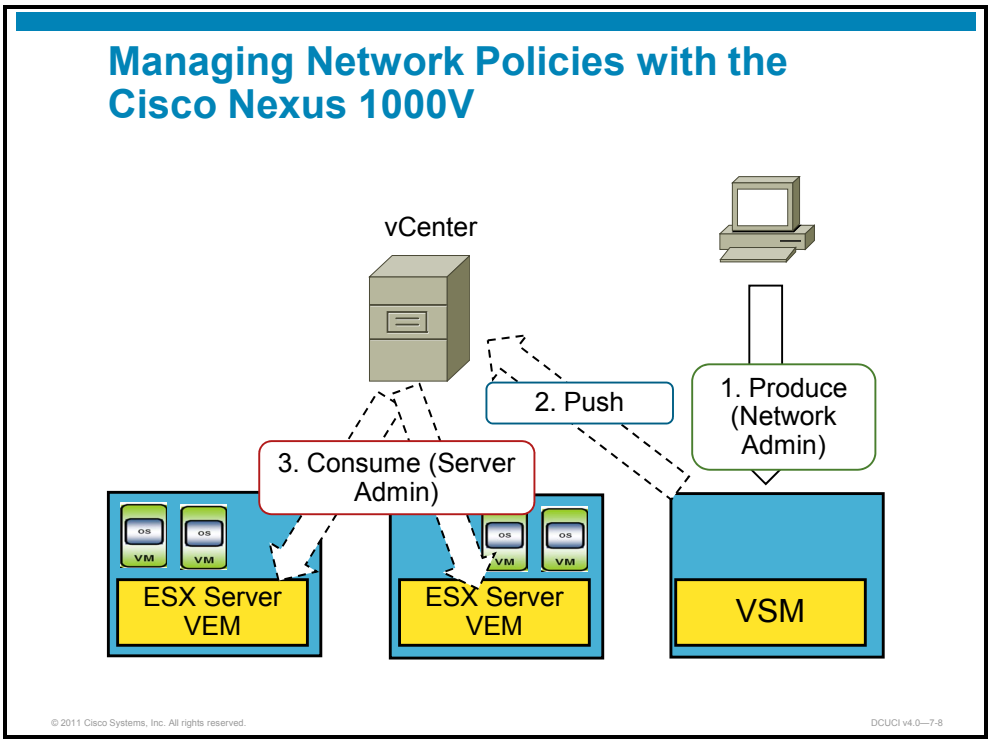

The Cisco Nexus 1000V provides an ideal model in which network administrators can define a network policy that virtualization or server administrators can use as new VMs are created. Policies that are defined on the Cisco Nexus 1000V are exported to vCenter. The server administrator then assigns specific network policies as new VMs require access. This concept is implemented on the Cisco Nexus 1000V by using a feature called port profiles. With the port profile feature, the Cisco Nexus 1000V eliminates the requirement for the server administrator to create or maintain vSwitch and port group configurations on ESX hosts.

Port profiles separate network and server administration. For network administrators, the Cisco Nexus 1000V feature set and the capability to define a port profile by using the same syntax that is used for existing physical Cisco switches helps ensure consistent policy enforcement, without the burden of managing individual switch ports. The Cisco Nexus 1000V solution also provides a consistent network management, diagnostic, and troubleshooting interface to the network operations team, allowing the virtual network infrastructure to be managed like the physical infrastructure.

Cisco Nexus 1000V Architectural Overview

This topic provides an overview of Cisco Nexus 1000V architecture.

Cisco Nexus 1000V Architecture

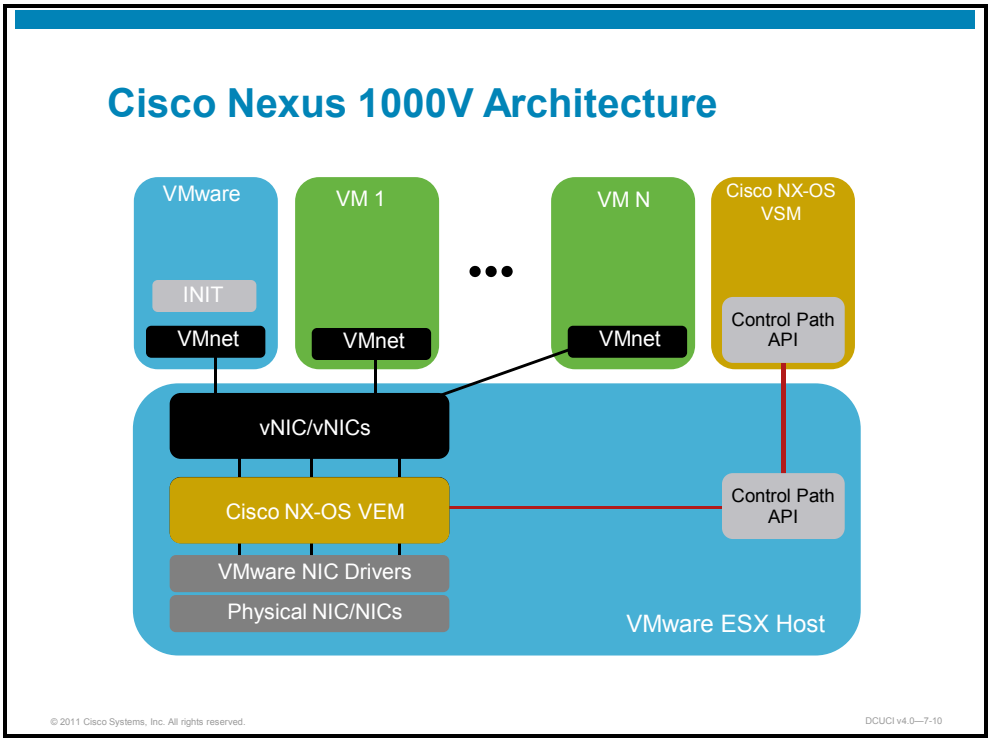

The figure depicts the basic architecture of the Cisco Nexus 1000V. The Cisco Nexus 1000V uses a control path application programming interface (API) to communicate with the data plane. This API simulates out-of-band (OOB) management and does not interfere with the data plane. In the figure, the Cisco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

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

Note	A NIC in VMware is represented by the VMNIC interface. The VMNIC number is allocated during VMware installation.

Cisco Nexus 1000V VLANs

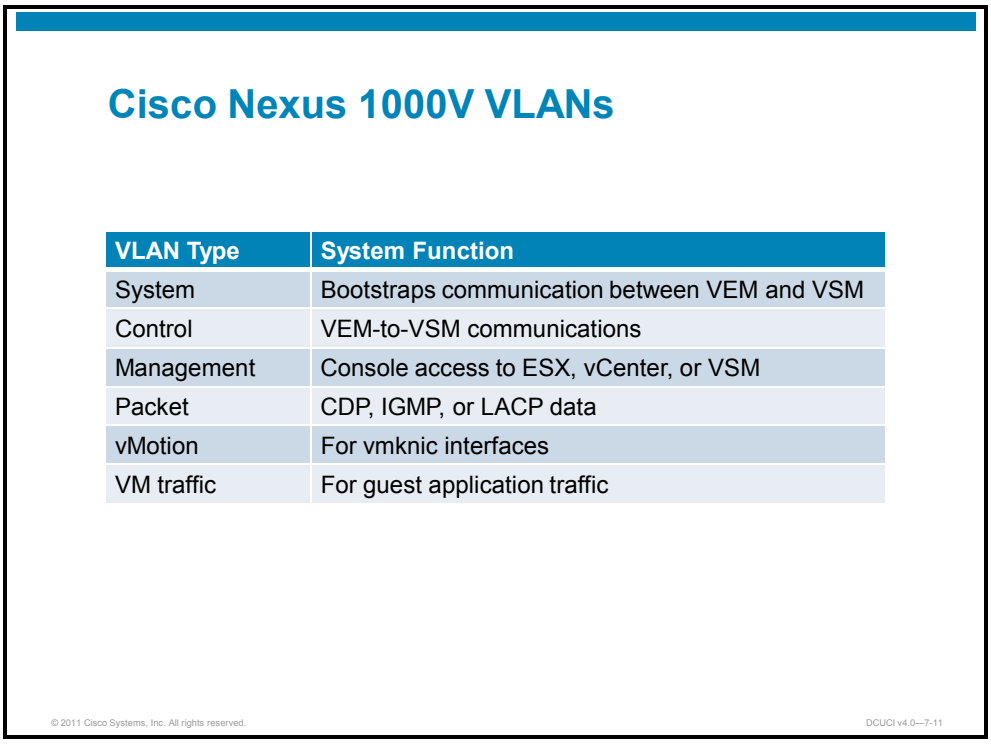

Classification of traffic types in any network is difficult to achieve. In a VMware environment, traffic varies based on the types of applications that are virtualized. However, some traffic types can be identified and general prioritization can be applied. The figure shows the general classifications of traffic for a typical VMware deployment.

Cisco Nexus 1000v Management VLAN

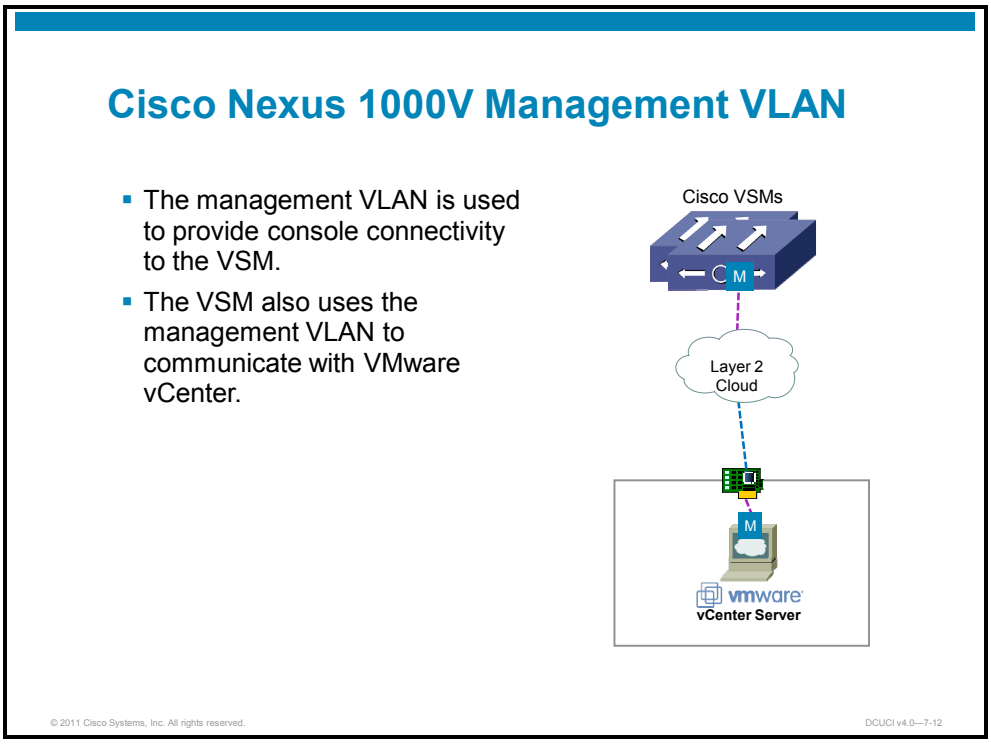

The management VLAN is one of three mandatory VLANs for Cisco Nexus 1000V communications services. The management VLAN has two primary purposes:

- Providing access to the VSM command-line interface (CLI) console via Secure Shell (SSH) or Telnet
- Communicating with vCenter

Cisco Nexus 1000V Control and Packet VLANs

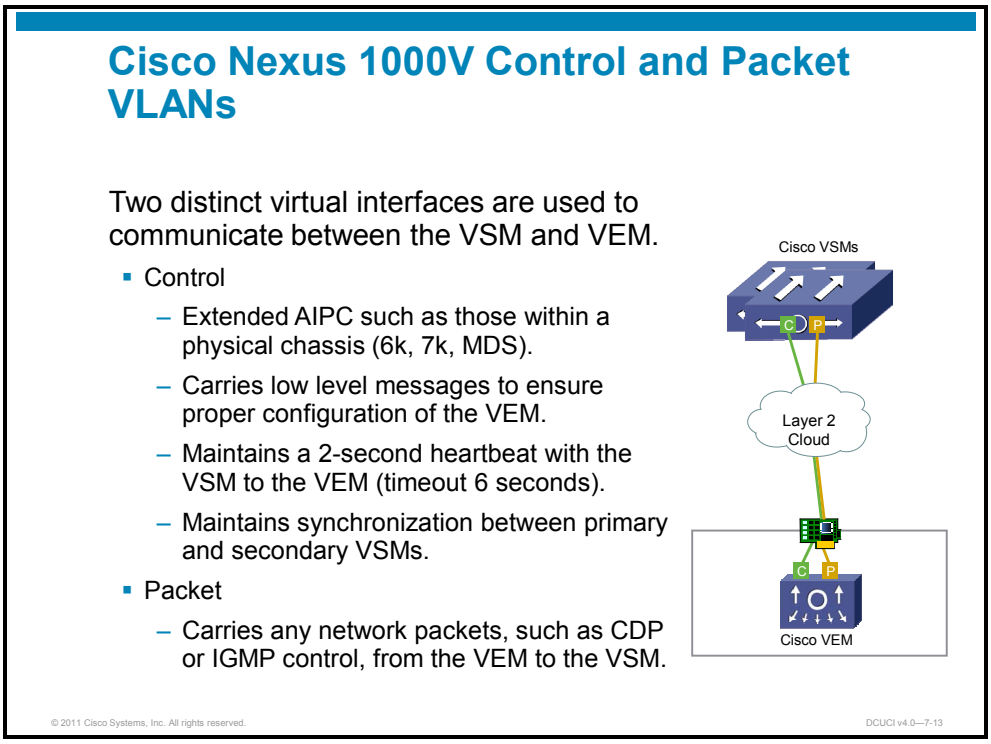

In a traditional modular switch, the supervisor has physical connections to its line cards. In the Cisco Nexus 1000V, the VSM uses IP packets to communicate with the VEMs as remote line cards. Two separate channels are required:

- Control VLAN
  - Advanced Interprocess Communications (AIPC)
  - Communication sync between primary and secondary VSMs
  - Low-level configuration messaging from VSM to VEM
  - Heartbeat from VSM to VEM
- Packet VLAN
  - Allowance for Cisco Discovery Protocol packets to be exchanged from the VSM and VEM
  - Allowance for port channel setup on VEM ports, via Link Aggregation Control Protocol (LACP)
  - Allowance for Internet Group Management Protocol (IGMP) join and leave messaging for multicast communications

Cisco Nexus 1000V Configuration Example

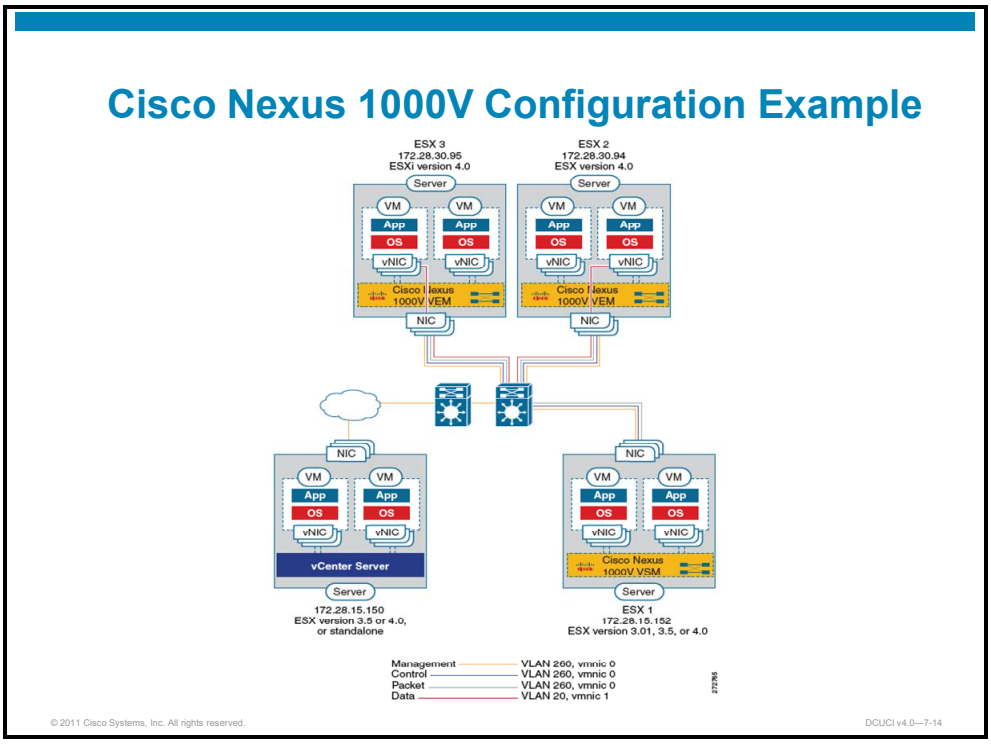

In this configuration example, three ESX servers are managed by the vCenter server that is shown in the lower left corner of the figure. A standalone VSM has been installed and configured on host ESX 1. Hosts ESX 2 and ESX 3 have VEMs that are installed to provide the data traffic forwarding for the host.

Four VLANs have been configured to support VSM-to-VEM, VSM-to-vCenter, and VSM-to-management network connectivity. VLAN 260 is used to support control, packet, and management functions, and VLAN 20 is used for data traffic forwarding.

VEM-to-VSM Communication

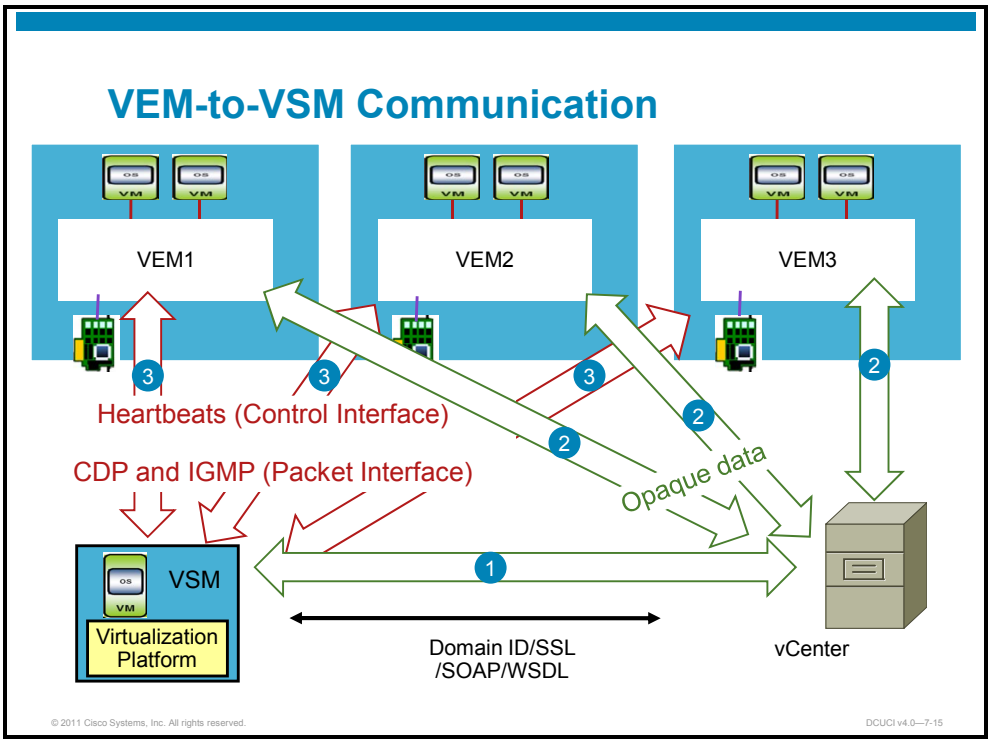

Like the VSM, each VEM has control and packet interfaces. The end user cannot manage or directly configure these interfaces.

1. The VEM uses the opaque data that vCenter provides to configure the control and packet interfaces with the correct VLANs. The VEM then applies the correct uplink port profile to the control and packet interfaces, to establish communication with the VSM. After the VSM recognizes the VEM, a new module is virtually inserted into the Cisco Nexus 1000V virtual chassis. The VSM CLI notifies the network administrator that a new module has powered on, much as happens with a physical chassis. The module assignment is sequential, meaning that the VEM is assigned the lowest available module number between 3 and 66. When a VEM comes online for the first time, the VSM assigns the module number and tracks that module by using the Universally Unique Identifier (UUID) of the ESX server. This process helps ensure that if the ESX host loses connectivity or is powered down for any reason, the VEM will retain its module number when the host comes back online.

2. Cisco Nexus 1000V implements a solution called domain IDs. A domain ID is a parameter of the Cisco Nexus 1000V and is used to identify a VSM and VEM that relate to each other. The domain ID of the Cisco Nexus 1000V is defined when the VSM is first installed and becomes part of the opaque data that is transmitted to vCenter. Each command the VSM sends to any associated VEMs is tagged with this domain ID. When a VSM and VEM share a domain ID, the VEM accepts and respond to requests and commands from the VSM. If the VEM receives a command or configuration request that is not tagged with the proper domain ID, then the VEM ignores that request. Similarly, if the VSM receives a packet that is tagged with the wrong domain ID from a VEM, the VSM ignores that packet.

3. The VSM maintains a heartbeat with its associated VEMs. This heartbeat is transmitted at 2-second intervals. If the VSM does not receive a response within 8 seconds, the VSM considers the VEM to be removed from the virtual chassis. If the VEM is not responding because of a connectivity problem, then the VEM continues to switch packets in its last known good state. When communication is restored between a running VEM and the VSM, the VEM is reprogrammed, causing a slight (1 to 15 second) pause in network traffic. All communication between the VSM and the VEM is encrypted with a 128-bit algorithm.

Cisco Nexus 1000V Opaque Data

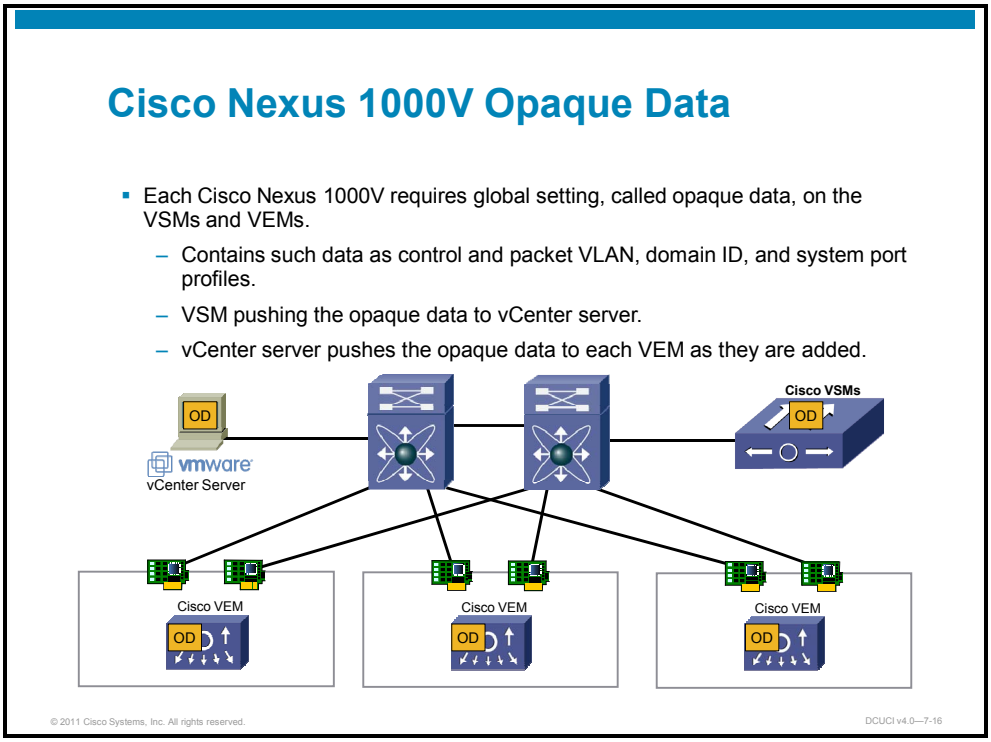

The term *opaque data* is applied to information that is encrypted with a Master Key that the VSM and vCenter hold. The vCenter server imports the VSM digital certificate, to establish a confidential channel to the VSM. As each VEM is enabled, vCenter exchanges a protected element that includes configuration information that is authenticated by the encrypted domain ID. This process ensures that only trusted vCenter servers push port profiles to the VEMs and that vCenter allows communications only to trusted VSMs.

Policy-Based VM Connectivity

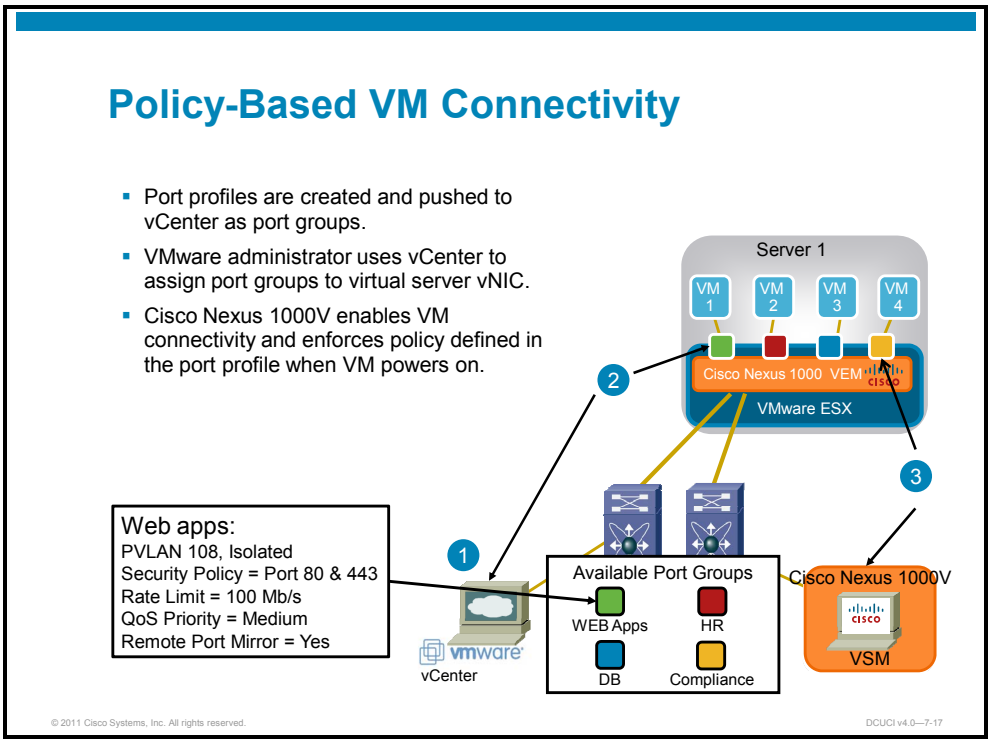

Network administrators configure VM port profiles that contain the policies for specific server types. These port profiles are passed to vCenter as port groups and are assigned to individual vNICs by the VMware administrator.

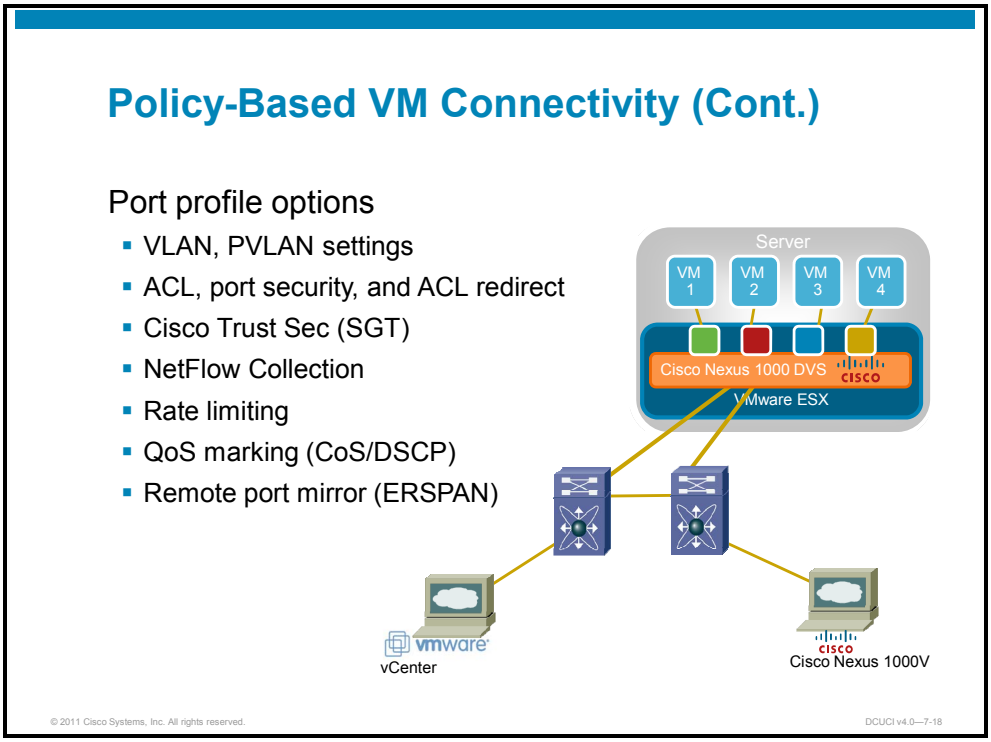

The figure summarizes the various security and traffic flow options that can be applied to port profiles:

- VLAN and private VLAN (PVLAN) settings
- Access control list (ACL), port security, and ACL redirect
- Cisco TrustSec security group tag (SGT)
- NetFlow collection
- Rate limiting
- Quality of service (QoS) marking (class of service [CoS]/ differentiated services code point [DSCP])
- Remote port mirror (Encapsulated Remote Switched Port Analyzer [ERSPAN])

Mobility of Security and Network Properties

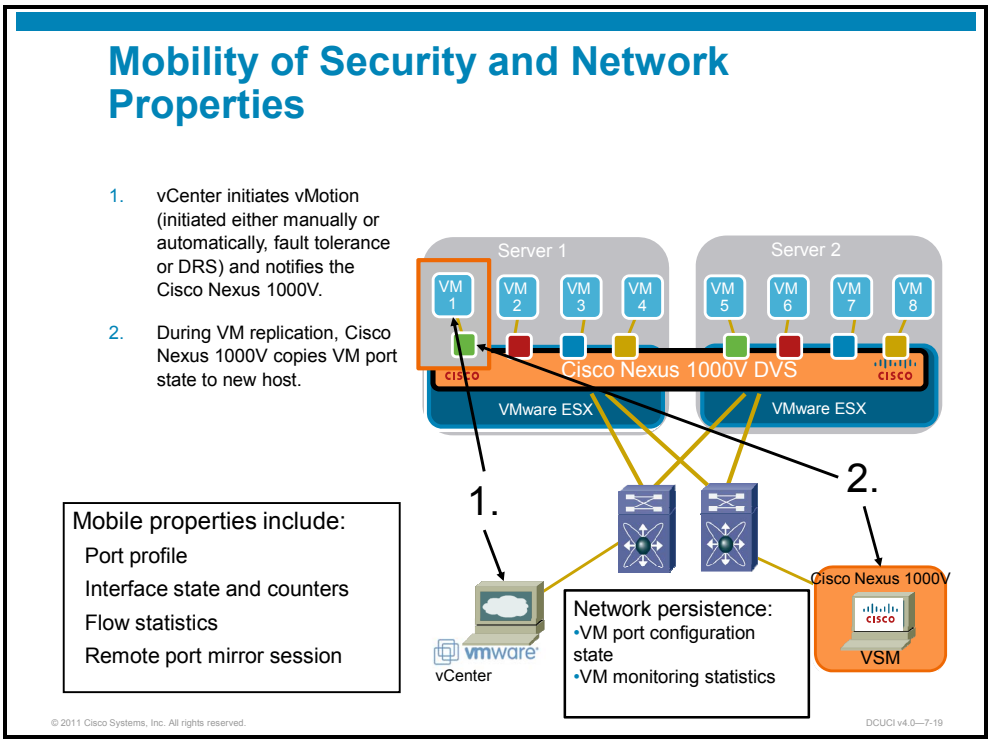

As VMs migrate between hosts through manual or automatic processes, vCenter and the Cisco Nexus 1000V VSM work together to maintain the state and port policies of the VMs. This cooperation helps to ensure network security and consistency in a rapidly adapting VMware environment.

Mobility of Security and Network Properties (Cont.)

- When vMotion completes, the port on the new VMware ESX host is brought up and the VM MAC address is announced to the network.

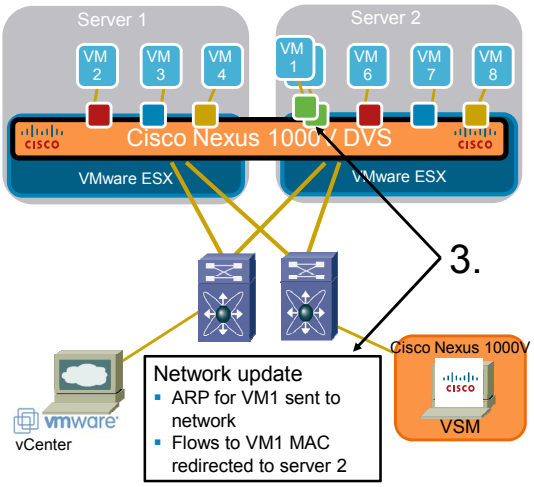

This function offers several benefits:

- Prevents disparate vSwitch configurations from affecting Distributed Resource Scheduler (DRS) or live VM migration
- Enables security setting, network policy, and connectivity state to move with vMotion
- Offers continuous traffic mirroring as VMs move between servers

Summary

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

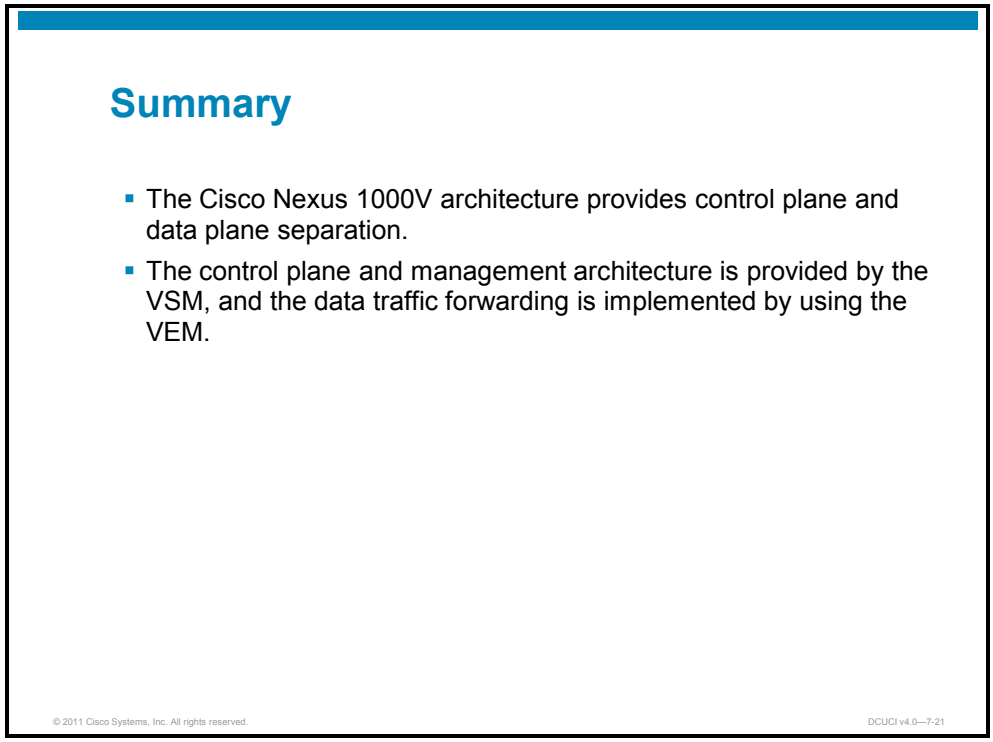

Lesson 4

Installing and Configuring the Cisco Nexus 1000V Switch

Overview

The installation of Cisco Nexus 1000V components can be a challenging aspect to delivering a Cisco Unified Computing System (UCS) solution. It is very important to understand the installation methods that are available for the Cisco Nexus 1000V Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM).

This lesson describes the process to install the Cisco Nexus 1000V VSM, provide connectivity to VMware vCenter, and perform initial configuration of the Cisco Nexus 1000V Distributed Virtual Switch. In addition, the Cisco Nexus 1010 Virtual Services Appliance is introduced and compared to the Cisco Nexus 1000V.

Objectives

Upon completing this lesson, you will be able to describe the installation, connection, and configuration of the Cisco Nexus 1000V VSM and compare the Cisco Nexus 1000V and the Cisco Nexus 1010 Virtual Services Appliance. This ability includes being able to meet these objectives:

- Describe how to configure the preinstallation vSwitch network
- Describe how to install a VSM on a VM
- Describe the initial configuration of the VSM
- Describe the configuration of the certificate exchange and connection from the VSM to vCenter
- Describe the Cisco Nexus 1000V high-availability configuration
- Describe the Cisco Nexus 1010 Virtual Services Appliance
- Differentiate between the capabilities of the Cisco Nexus 1010 Virtual Services Appliance and the Cisco Nexus 1000V software VSM

Configure VSM vSwitch Networks

This topic discusses how to configure VSM VMware vNetwork Standard Switch (vSwitch) networks.

Preparing the ESX Servers

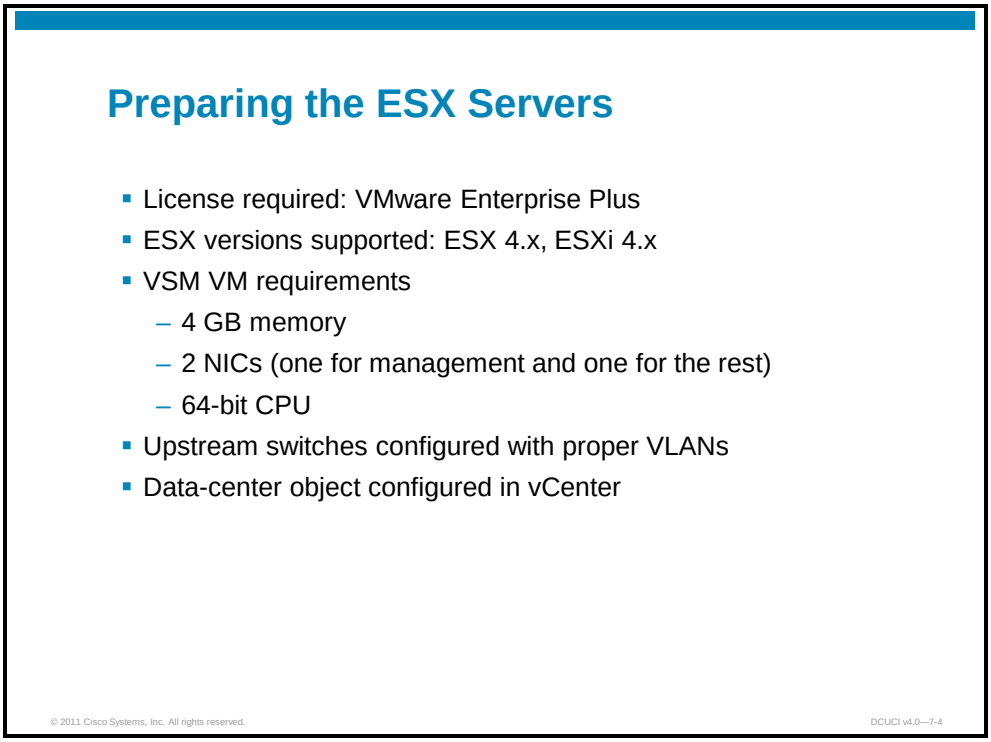

You must install the VMware Enterprise Plus license on each VMware ESX server, to run a VEM on that host.

You can run the VSM in a virtual machine (VM) that is running either VMware ESX 3.5 or 4.0 or ESXi 4.0. You must run ESX or ESXi 4.0 to support the VEMs.

The VSM VMs require at least 4 GB of memory, two physical network interface cards (NICs) on the ESX host (one NIC for management and one for the rest of the traffic), and a 64-bit processor.

The upstream switches must be configured to trunk the VLANs that you intend to use for control, management, packet, VMware vMotion, and VM traffic networks.

Last, you must configure a data-center object to contain your ESX servers in vCenter.

VSM Port Group Requirements

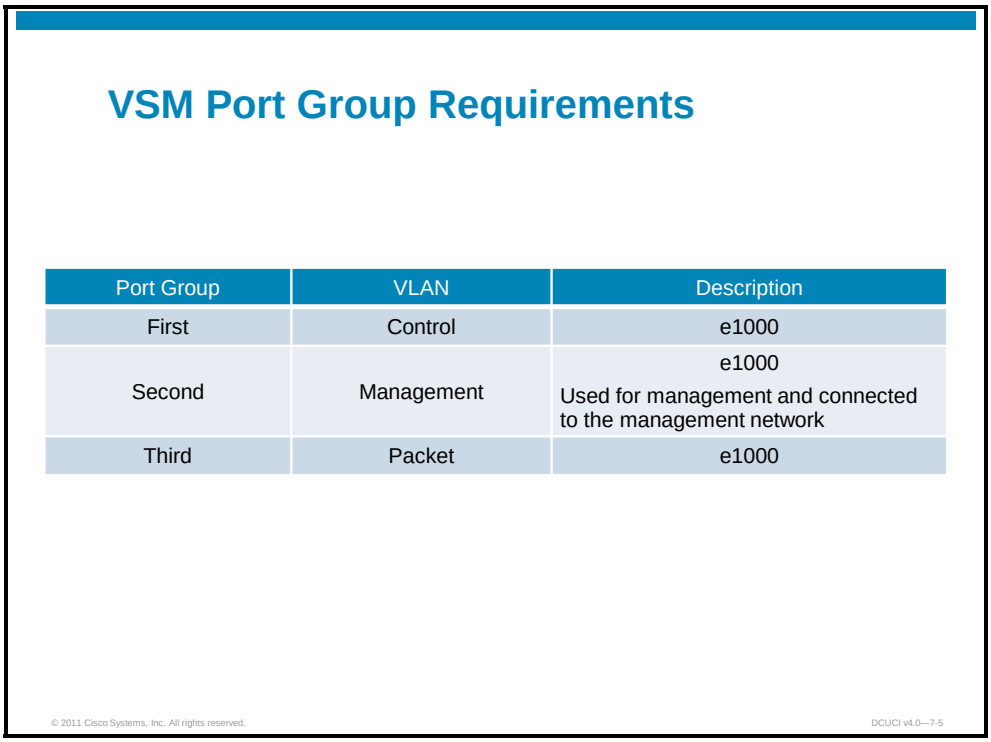

The Cisco Nexus 1000V VSM requires at least three types of port groups to be created on an existing or new vSwitch. These port groups correspond to the VSM communication VLANs, as described previously.

VSM Port Group Creation

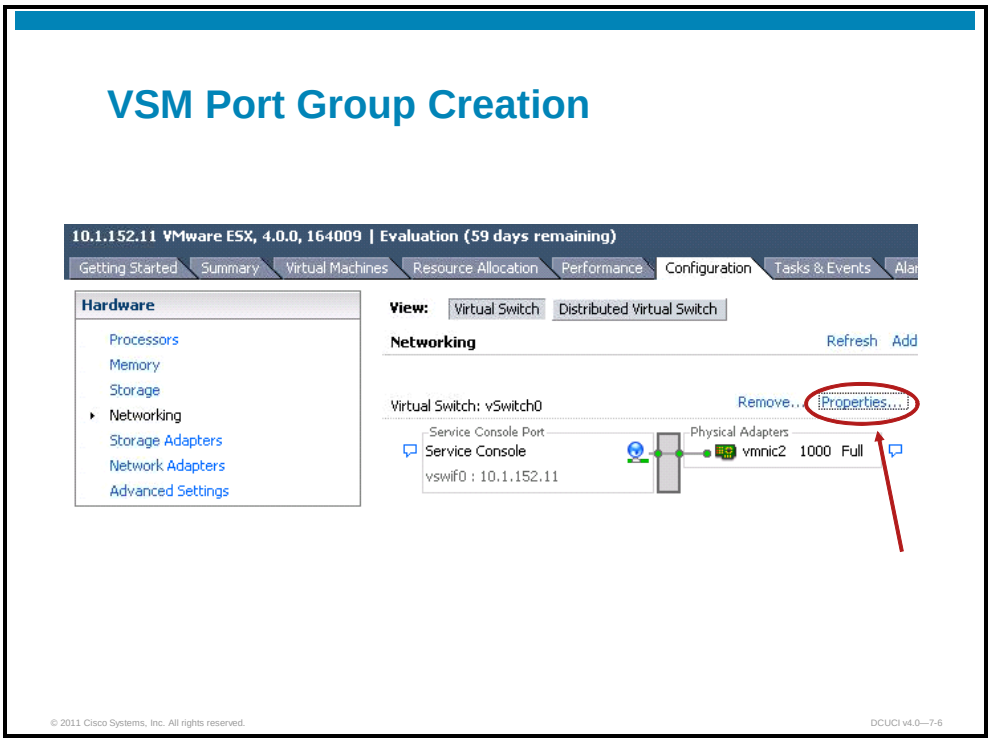

To create a VMware port group that the VSM can use, follow these steps:

1. Choose the host on which the VSM will reside.
2. Choose the **Configuration** tab.
3. Choose **Networking**.
4. Choose **Properties**.

VSM Port Group Creation (Cont.)

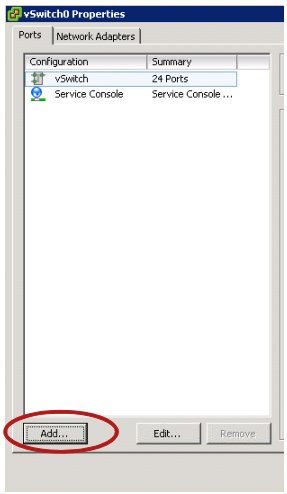

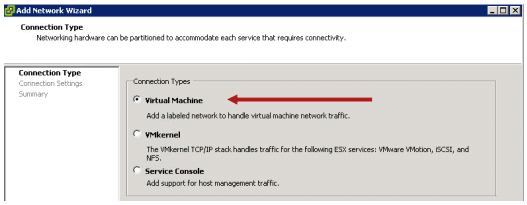

5. In the Properties pane, choose **Add**.
6. Choose **Virtual Machine**.
7. Click **Next**.

VSM Port Group Creation (Cont.)

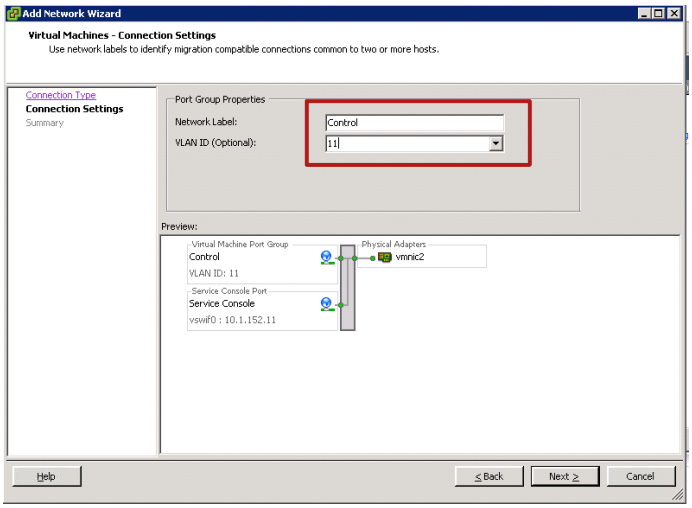

After you choose Add, configure the following parameters:

8. For the Network Label, type **Control**.
9. For the VLAN ID, provide the VLAN number that was chosen for the VSM control traffic.
10. Click **Next**.

VSM Port Group Creation (Cont.)

- Repeat this procedure for the management and packet port groups.

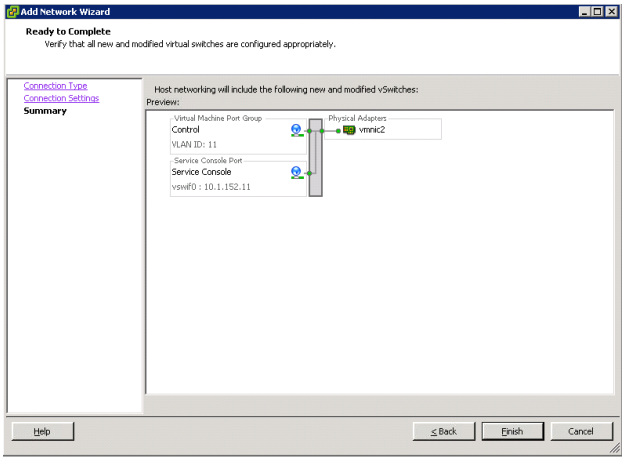

After you choose Next, verify the configuration of the control VLAN for the VSM. Click **Finish**. Repeat for the management and packet port groups.

VSM vSwitch Configuration Showing Port Groups

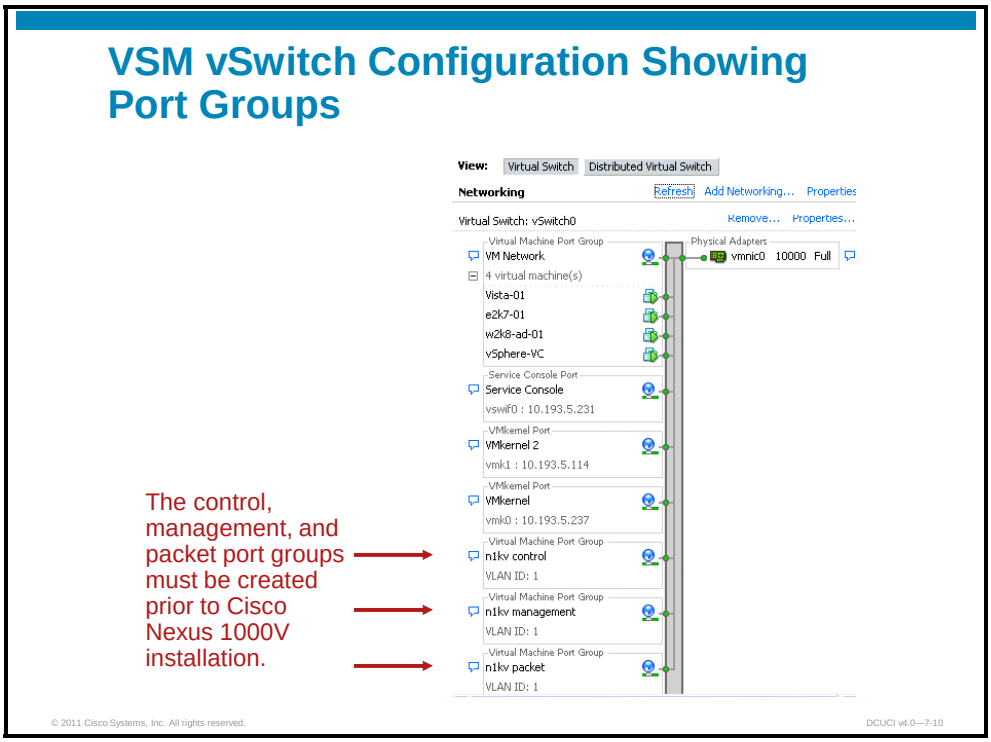

In the traditional Cisco Nexus 1000V installation, you create different port groups and VLANs for the control, management, and packet interfaces, as shown in the figure. These port groups must exist before you install the Cisco Nexus1000V.

Install the VSM on a VM

This topic discusses how to install the VSM on a VM.

Cisco Nexus 1000V VSM Installation Methods

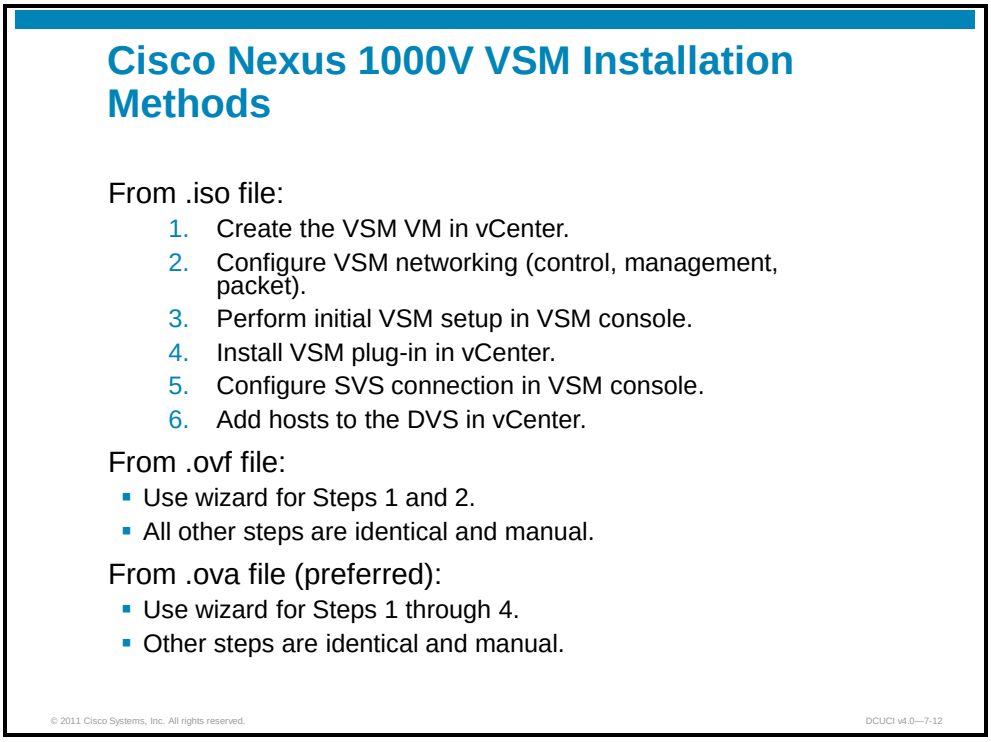

There are several ways to install and deploy the VSM. The preferred method is to use an Open Virtual Appliance (OVA) file. This method provides the highest degree of guidance and error-checking for the user.

Open Virtualization Format (OVF) files are standardized file structures that are used to deploy VMs. You can create and manage OVF files by using the VMware OVF Tool.

OVA files are like OVF files.

Creating a VSM VM—Choose VM Configuration

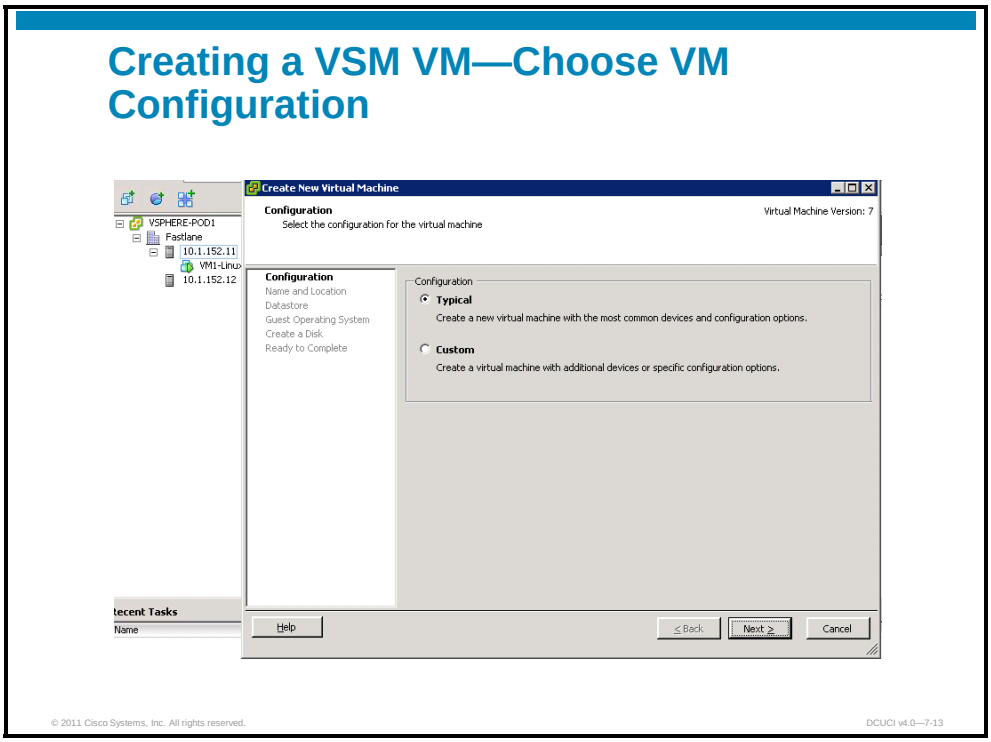

The VSM VM requirements include these specifics:

- **Type:** Other Linux (64 bit)
  - 3 GB disk
  - 2 GB memory
  - 3 NICs, type E1000

To create the VM, right-click the host on which you want to install the VSM, then choose **New Virtual Machine**.

Choose **Typical**, then click **Next**.

Creating a VM—Name the VM and Inventory Location

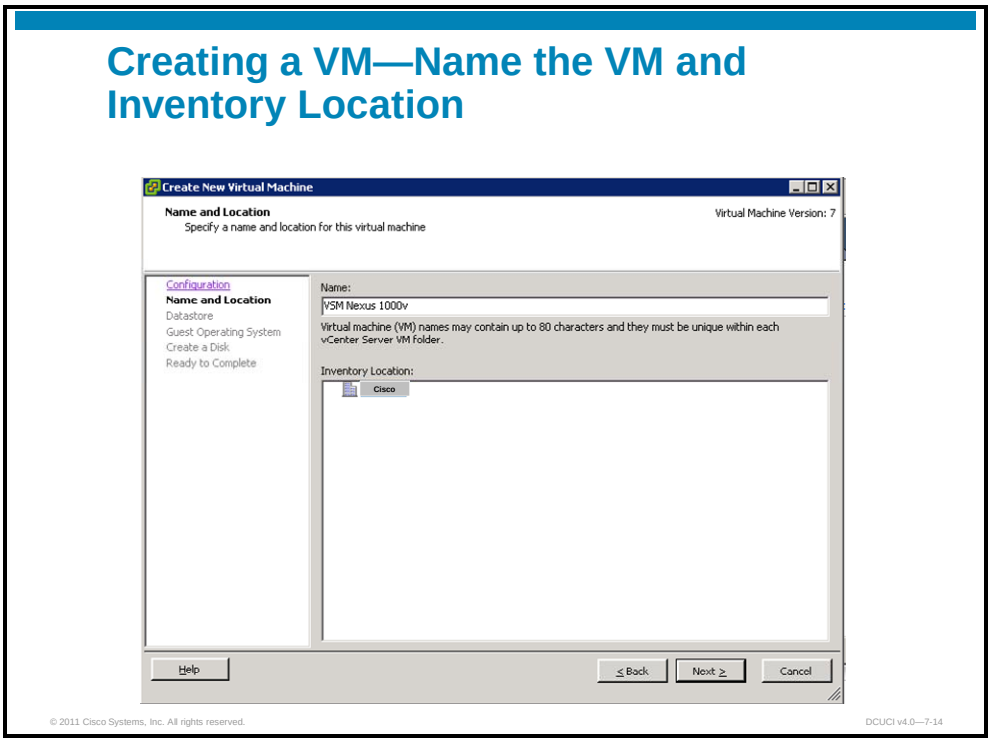

In the dialog box, name the VM and choose a data center. Click **Next**.

Creating a VM—Choose Shared Storage for VSM VM Files

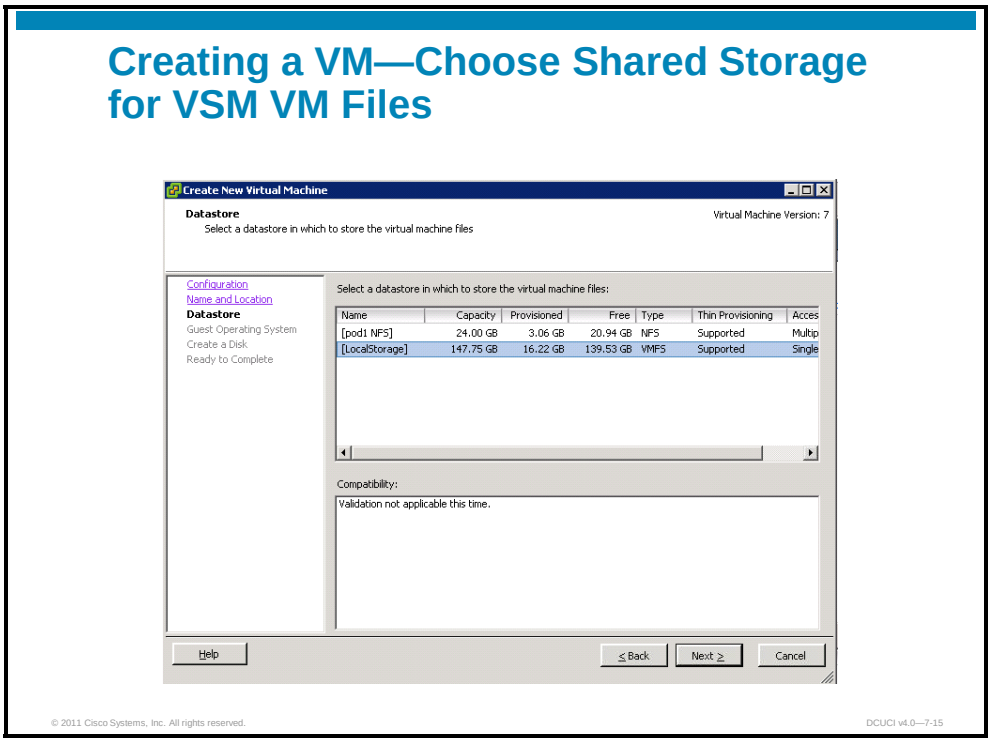

At the next stage, choose a datastore. Click **Next**.

Creating a VM—Choose VSM Guest Operating System

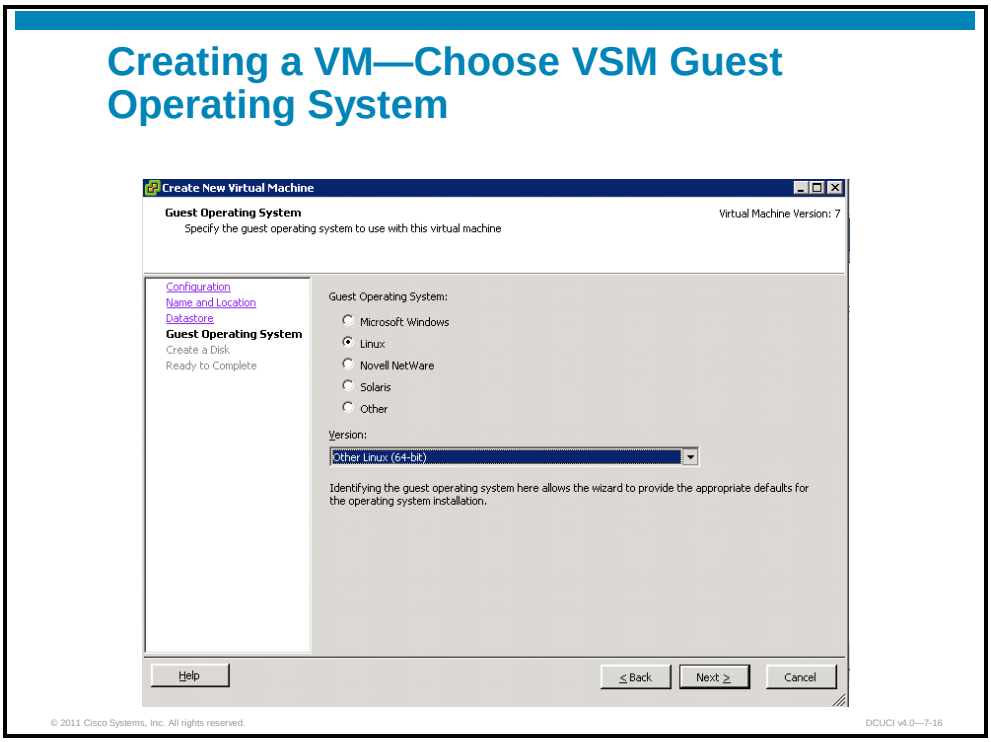

Choose the correct operating system, per VSM requirements.

Creating a VM—Create a VSM Local Disk

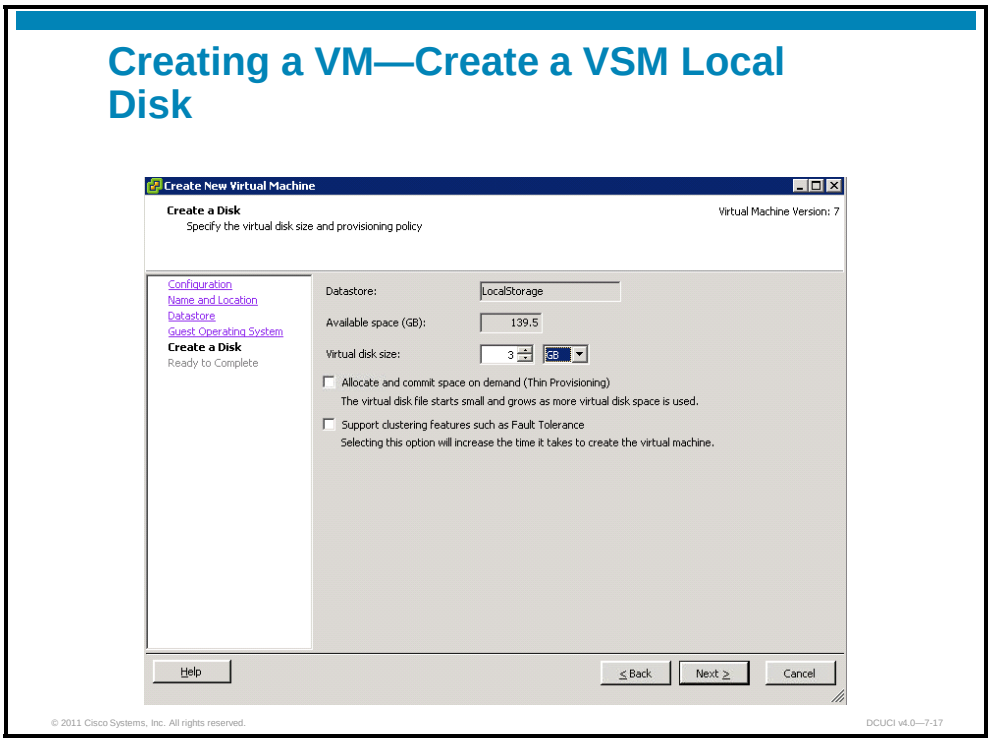

Ensure that you have configured the correct virtual disk size, per VSM requirements. Click **Next**.

Creating a VM—Review VSM Options

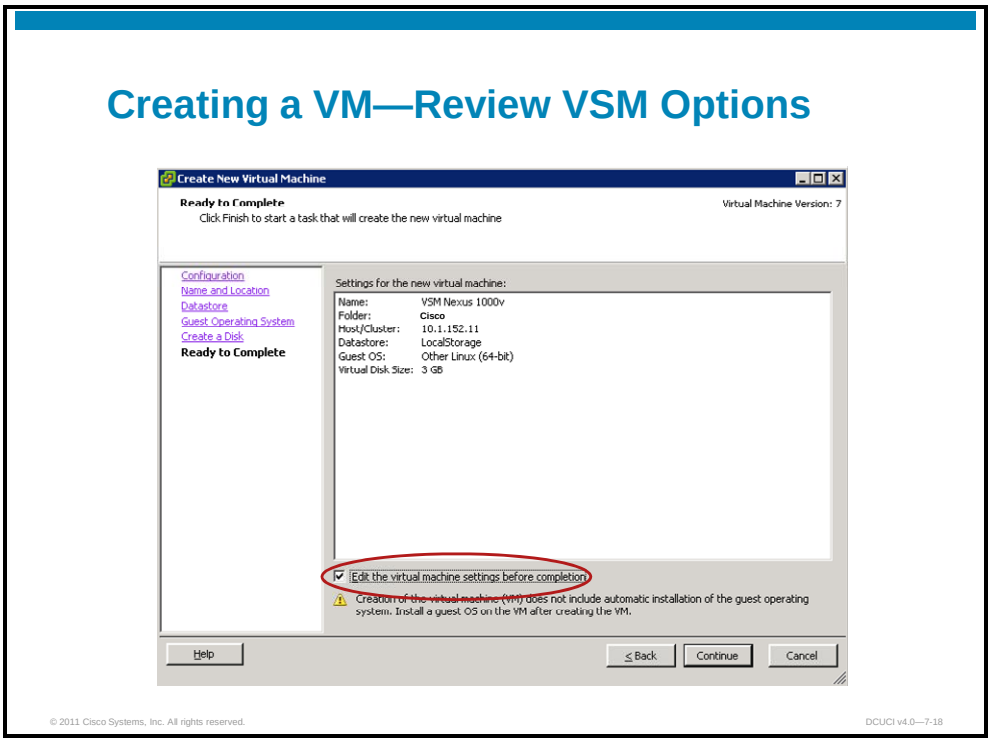

So that you can modify the remainder of the settings for the VSM, check **Edit the Virtual Machine Settings** before completing the configuration. Click **Continue**.

Creating a VM—Adjust the VSM Memory Size

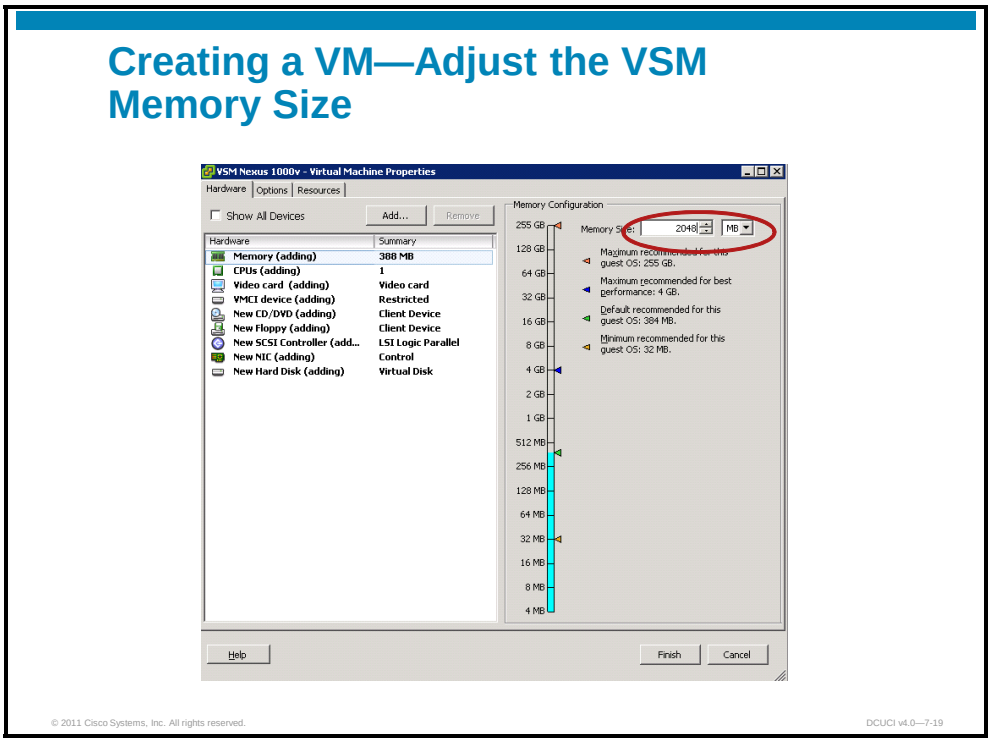

Change the memory size to **2048**; the VSM requires a minimum of 4 GB of memory. Then choose **New NIC**.

Creating a VM—Add the VSM Port Groups

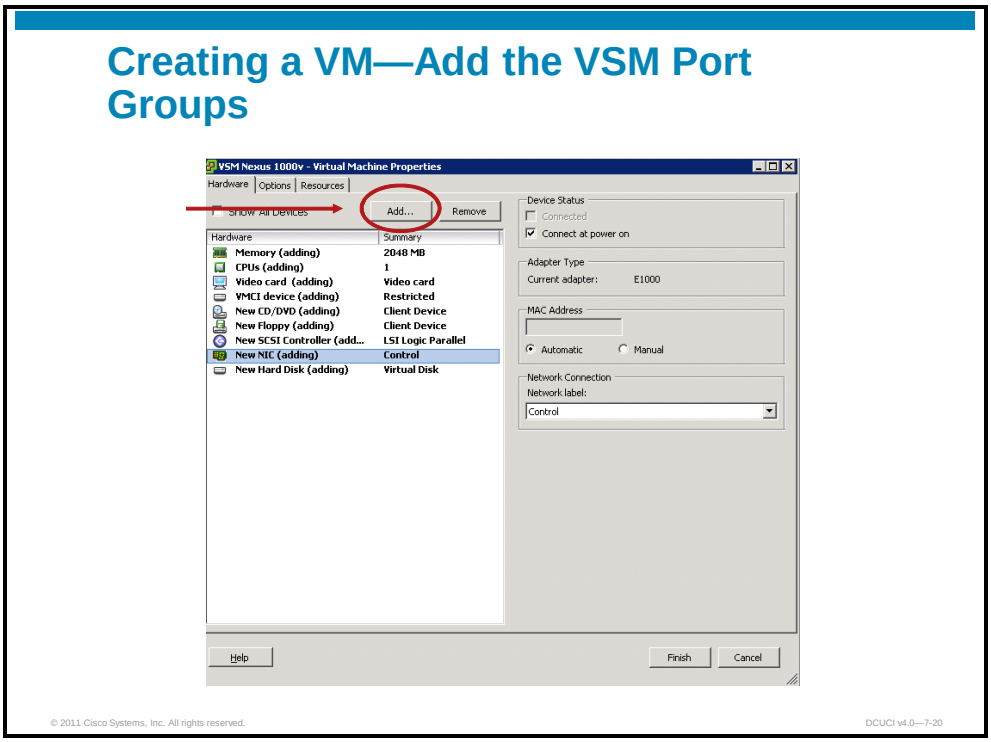

Click **Remove** to remove the new NIC. Then choose **Add**.

Creating a VM—Add the Port Group Adapters

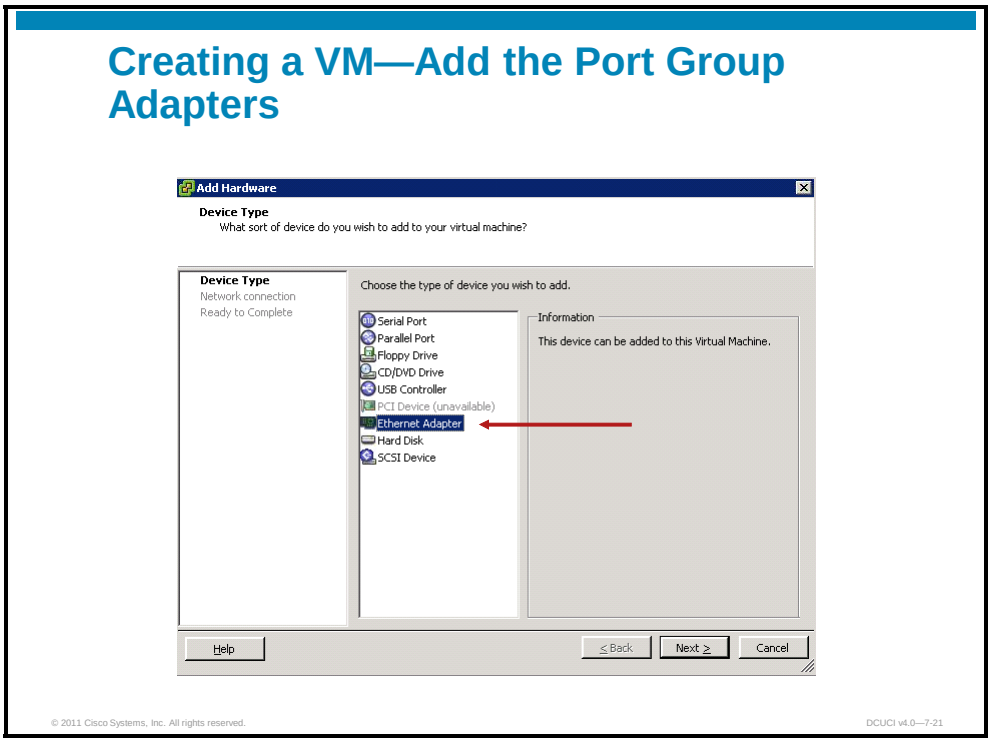

Choose **Ethernet Adapter**. Click **Next**.

Creating a VM—Choose Adapter Driver

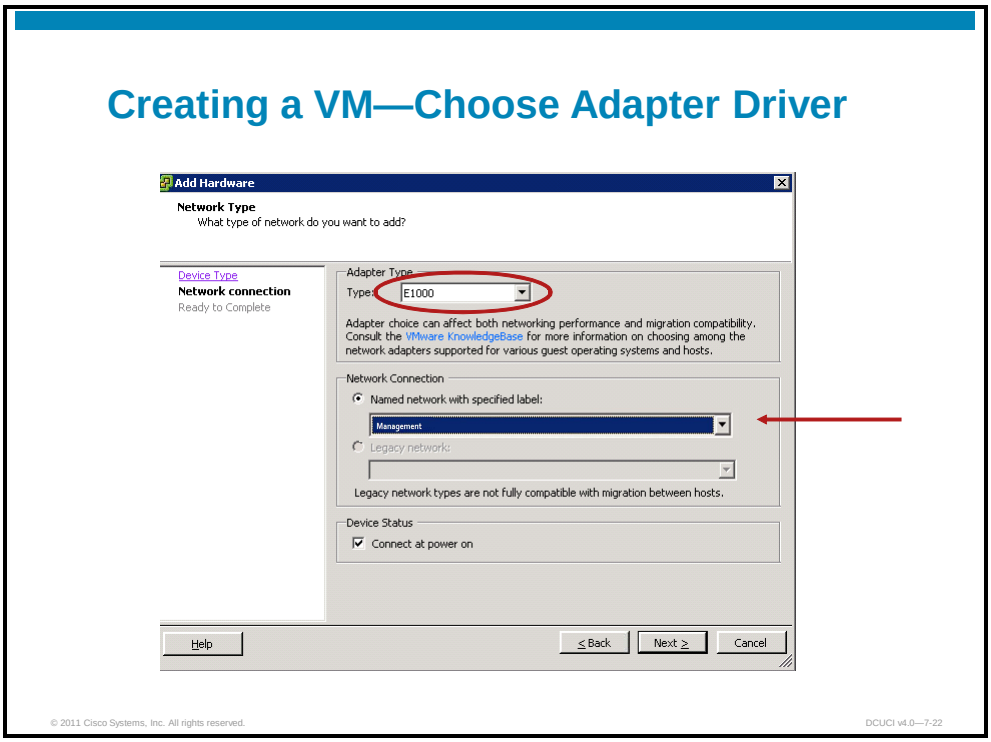

Ensure that adapter type E1000 is specified. Click **Next**.

Creating a VM—Review Adapter Options

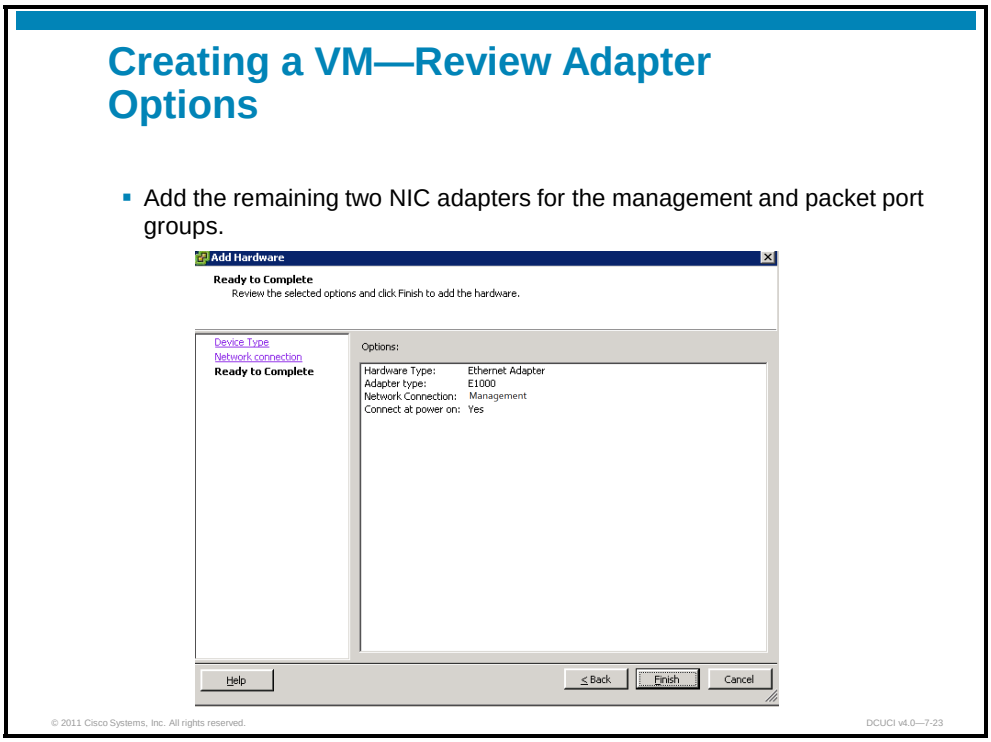

Verify the configuration and click **Finish**. Repeat this process to add two additional NIC adapters, as required by the VSM.

The NIC adapters should be assigned in ascending order:

- **Control:** NIC0
- **Management:** NIC1
- **Packet:** NIC2

Verify Port Group Configuration

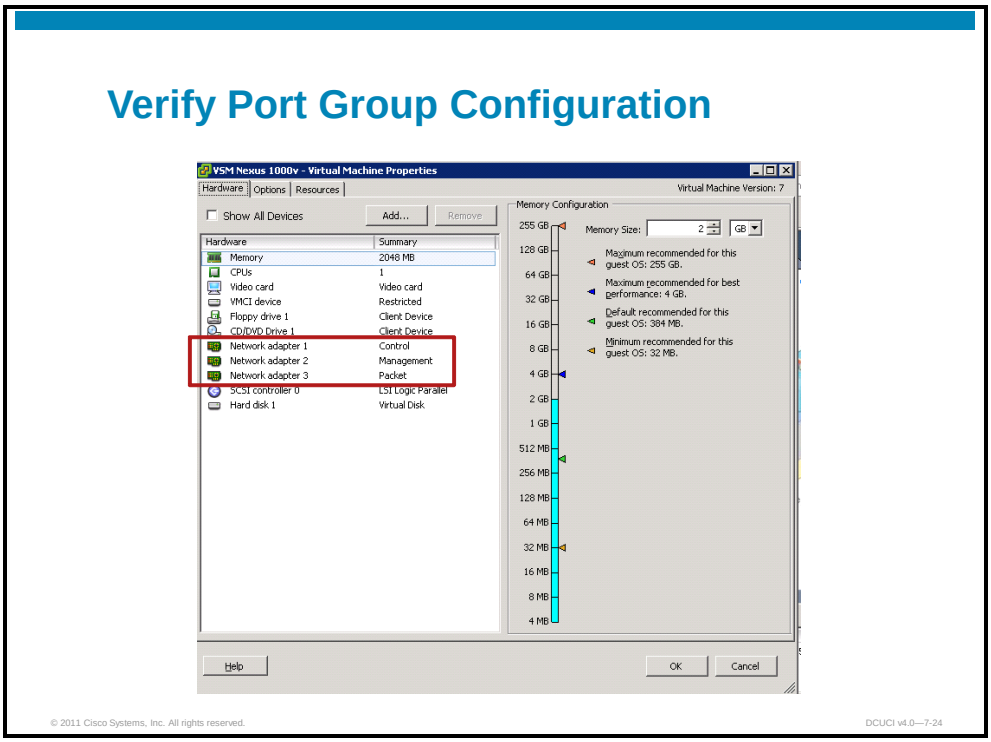

After the VM has been created for the VSM, choose the VSM. Click **Edit Settings**. Verify that the control, management, and packet adapters are assigned in ascending order. The VSM operating system will ensure that the appropriate traffic is assigned to each adapter, based on lowest-to-highest numbering.

Creating a VM—Choose .iso Boot File

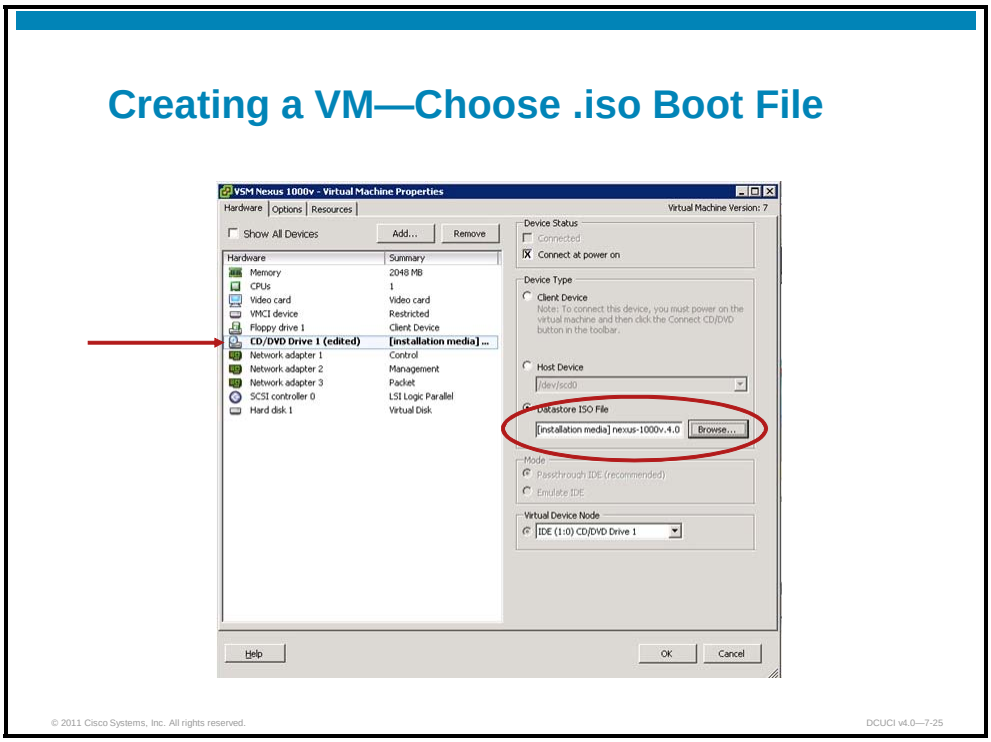

To attach the VSM to a bootable Cisco Nexus1000V, install the .iso boot file or disk. Follow these steps:

1. Choose **CD/DVD Drive**.
2. Choose **Datastore ISO File**.
3. Browse to the Cisco Nexus 1000V VSM .iso file.
4. Click **OK**.

Initial VSM Configuration

This topic discusses the initial VSM setup configuration.

Access the VSM Console

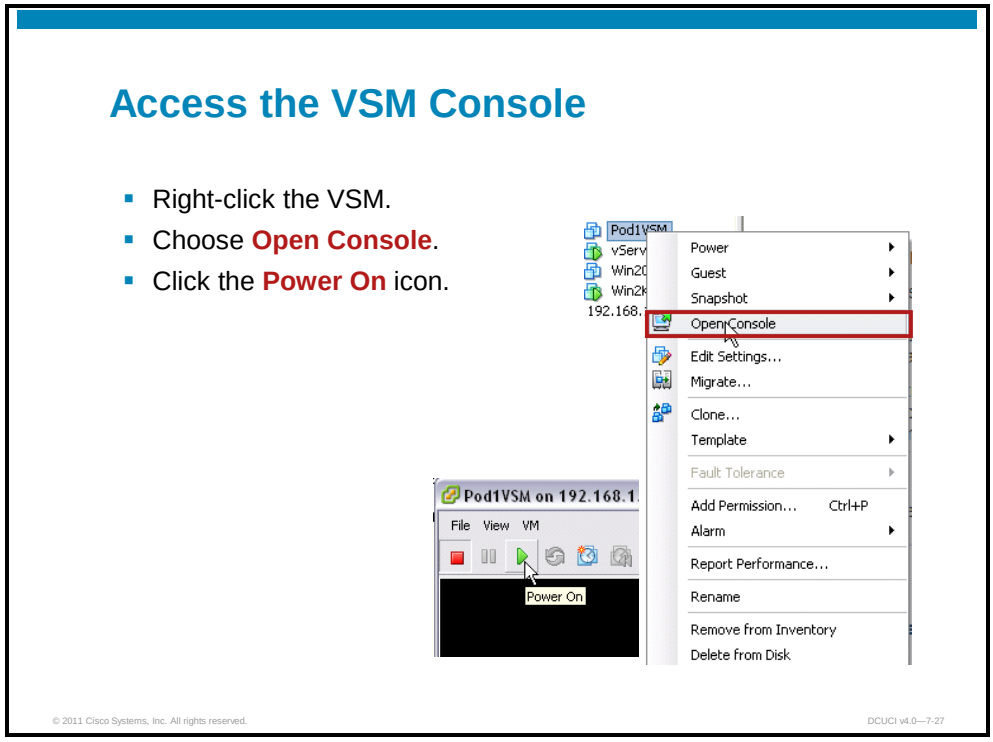

To access the VSM console and to power up the VM, right-click the VSM VM. Choose **Open Console**. When the console screen opens, click the **Power On** icon (the green triangle).

Access the VSM Console (Cont.)

- When the boot menu appears, choose **Install Cisco Nexus 1000V and bring up the new image**.
- This process can take as much as 5 minutes.

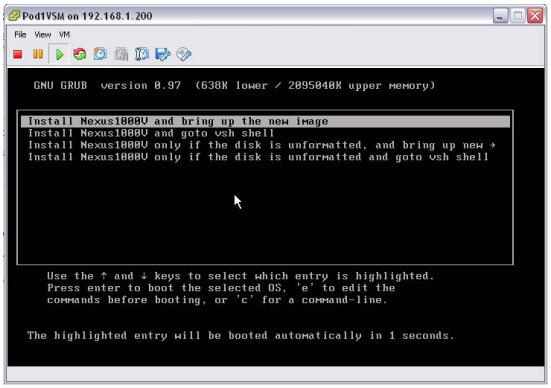

After the virtual machine is powered up, a boot menu appears. Choose **Install Nexus 1000V**. This choice will bring up the new image.

Note	This process can take as much as 5 minutes.

Initial Setup

Initial Setup

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

```
---- System Admin Account Setup ----
Enter the password for "admin": Qwer12345
Confirm the password for "admin": Qwer12345
[########################################] 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
```


The Cisco Nexus 1000V runs the Cisco Nexus Operating System (NX-OS). As with most Cisco devices, a basic setup script is provided when booting a new switch; issuing a write, erase, or reload; or issuing the command setup at the command-line interface (CLI).

The basic setup script leads the administrator through basic connectivity options and switch parameters.

Initial Setup (Cont.)

Enter the configuration for your VSM.

```
Create another login account (yes/no) [n]: no
Configure read-only SNMP community string (yes/no) [n]: no
Configure read-write SNMP community string (yes/no) [n]: no
Enter the switch name: VSM-1
Continue with Out-of-band management configuration? [yes/no]: yes
Mgmt0 IPv4 address: 192.168.110.31
Mgmt0 IPv4 netmask: 255.255.255.0
Configure the default-gateway: (yes/no) [y]: yes
IPv4 address of the default-gateway: 192.168.110.1
Configure Advanced IP options (yes/no)? [n]: no
Enable the telnet service? (yes/no) [y]: no
Enable the ssh service? (yes/no) [y]: yes
Type of ssh key you would like to generate (dsa/rsa) : rsa
Number of key bits <768-2048> : 1024
```

Assign an IP address, netmask, and default gateway for the management interface. By default, Telnet is enabled and Secure Shell (SSH) is disabled. If you enable SSH, then you must configure the security settings according to your needs.

You can configure Network Time Protocol (NTP) to ensure that the VSM clock is synchronized with an external time source.

You will also assign VLAN IDs for the control and packet interfaces. Be careful to use the control and packet VLANs that you configured in vCenter before the Cisco Nexus 1000V installation. This simplistic example uses VLAN 1 for all the traffic (control, management, and packet); however, in a production environment, you should use individual VLANs for each type of traffic (control, management, and packet).

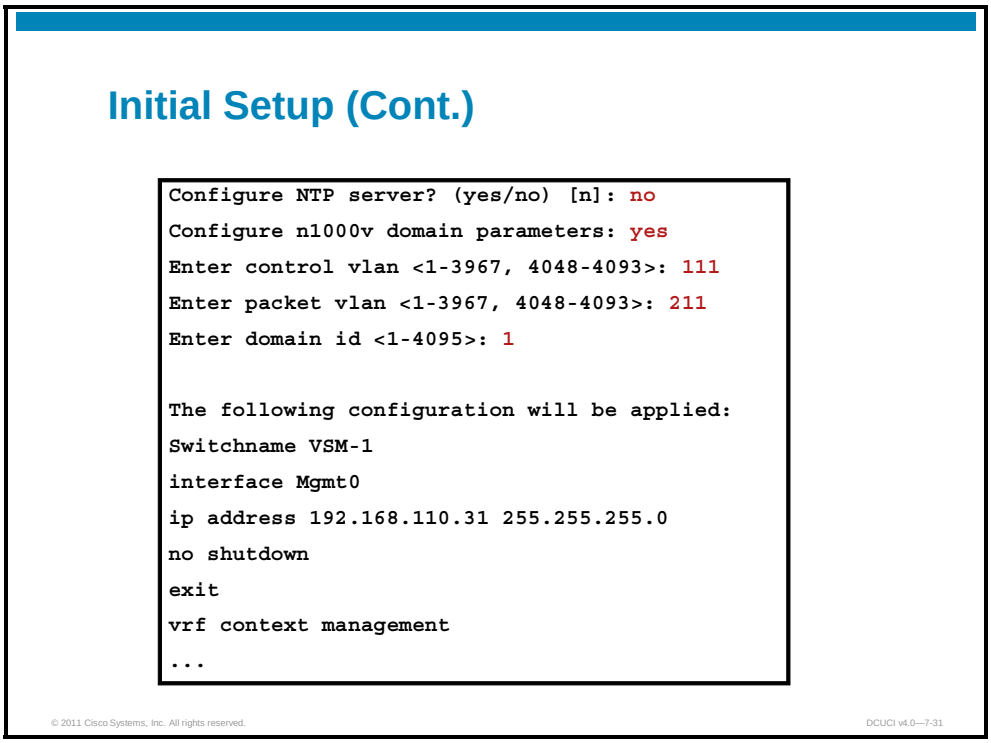

After the dialog is complete, you are presented with a summary of all the configurations that you entered. Review the configurations and enter “n” to quit the wizard and save the configurations or “y” to make changes.

Configure the VSM-to-vCenter Connection

This topic discusses how to configure the VSM-to-vCenter connection.

Install and Register the Plug-In for the New VSM

Install and Register 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.

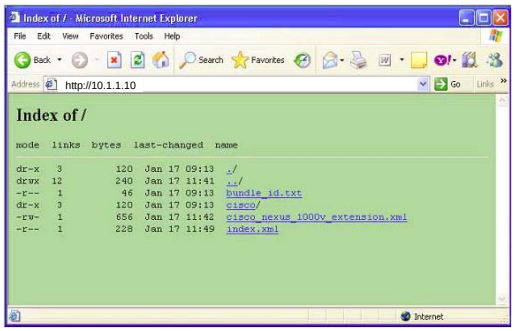

For the VSM to connect to the vCenter server, a plug-in must be installed. You can download the plug-in for the vCenter server from the Cisco Nexus 1000V. After installing the Cisco Nexus 1000V, open an Internet browser. Navigate to the IP address of the Cisco Nexus 1000V VSM and download the file cisco_nexus1000v_extension.xml.

VSM Plug-In

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.

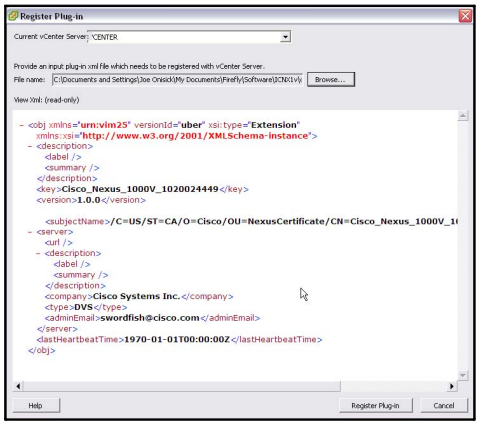

After you download the file, follow these steps:

1. Open the vSphere client. The local host VMware Infrastructure Client dialog box opens.
2. From the Plug-ins menu, choose **Manage Plug-ins**. The Plug-in Manager dialog box opens.
3. Right-click the white space within the dialog box. Choose **New Plug-in** from the pop-up menu.
4. Choose **Browse** and choose **cisco_nexus1000v_extension.xml** file, which you downloaded.

Install the Plug-In for the New VSM

Install the Plug-In for the New VSM

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

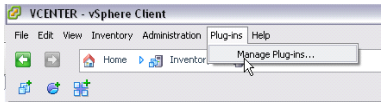

Step 1

- Open the **vCenter Plug-ins** menu.

Step 2

- Right-click the white space and click **New Plug-in**.

Step 3

- Choose the plug-in downloaded from the VSM.
- Click **Open**.

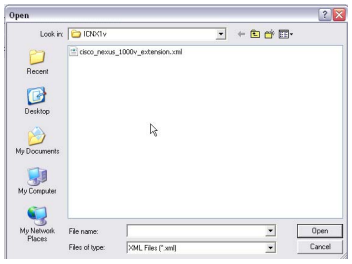

Choose **Register Plug-in**. The plug-in is created and registered.

The VSM plug-in is unique for each VSM and must be installed manually. The plug-in acts as the authentication method for each VSM to prevent unauthorized distributed virtual switch (DVS) configuration from an unregistered VSM.

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

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

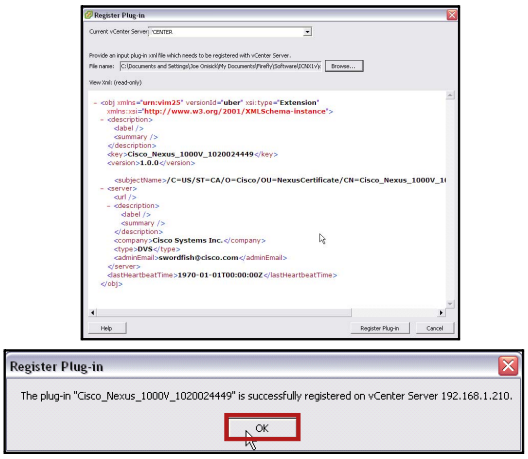

Choose **Register Plug-in**, to complete the process of making the Cisco Nexus 1000V a plug-in of vCenter.

Verify Connectivity

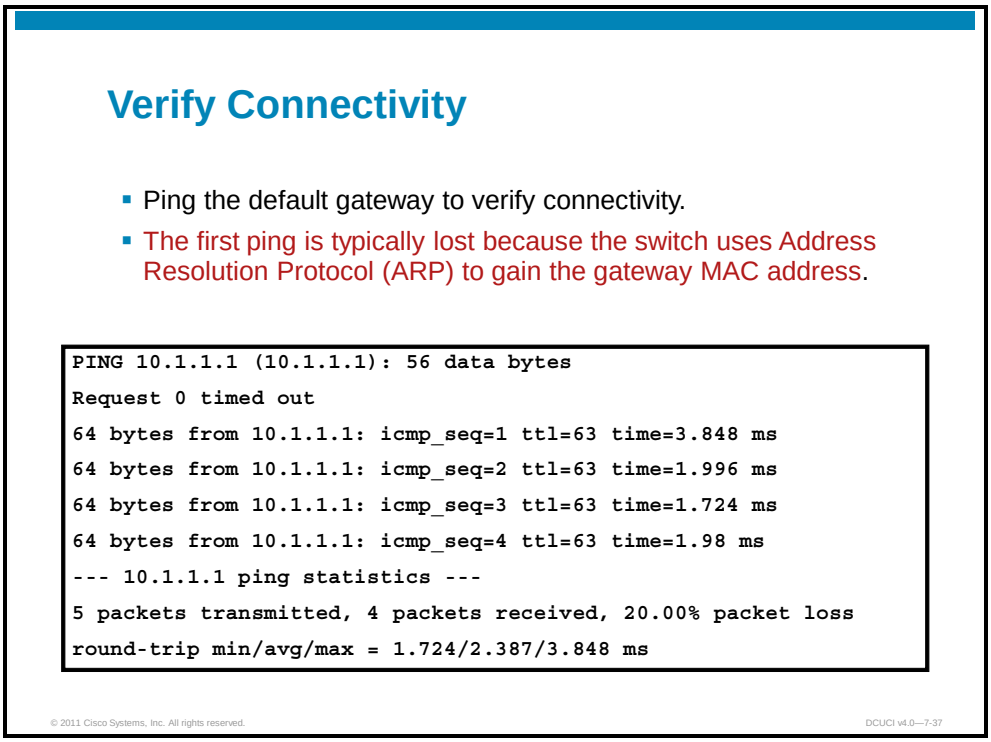

After you complete the initial VSM setup and vCenter plug-in registration, you should verify Layer 3 connectivity to the default gateway of the management network, by issuing the **ping** command.

You can use additional pings to validate connectivity to the vCenter server and the ESX hosts that will be added to the Cisco Nexus 1000V switch instance.

Configure VSM Connectivity

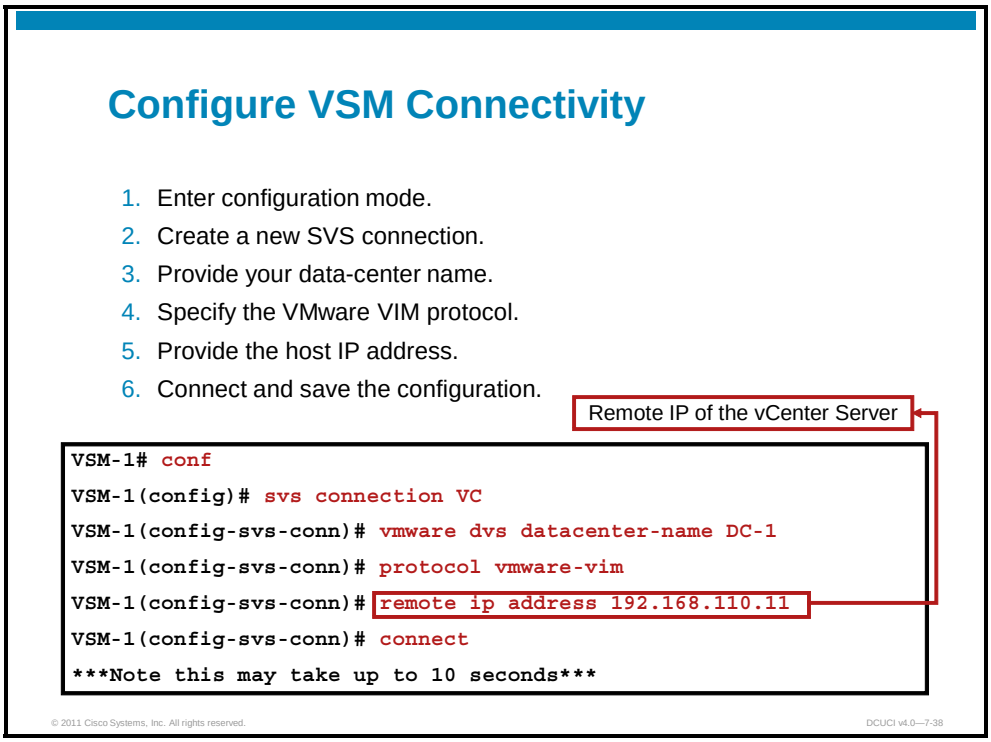

To connect from the VSM to the vCenter server, an SVS connection must be created. The SVS connection requires these parameters:

- SVS connection name
- VMware datacenter name (case sensitive)
- Protocol (currently, only VMware VIM is supported)
- IP address of the vCenter server

Verify VSM Connectivity

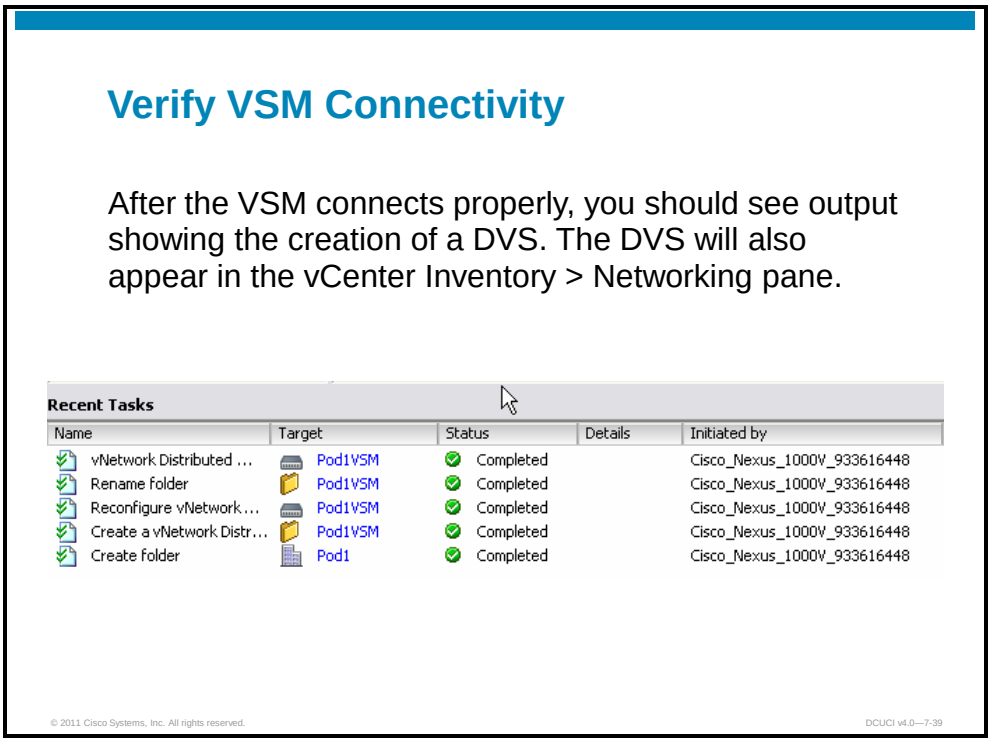

After the VSM is connected to vCenter, the DVS appears in the networking inventory panel of vCenter. You should see the port groups that you configured for control, management, and packet traffic. Some other port groups are also created by default: One is the Unused_Or_Quarantined DVUplinks port group (which connects to physical NICs) and the other is the Unused_Or_Quarantined VMData port group (which is VM facing).

Cisco Nexus 1000V High-Availability Configuration

This topic discusses high-availability configuration of the Cisco Nexus 1000V.

Deploy the Secondary VSM

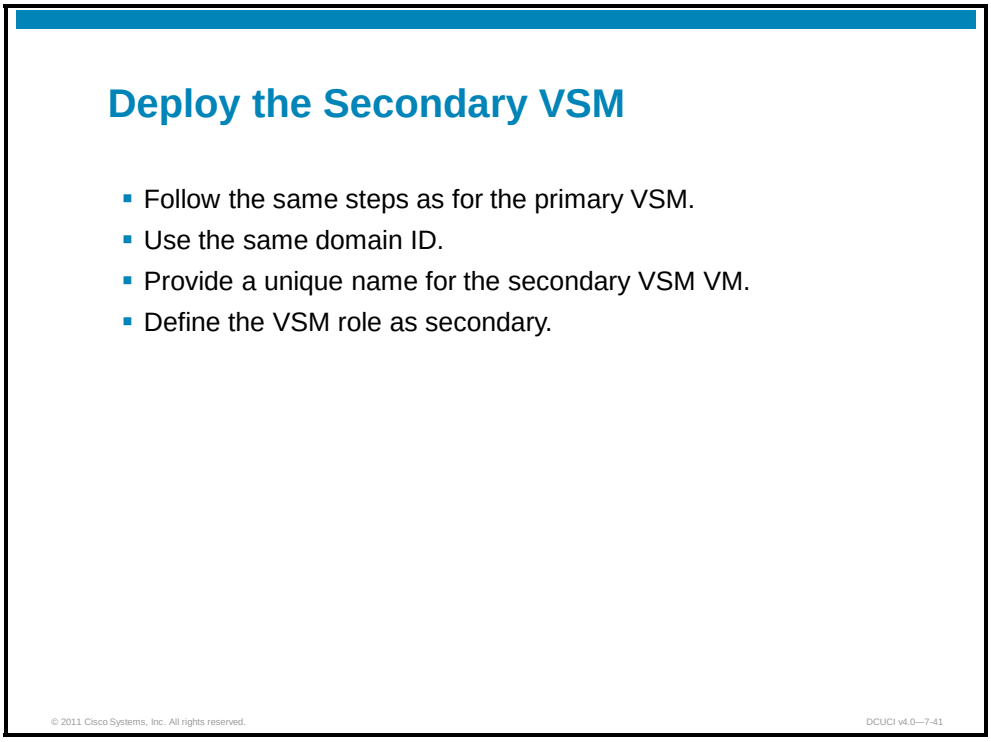

If you run your VSM in a high-availability configuration, then you can configure the secondary VSM as you did the primary VSM. There are two exceptions: You need a unique name for the secondary VSM, and you must define its role as secondary to the primary VSM.

Cisco Nexus 1000V High Availability

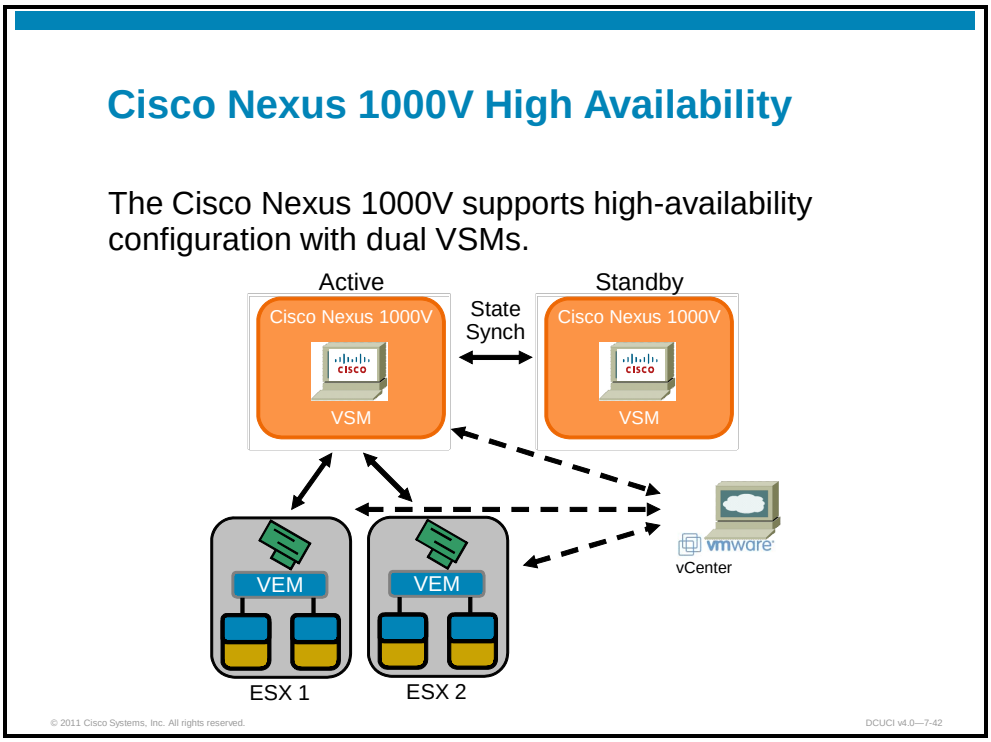

The Nexus 1000V supports high availability by using primary and standby supervisors. Only one active running configuration is managed by connecting to the primary VSM. The primary and standby VSMs are state synchronized. In the event of a manual or automatic switchover, the standby supervisor assumes the role of the primary, whereas the former primary (now the standby) attempts to complete the state synchronization process after it reboots.

During initial setup of the Cisco Nexus 1000V, you need to configure only one management IP address for remote management. The primary supervisor seizes the management IP address because that supervisor is responsible for the active running configuration.

Supervisor Modes

Supervisor Modes

- In Dual Supervisor mode, the active (primary) supervisor is the supervisor in slot 1; the standby (secondary) supervisor resides in slot 2.
- All services are run on the active supervisor and exist in 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.

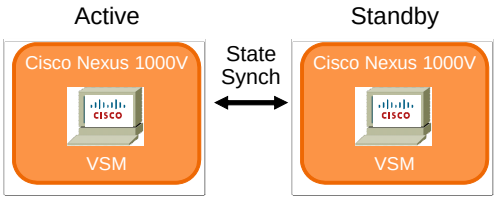

After the initial setup, the primary supervisor appears in slot 1 (or module 1, if the **show module** command is issued from the CLI). The standby, if one is created, appears in slot 2 (or module 2, if the **show module** command is issued from the CLI).

If only a standalone supervisor has been configured, then module 2 does not appear when the **show module** command is issued from the CLI.

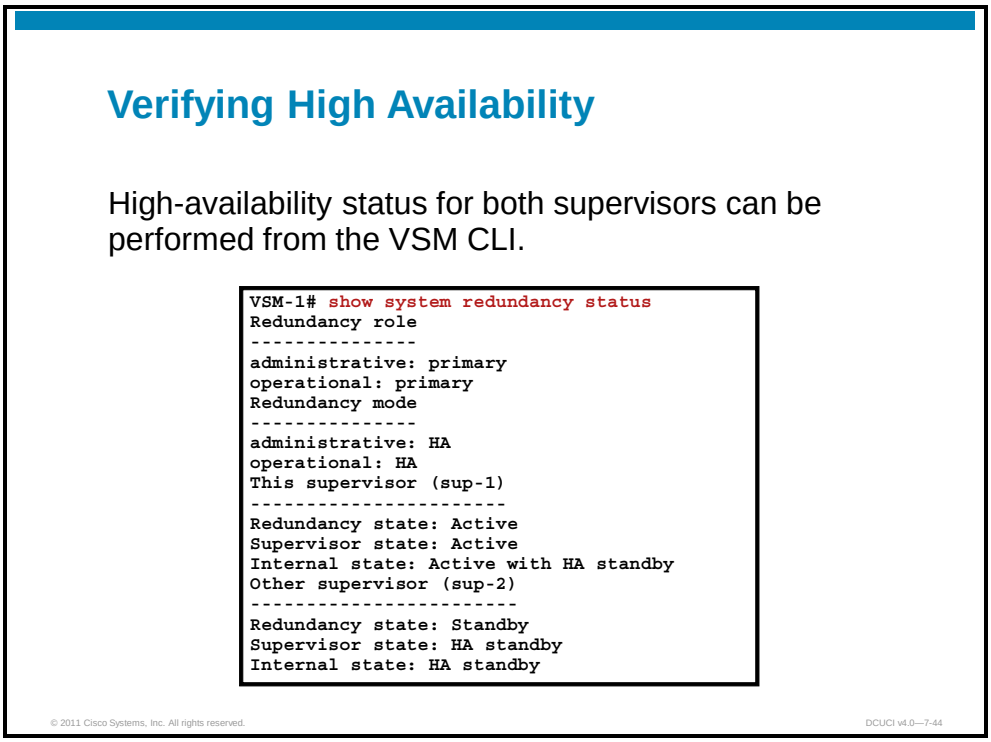

The output of the **show system redundancy status** command shows the primary VSM as the supervisor that is associated with the running configuration. The standby supervisor appears in the active state as the high-availability standby. This output confirms the state synchronized mode of operation when dual VSMs are installed and configured.

Verifying High Availability (Cont.)

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

```
VSM-1# 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
...
```


The output of the **show module** command shows the primary supervisor in slot 1 and the high-availability standby in slot 2. The first ESX host that is added to the Cisco Nexus 1000V switch instance appears in slot 3.

Cisco Nexus 1010 Virtual Services Appliance

This topic discusses the Cisco Nexus 1010 Virtual Services Appliance.

Cisco Nexus 1010 Virtual Services Appliance

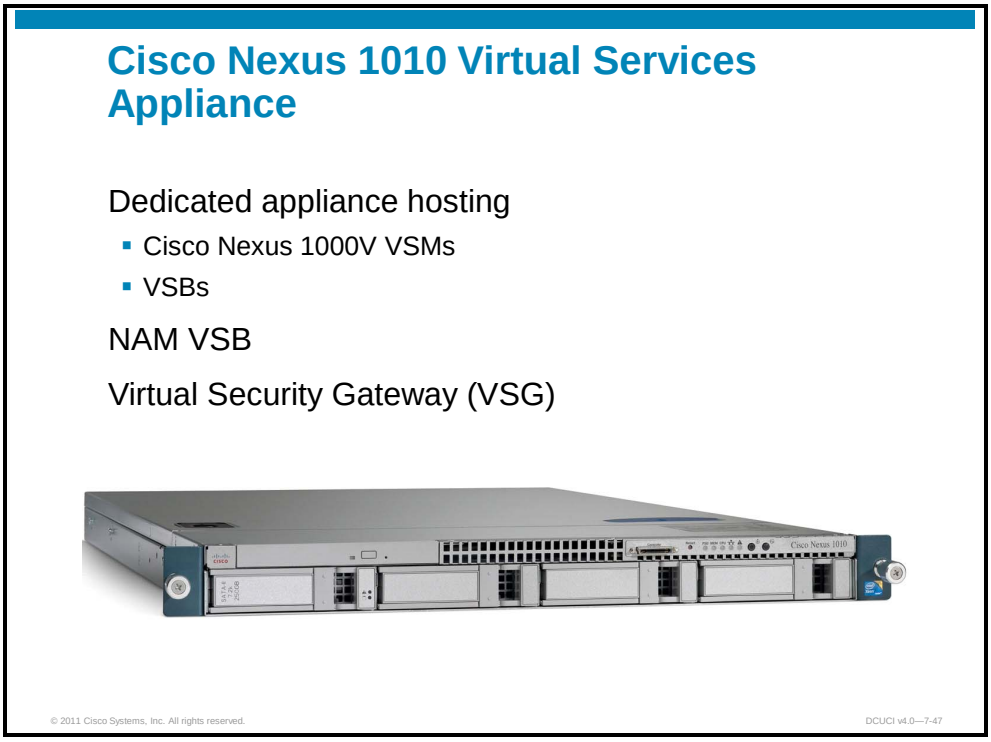

The Cisco Nexus 1010 Virtual Services Appliance is a member of the Cisco Nexus 1000V Series Switches. This appliance hosts the Cisco Nexus 1000V VSM and supports the Cisco Nexus 1000V Network Analysis Module (NAM) Virtual Service Blade (VSB), to provide a comprehensive solution for virtual access switching. Because the Cisco Nexus 1010 provides dedicated hardware for the VSM, the platform makes virtual access switch deployment much easier for the network administrator. With its support for additional VSBs, such as the Cisco Nexus 1000V NAM VSB, the Cisco Nexus 1010 is a crucial component of a virtual access switch solution.

Cisco Nexus 1010 Hardware Configuration

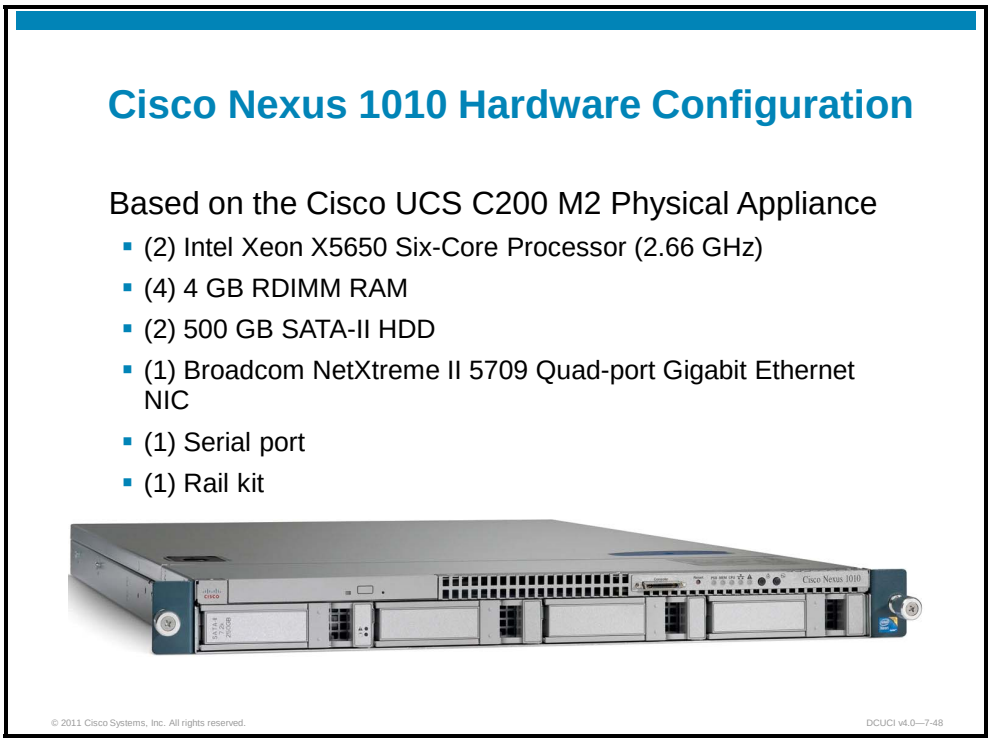

With the introduction of the Cisco Nexus 1010, the VSM now has dedicated hardware. Therefore, network administrators can install and configure virtual access switches like they install physical switches. The dedicated VSM hardware is especially helpful in a data-center power-up, because there is no dependency in finding server resources for the VSM. Thus, the Cisco Nexus 1010 enables network administrators to manage the Cisco Nexus 1000V virtual access switch just like other physical switches and to scale server-virtualization deployments.

The figure summarizes the hardware characteristics, based on the Cisco UCS C200 Series Rack-Mount Server architecture.

Cisco Nexus 1010 Virtual Service Appliance (Cont.)

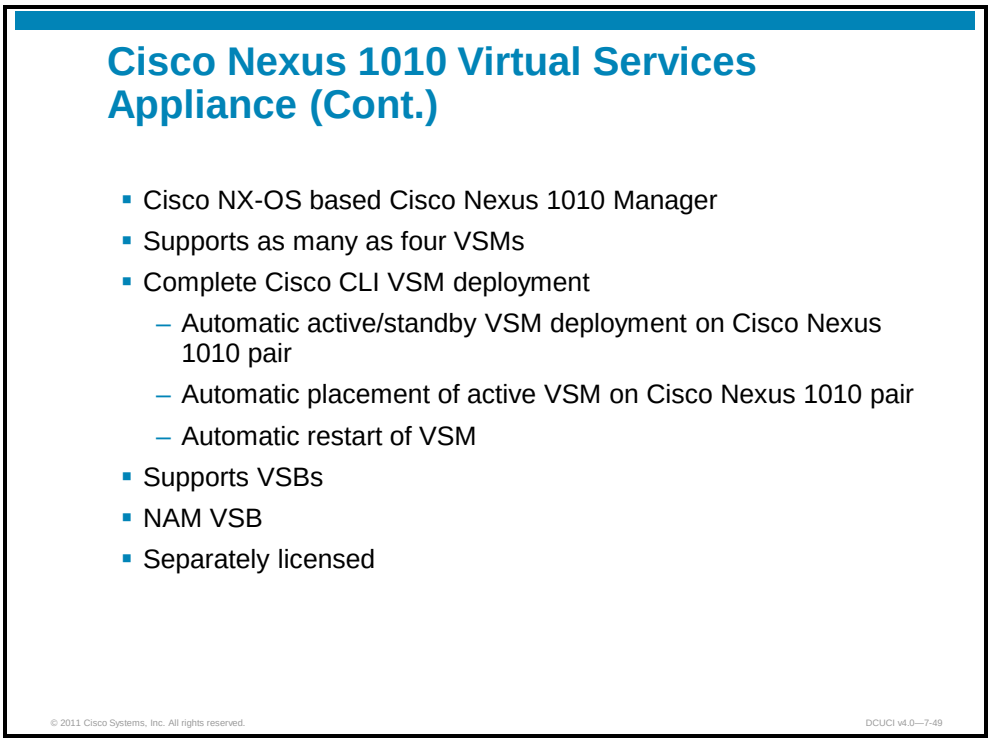

The Cisco Nexus 1010 runs Cisco NX-OS. The appliance supports as many as four instances of the VSM, running in high-availability mode (active-standby pairs). The appliance also supports automatic restart of the VSM, as well as automatic placement of the active VSM on the high-availability pair.

The Cisco Nexus 1010 also supports the VSB and NAM, both of which are licensed separately.

Cisco Nexus 1010 Internal Architecture

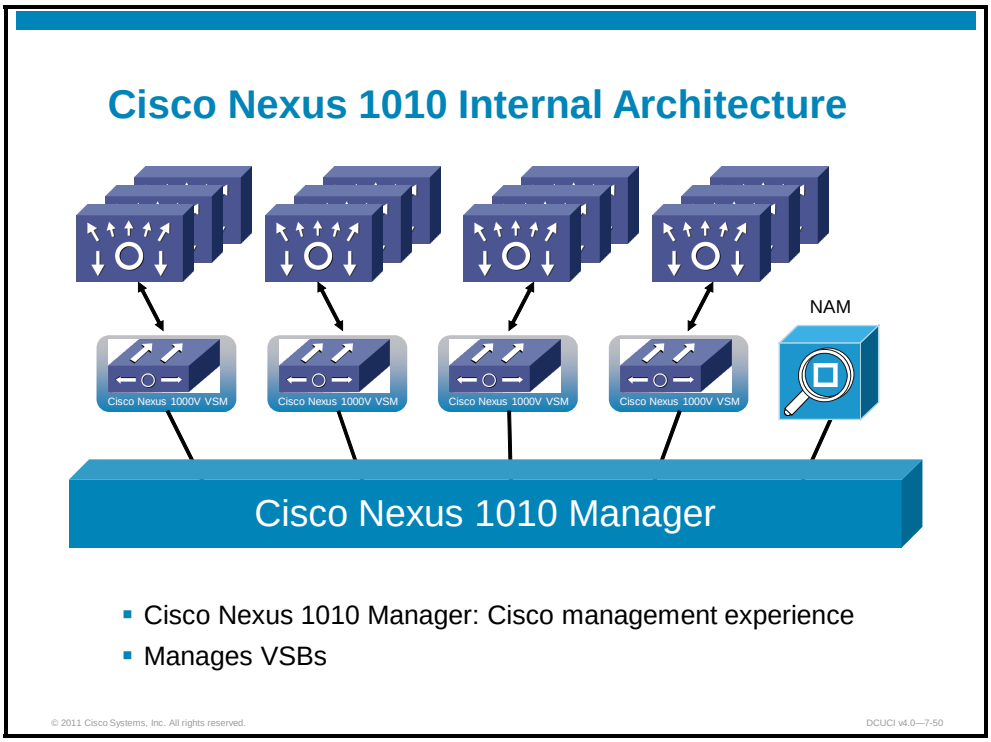

This figure shows the internal architecture of the Cisco Nexus 1010. The Cisco Nexus 1010 contains Cisco Nexus 1010 Manager, which is based on Cisco NX-OS. The appliance can host as many as four VSMs and supports the Cisco Nexus 1000V NAM VSB. Therefore, in addition to hosting Cisco Nexus 1000V VSMs, the Cisco Nexus 1010 becomes a platform for other networking services. The appliance will also support other VSBs in the future.

Because the Cisco Nexus 1010 uses the same VSM as the Cisco Nexus 1000V Series, the Cisco Nexus 1000V Series solution with Cisco Nexus 1010 has all the features of the Cisco Nexus 1000V Series. Because Cisco Nexus 1010 Manager is based on Cisco NX-OS, the network administrator has a familiar interface for installing and configuring Cisco Nexus 1010. Cisco Nexus 1010 Manager also supports Cisco NX-OS high availability, allowing a standby Cisco Nexus 1010 appliance to become active if the primary Cisco Nexus 1010 appliance fails.

Cisco Nexus 1010 Benefits

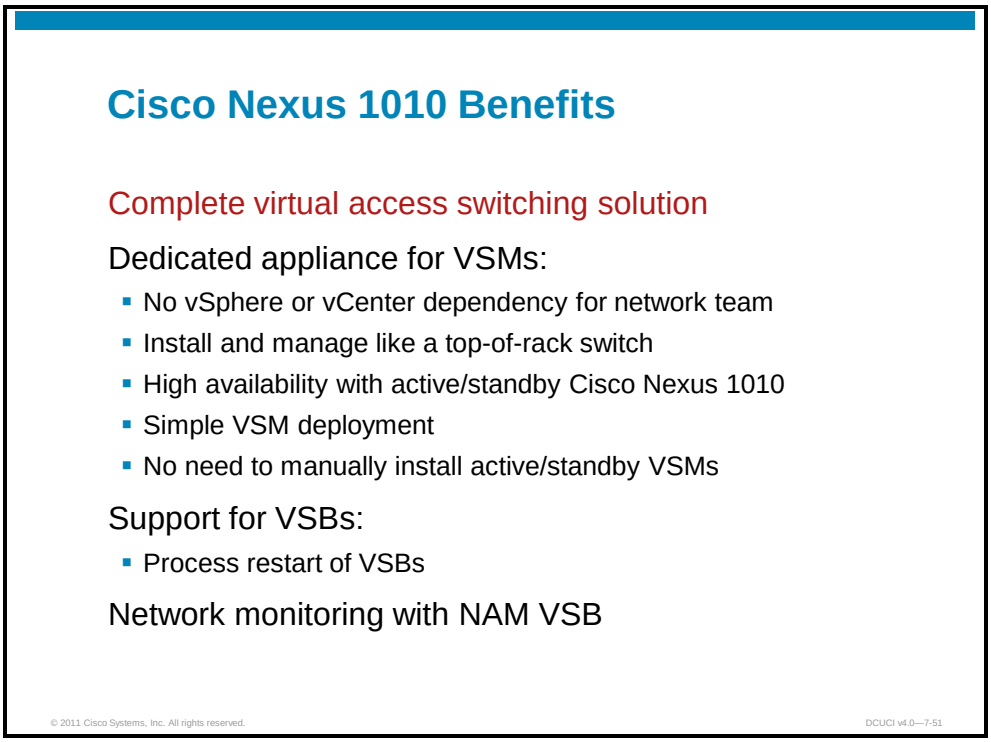

Cisco Nexus 1000V Series Switches are intelligent VM access switches that are designed for VMware vSphere environments running Cisco NX-OS. Operating inside the VMware ESX hypervisor, the Cisco Nexus 1000V Series supports Cisco Virtual Network Link (VN-Link) server-virtualization technology. This technology provides the following:

- Policy-based VM connectivity
- Mobile VM security and network policy
- Nondisruptive operational model for server virtualization and networking teams

When server virtualization is deployed in the data center, virtual servers typically are not managed like physical servers. Server virtualization is treated as a special deployment. This treatment leads to longer deployment times and requires a greater degree of coordination among server, network, storage, and security administrators. With the Cisco Nexus 1000V Series, you can have a consistent networking feature set and provisioning process, all the way from the VM access layer to the core of the data-center network infrastructure. Virtual servers can now use the same network configuration, security policy, diagnostic tools, and operational models as their physical-server counterparts that are attached to dedicated physical network ports. Virtualization administrators can access predefined network policy that follows mobile VMs to ensure proper connectivity. This approach saves valuable time that administrators can use to focus on VM administration.

Cisco Nexus 1010 Connectivity

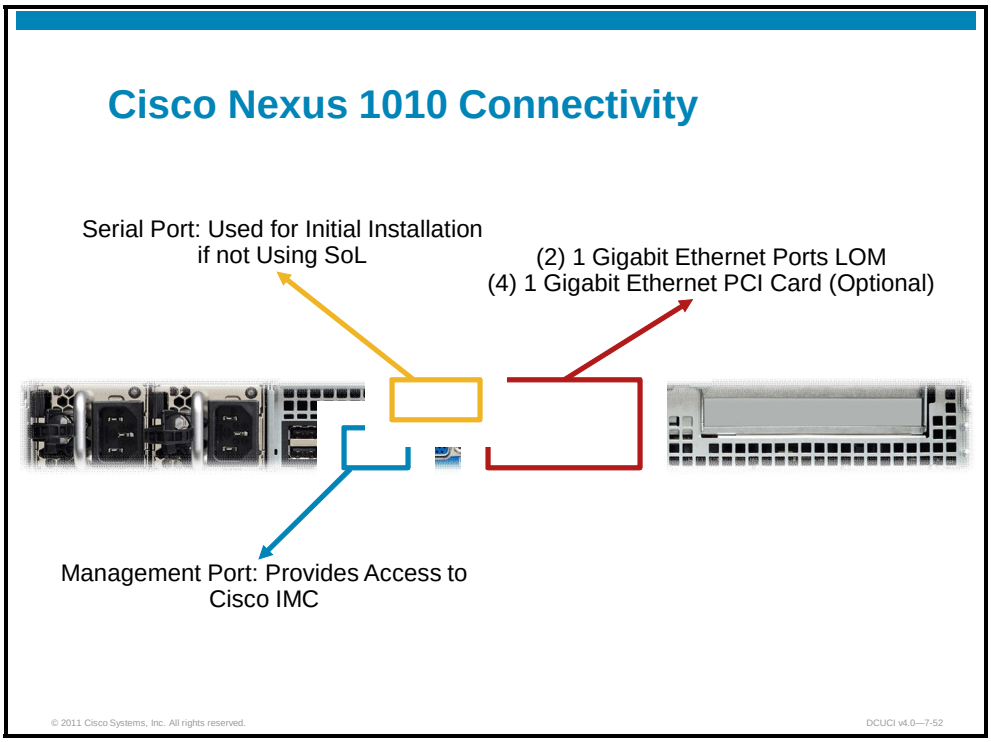

The rear panel image of the Cisco Nexus 1010 appliance illustrates the management connectivity that is provided by the serial port that is used for serial over LAN (SoL), as well as the management port, which provides access to the Cisco Integrated Management Controller(IMC).

Network connectivity may be provided by using the two Gigabit Ethernet LAN on Motherboard (LOM) ports or the four Gigabit Ethernet ports on the Broadcom PCI NIC.

Cisco Nexus 1010 Connectivity (Cont.)

There are four options to connect the Cisco Nexus 1010 to the physical network by using the six available 1 Gigabit Ethernet interfaces.

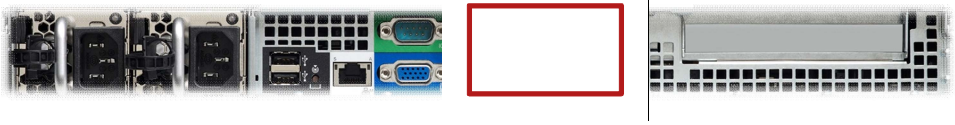

Within the Cisco Nexus 1010 CLI, these options are called network options and influence on which interface different traffic will be configured.

There are four types of traffic available on the system:

Management | Control | Packet | Data

The figure shows the rear panel of the Cisco Nexus 1010, with the physical connectivity provided by the four-port, Broadcom Gigabit Ethernet NIC. The administrator has flexibility in configuring connectivity of the control, management, packet, and data traffic types, by using the **network options** command within the Cisco Nexus 1010 CLI.

Cisco Nexus 1010 High Availability

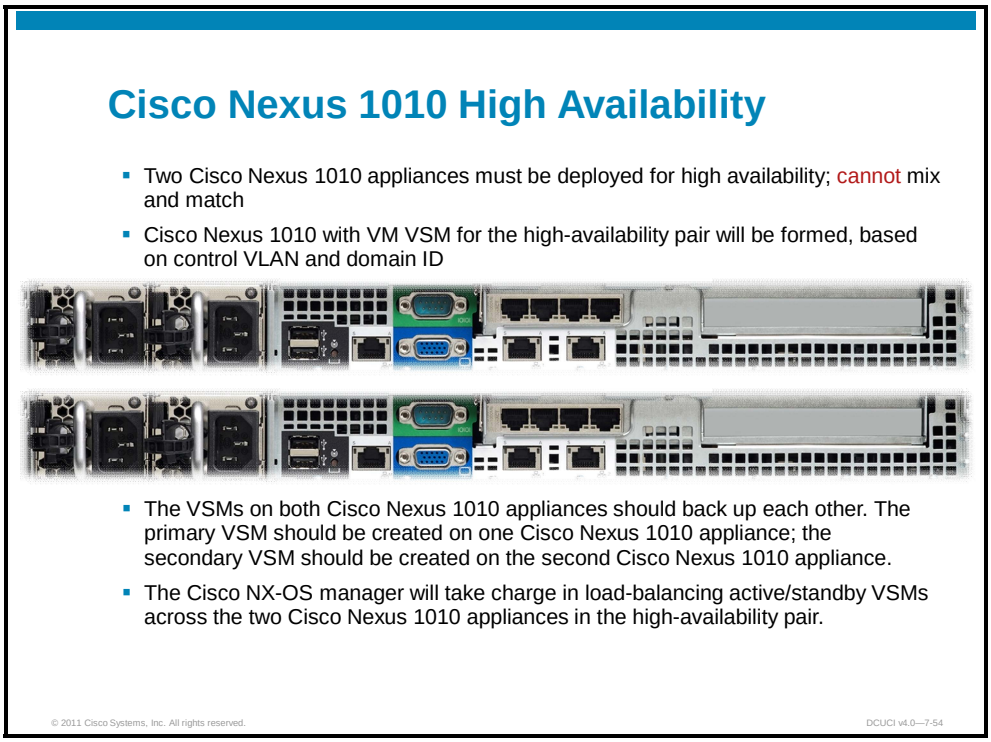

The figure shows the high availability that is built into Cisco Nexus 1010 Manager. Because Cisco Nexus 1010 Manager is built on Cisco NX-OS, it offers active-standby redundancy. If one Cisco Nexus 1010 Manager instance fails, the other Cisco Nexus 1010 Manager instance automatically takes over and continues operation. In addition, Cisco Nexus 1010 Manager automatically places the active VSM to balance the distribution and reduce the potential fault domain.

Cisco Nexus 1010 VSB

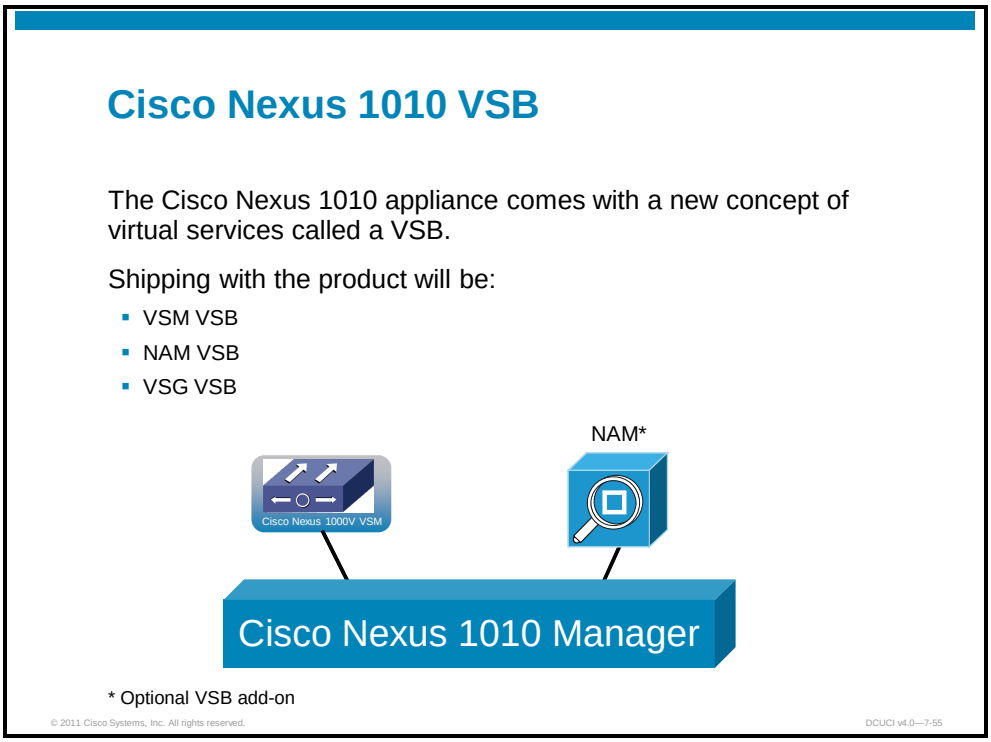

The VSB provides expansion capabilities so that new features can be added in the future. Cisco Nexus 1010 Manager enables customers to install, configure, and manage VSBs. If the VSB stops, Cisco Nexus 1010 Manager automatically restarts it. The Cisco Nexus 1000V NAM VSB takes advantage of these capabilities to provide a robust, complete solution for the virtualized data center.

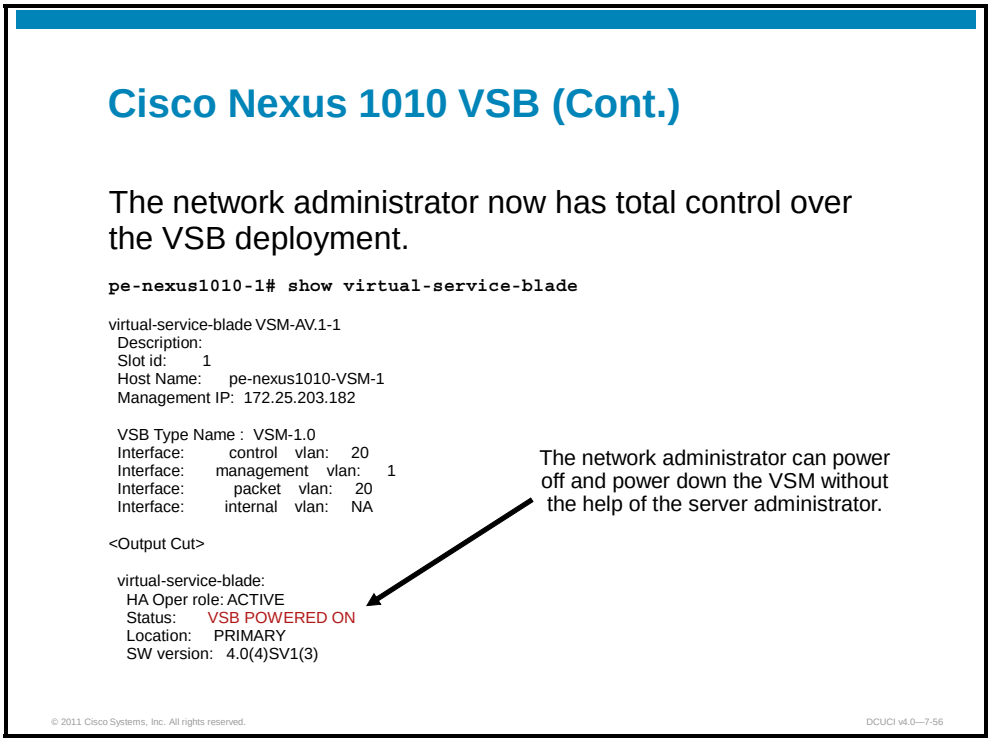

The output of the **show virtual-service-blade** command that is issued from the Cisco Nexus 1010 primary VSM CLI shows parameters that are associated with the VSB. The hardware version, hostname, and management IP address of the primary VSM is listed, as are the control, management, and packet VLANs.

The operational status is powered on and the high-availability operational role is active.

Comparing the Cisco Nexus 1000V and Cisco Nexus 1010 Virtual Services Appliance

This topic compares the Cisco Nexus 1000V and the Cisco Nexus 1010 VSM Appliance.

Architectural Comparison

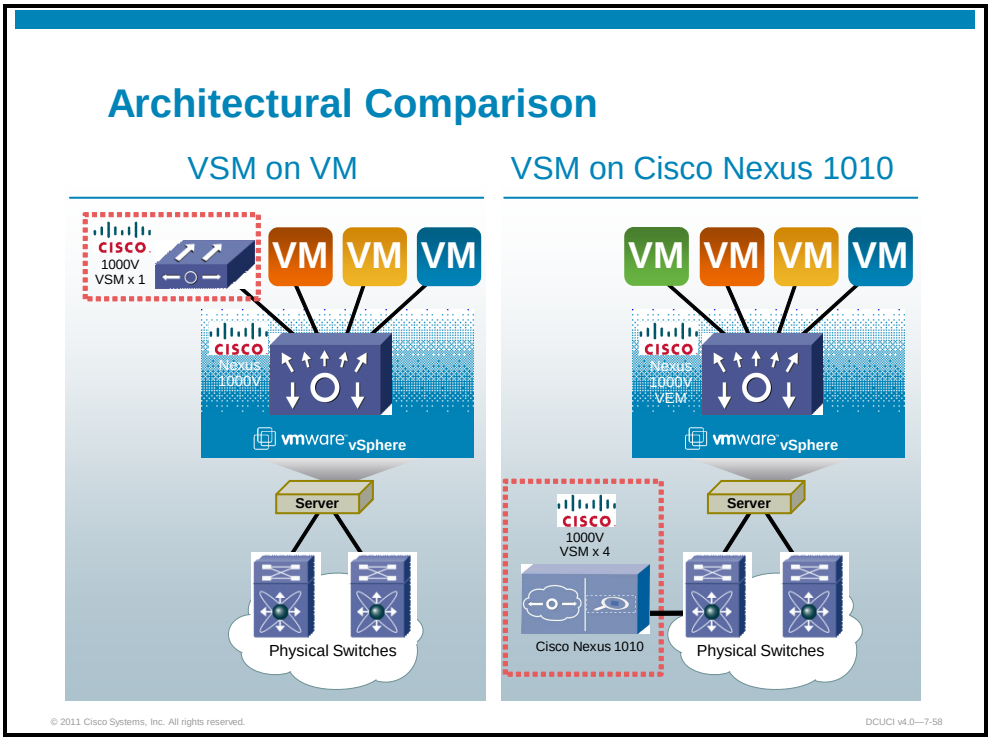

The Cisco Nexus 1000V VSM installs as a VM on an ESX or ESXi host and supports high availability. The VSM communicates directly with the VEM that is installed on every ESX or ESXi host that is a part of the Cisco Nexus 1000V instance. A single instance of the Cisco Nexus 1000V VSM can support as many as 64 ESX hosts that are managed within a single data-center object within vCenter. Additional Cisco Nexus 1000V instances may be added to expand beyond 64 hosts.

The Cisco Nexus 1010 VSM is supported in hardware as a virtual appliance. The VSM connects directly to the physical access switch and requires the proper configuration of control, management, and packet VLANs. The Cisco Nexus 1010 VSMs communicate directly with the VEMs, which are installed on every ESX or ESXi host that is a part of the Cisco Nexus 1010 instance.

The Cisco Nexus 1010 can support as many as four high-availability instances (primary and standby VSMs), scaling to manage as many as 256 ESX or ESXi hosts.

Feature Comparison

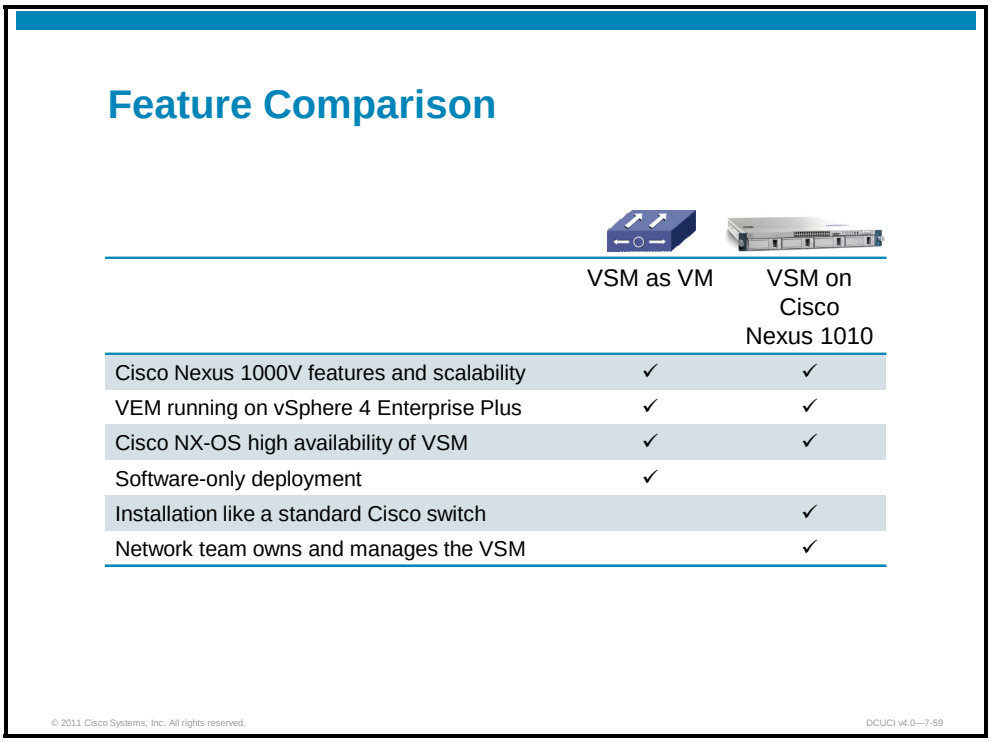

The figure shows a comparison deployment of a VSM as a VM and of a VM on the Cisco Nexus 1010 appliance. For customers who want a complete software deployment of the Cisco Nexus 1000V Series, deployment of the VSM as a VM provides flexibility in VSM placement and mobility with VMware vMotion.

For network administrators who want greater control over the management of the VSM, the Cisco Nexus 1010 solution provides a complete Cisco NX-OS experience in installing the Cisco Nexus 1000V virtual access switch. In addition, the Cisco Nexus 1010 appliance offers fewer dependencies when the data center is powered on because the VSM can be initiated at the same time as the VEMs.

Summary

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

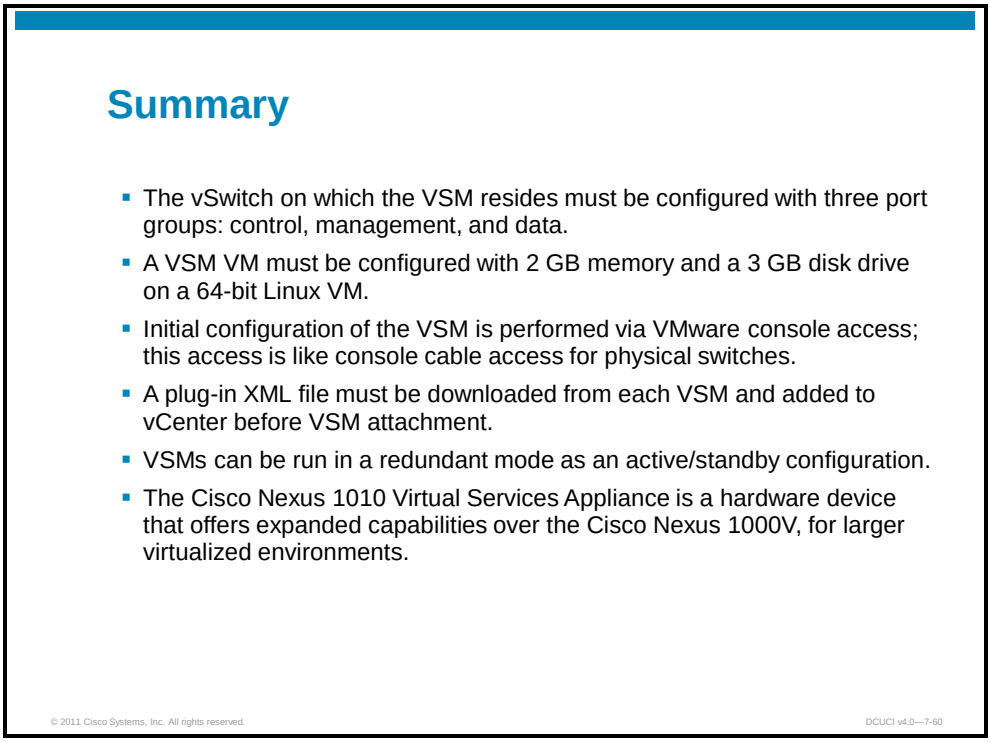

Lesson 5

Configuring Basic Cisco Nexus 1000V Networking

Overview

When the Virtual Supervisor Module (VSM) is installed and linked to the Virtual Ethernet Modules (VEMs), implementers are faced with a vast array of features that need to be integrated to meet the requirements of a customer's design. This lesson presents the basic networking capabilities of the Cisco Nexus 1000V and describes the use of port profiles to define VLANs, private VLANs (PVLANs), uplink and virtual Ethernet (vEthernet) port profiles, and port channels. The lesson describes how to install VEMs by using the VMware vCenter Update Manager, as well as how to back up a VSM configuration. Finally, the lesson describes VMware vMotion in relation to the port profile mobility that Cisco Nexus 1000V offers.

Objectives

Upon completing this lesson, you will be able to describe the networking capabilities of the Cisco Nexus 1000V, including the use of port profiles, VEM installation, and VSM configuration backup. This ability includes being able to meet these objectives:

- Describe port profiles within the VSM
- Describe VLAN configuration and assignment to port profiles
- Describe PVLAN configuration
- Describe the configuration of uplink port profiles
- Describe the configuration of vEthernet data port profiles
- Describe Cisco Nexus 1000V port channel configuration
- Describe the process of adding hosts to a VSM
- Describe how to back up a VSM configuration to a TFTP server
- Describe the vMotion process and Cisco Nexus 1000V port profile mobility

Port Profiles Within the VSM

This topic discusses port profiles within the VSM.

Cisco Nexus 1000V Port Profiles

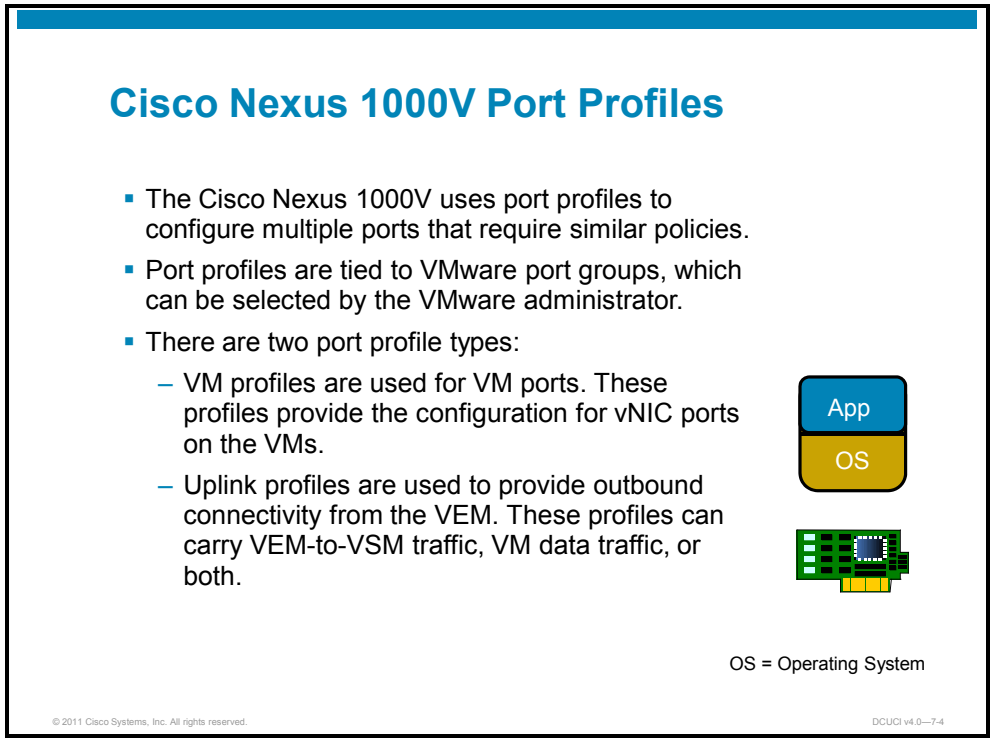

The Cisco Nexus 1000V uses a new concept called port profiles. Port profiles are used to provide port configuration for the ports that are assigned to the profile. There are two types of port profiles: virtual machine (VM) and uplink. Uplink profiles provide all outbound communication for VMs, as well as all VSM-to-VEM communication on the control and packet 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 using port profiles, groups of ports that require the same network configuration can be configured rapidly. When a port profile has been created and configured, ports can be assigned to the profile and receive the configuration that is contained in that profile.

Port profiles tie directly to VMware port groups on the distributed virtual switch (DVS) that the Cisco Nexus 1000V controls. After a profile is created, a VMware port group name can be assigned to the profile. By default, the port group name is the same as the port profile name, but this setting is configurable. When a port profile is enabled, the corresponding port group is created on the DVS within VMware vCenter Server.

VMware administrators assign ports to a port group from the VMware 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 VEM. Any configuration changes to the port profile are immediately assigned to member ports and are stored in the Cisco Nexus 1000V running configuration file.

Port Profiles and Port Groups

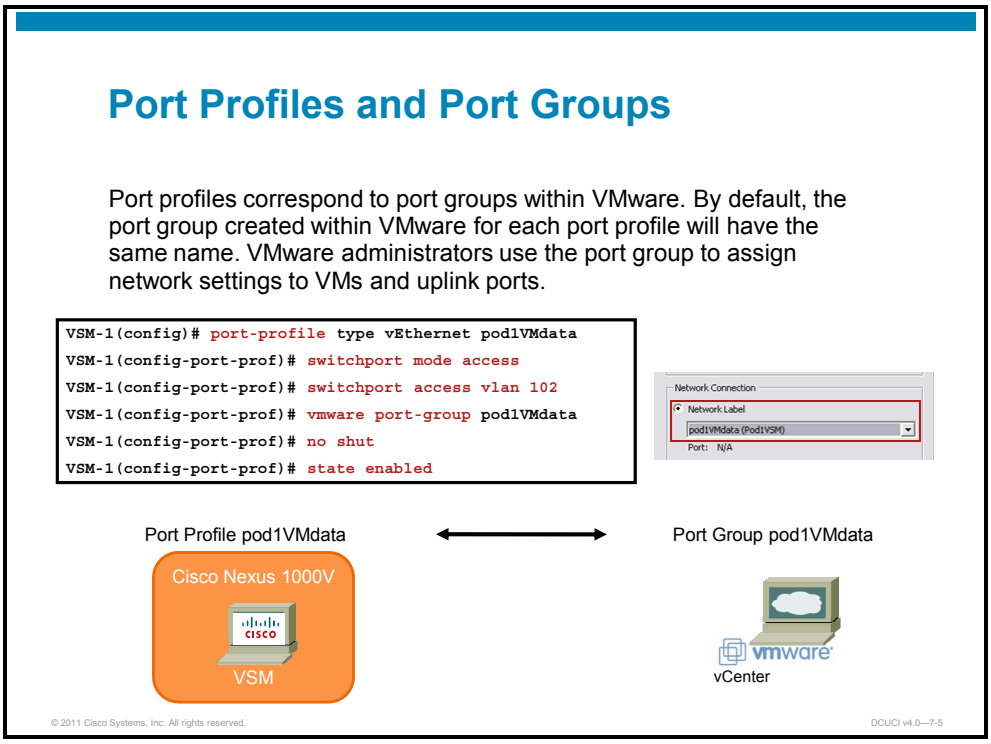

When a port profile is created and enabled, a corresponding port group is created in VMware. By default, this port group has the same name as the profile, but this name is configurable. VMware administrators use the port profile to assign network settings to VMs and uplink ports. When a VMware ESX host port (a physical VM NIC, or VMNIC) is added to a DVS that the Cisco Nexus 1000V switch controls, an available uplink port group is assigned and those settings are applied. When a NIC is added to a VM, an available VM port group is assigned and the network settings that are associated with that profile are inherited.

A NIC in VMware is represented by a VMNIC interface. The VMNIC number is allocated during VMware installation.

Uplink Port Profiles

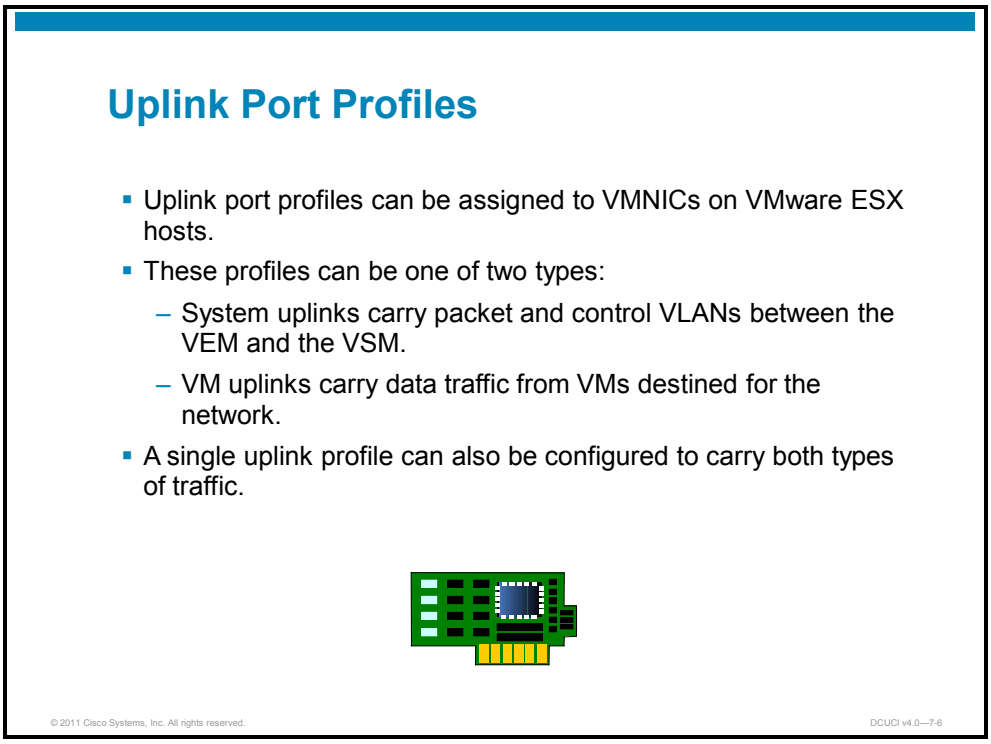

Uplink profiles are used to provide outbound connections from the VEM to the VSM, as well as to carry VM data traffic to the network. Uplink profiles are used to configure VMNICs or physical host ports. These profiles typically contain information such as trunking and port channel behavior.

VM Profiles—Type vEthernet

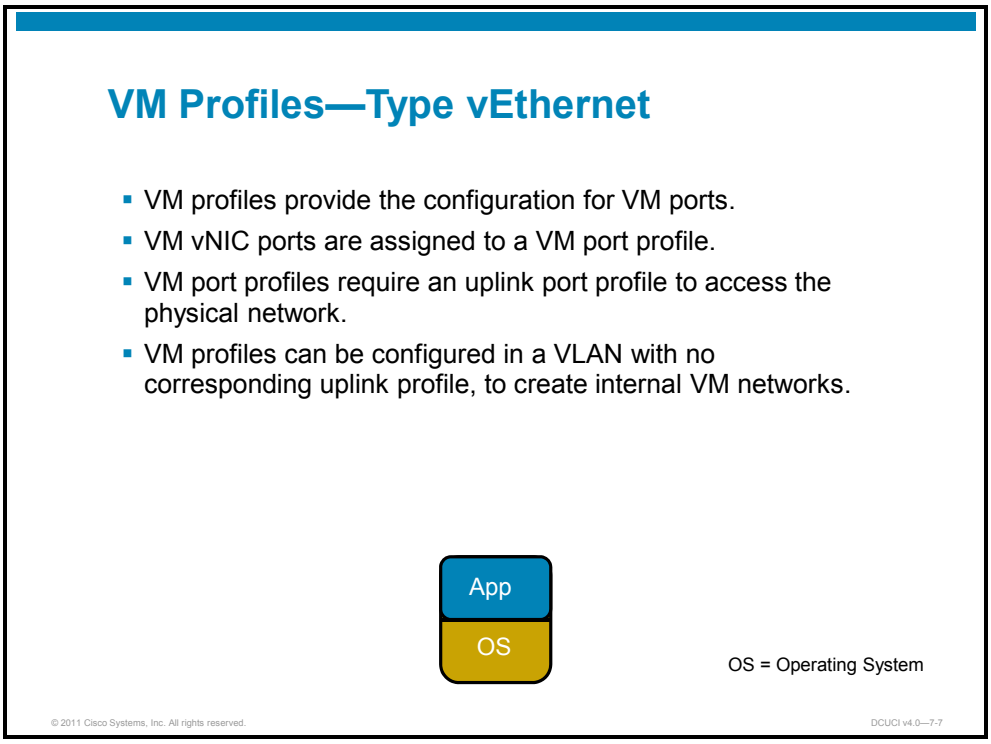

VM profiles are used to provide configuration for VMs. These profiles require an uplink profile to access the physical network.

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

Port Profile States

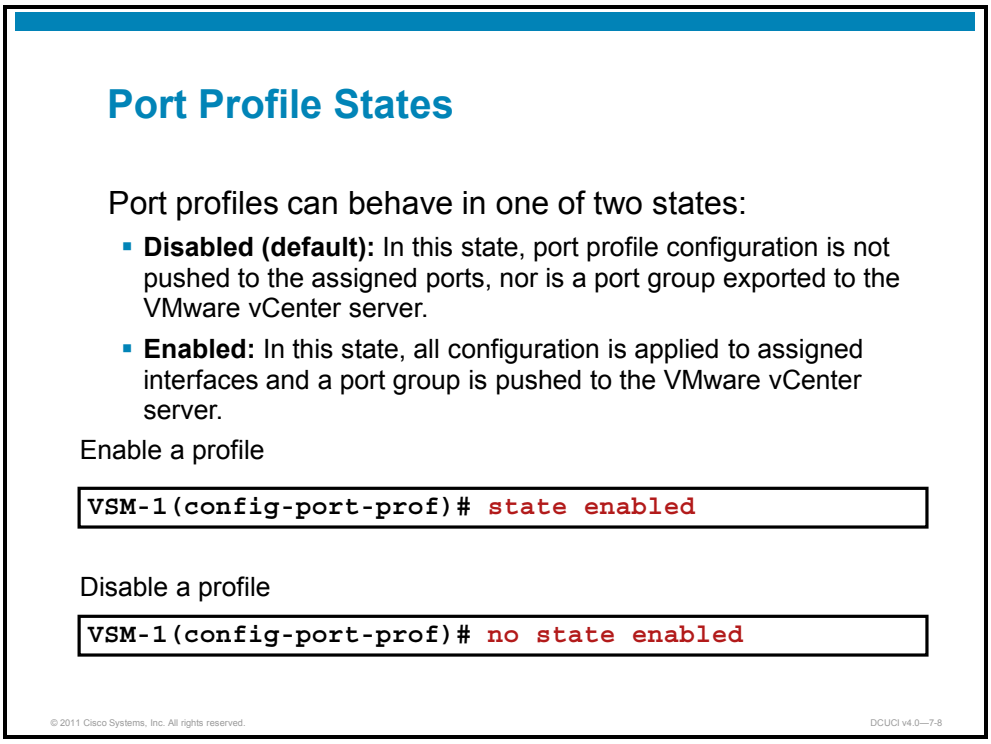

Port profiles can be enabled or disabled as a whole. The configuration of a port profile is applied to assigned ports only if the profile is in an enabled state. Additionally, no VMware port group is created for disabled port groups.

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

Port Profile Usage

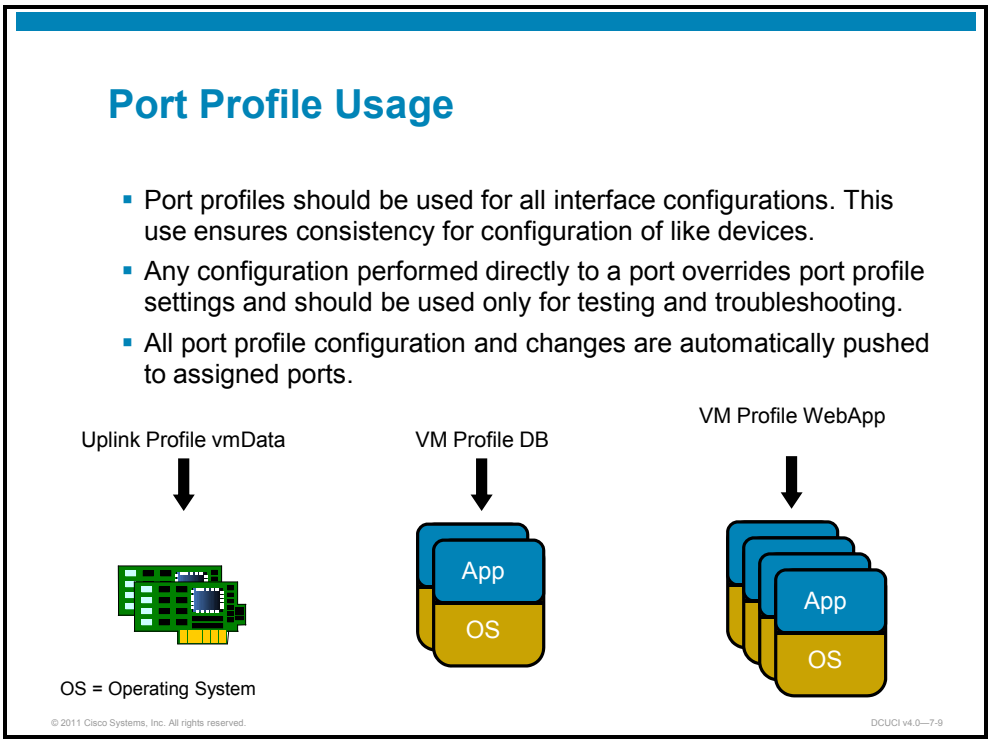

You should use port profiles for all port configurations on the Cisco Nexus 1000V. Using port profiles for port configurations helps to ensure consistency among ports that have similar characteristics and speeds the deployment process. Any configuration that is performed on an individual port overrides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

Port Profile Properties

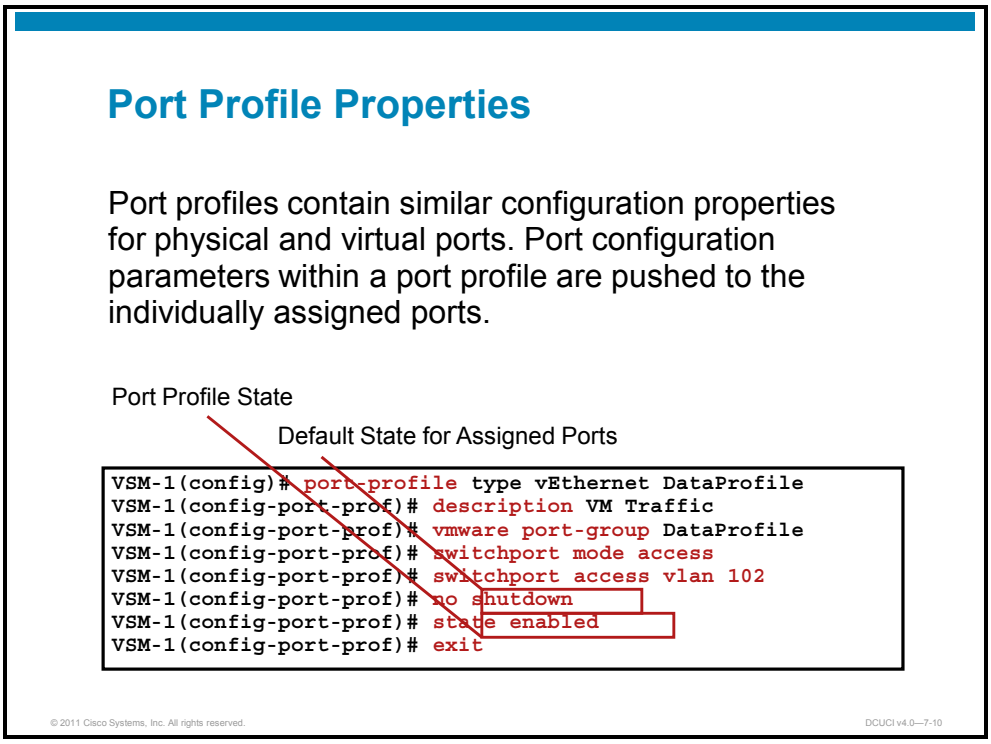

Port profiles are a key element of port configuration on the Cisco Nexus 1000V. VM ports and uplink ports are configured as port profiles. Port profiles are then assigned to virtual machines or virtual switches.

Preventing VM Sprawl

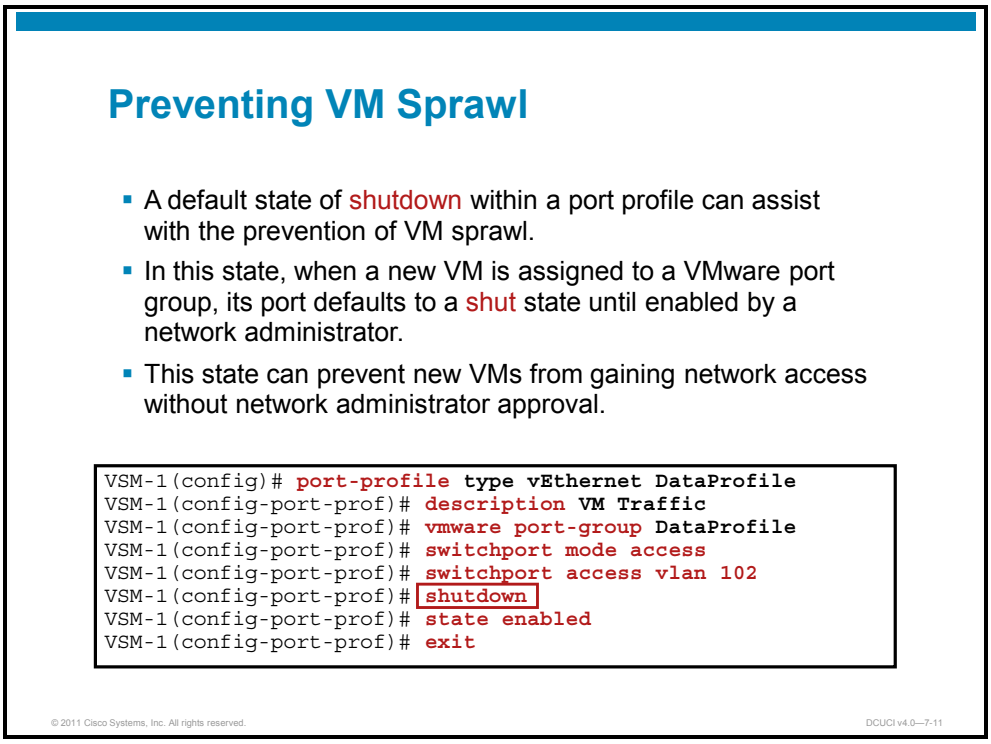

VM sprawl is a new data-center challenge that has come about with virtualization. Virtual servers typically make new server deployment very easy. As such, servers may be deployed more than would be typical in a standard physical architecture. By using the port profiles on the Cisco Nexus 1000V, VM sprawl can be minimized by preventing VMs from accessing physical network resources without network administrator intervention.

By dictating a port profile default port state of down, a network administrator can ensure that new VMs have access to the network only with 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 exists but has no outbound network connectivity.

VLAN Configuration

This topic discusses configuring VLANs.

VLAN Configuration

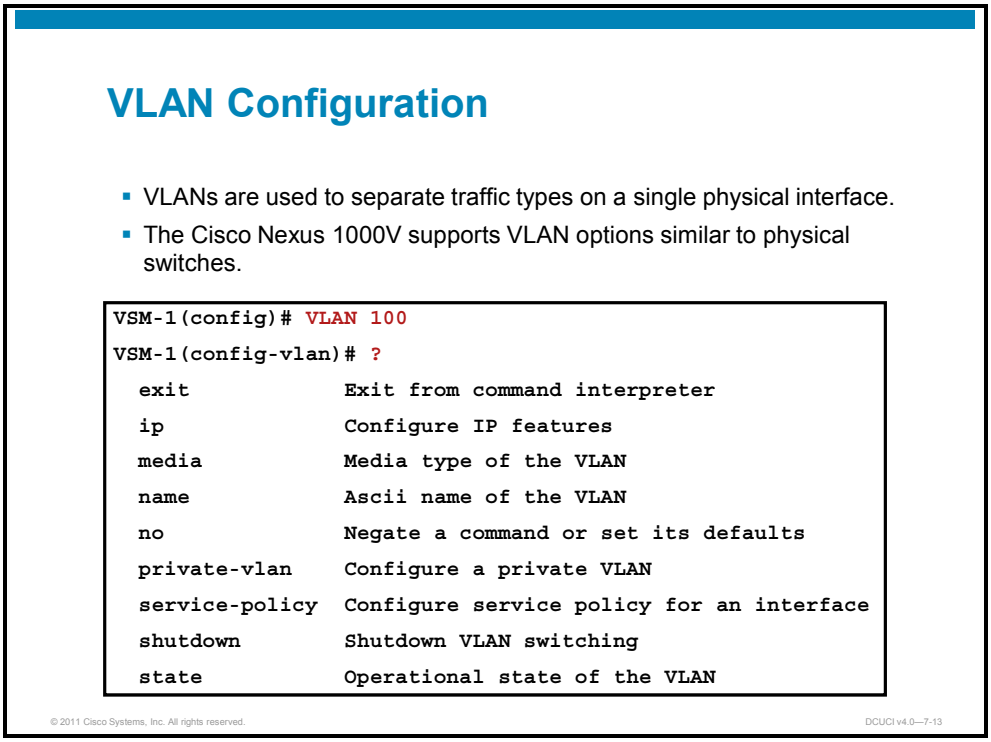

The VLAN command is used to define a VLAN on the Cisco Nexus 1000V Virtual Supervisor Module in a manner similar to VLAN creation on a physical switch. When in VLAN configuration mode, additional options such as private VLANs, an IP address for a switched virtual interface (SVI), or operational state can be configured.

VLAN Port Profile Configuration

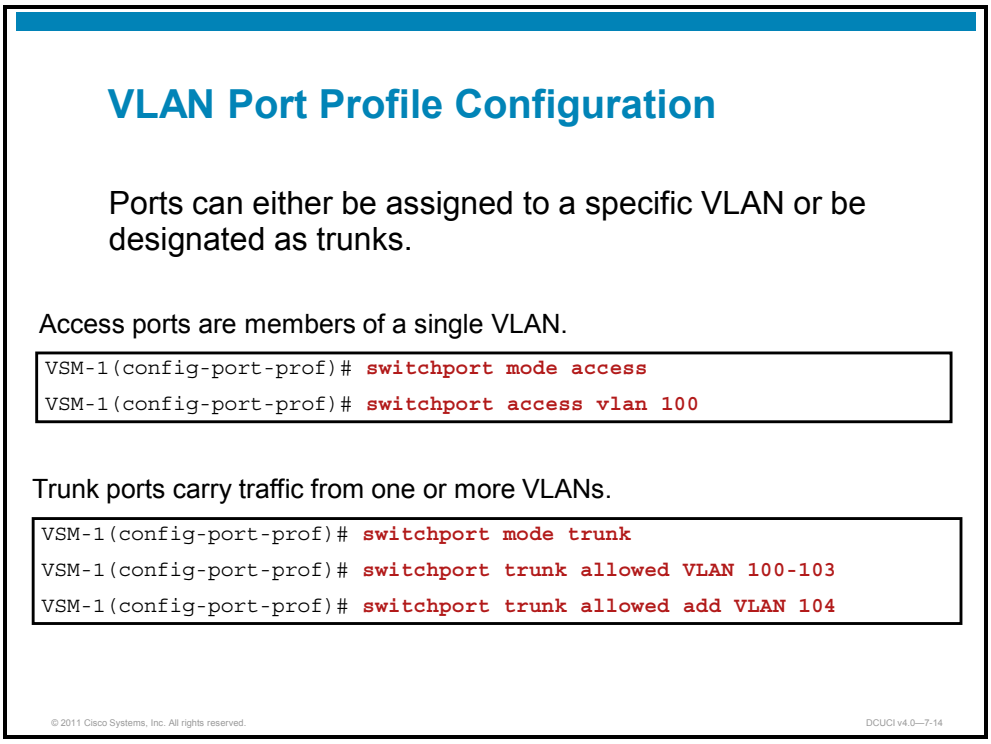

After you configure VLANs on the switch, you can assign port profiles to use them. Ports inherit VLAN settings from the port profile. Two port modes are typically used for VLANs:

- **Trunk ports:** Trunk ports carry traffic for multiple VLANs. By default, these ports carry data for all VLANs, but the list of allowed VLANs 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 of which they are members. In most cases, single server instances (one operating system and one application) running on bare metal or virtual hardware are assigned to access ports.

Create VLANs on the VSM

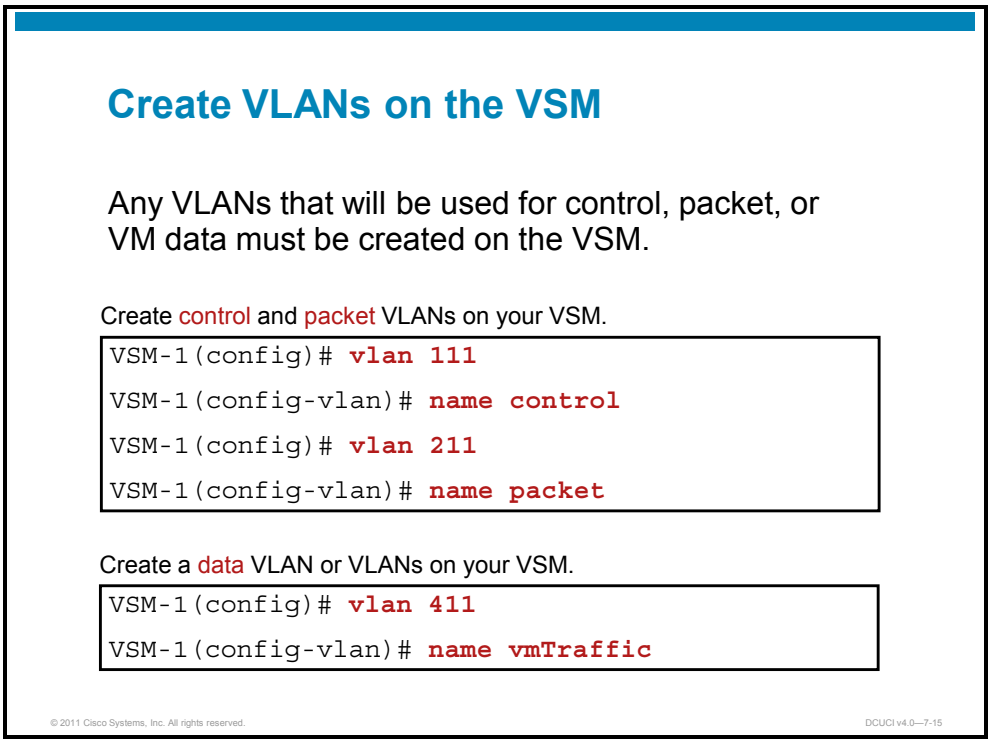

Because the Cisco Nexus 1000V acts as a fully functional switch within the virtual environment, any VLANs that the attached ports use need to be created on the Cisco Nexus 1000V, to facilitate switching for those ports. VLAN creation is done from the configuration level of the Cisco Nexus Operating System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

Private VLAN Configuration

This topic discusses configuring private VLANs.

Create Private VLANs and Assign Them to a Port Profile

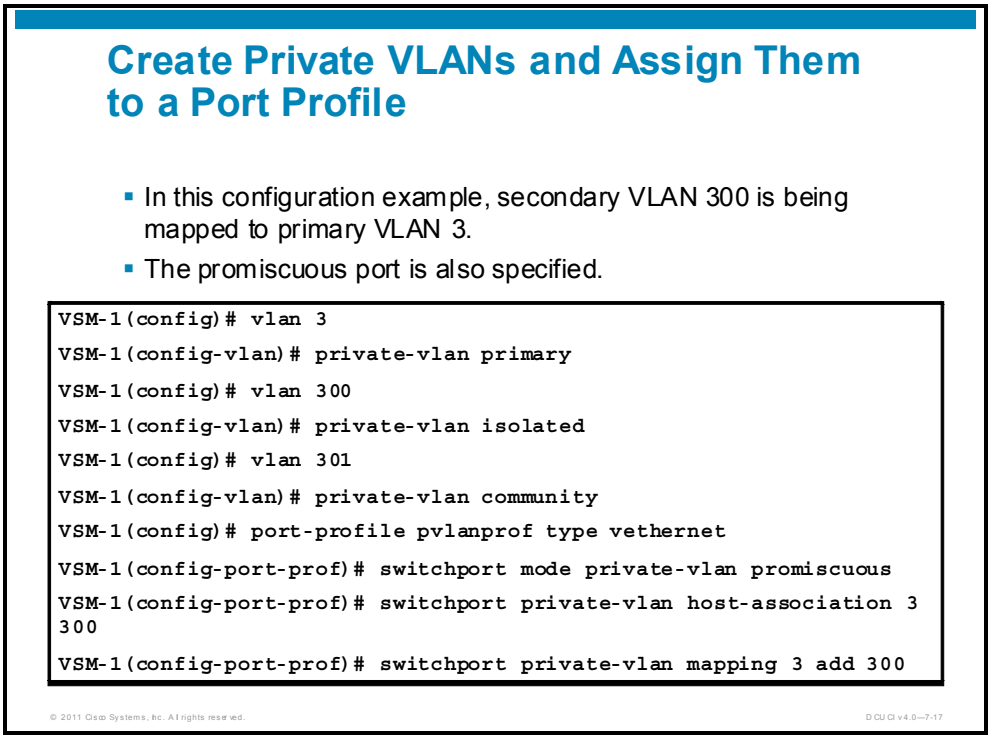

In the configuration example that the figure shows, the secondary VLANs 300, 301, and 302 are associated with the primary VLAN 3. In addition, the port profile specifies a promiscuous port, which will manage traffic flow into and out of the private VLAN association.

The configuration has been applied to the port profile named pvlanprof and may be assigned to vNIC interfaces.

Validate the Private VLAN Port Profile

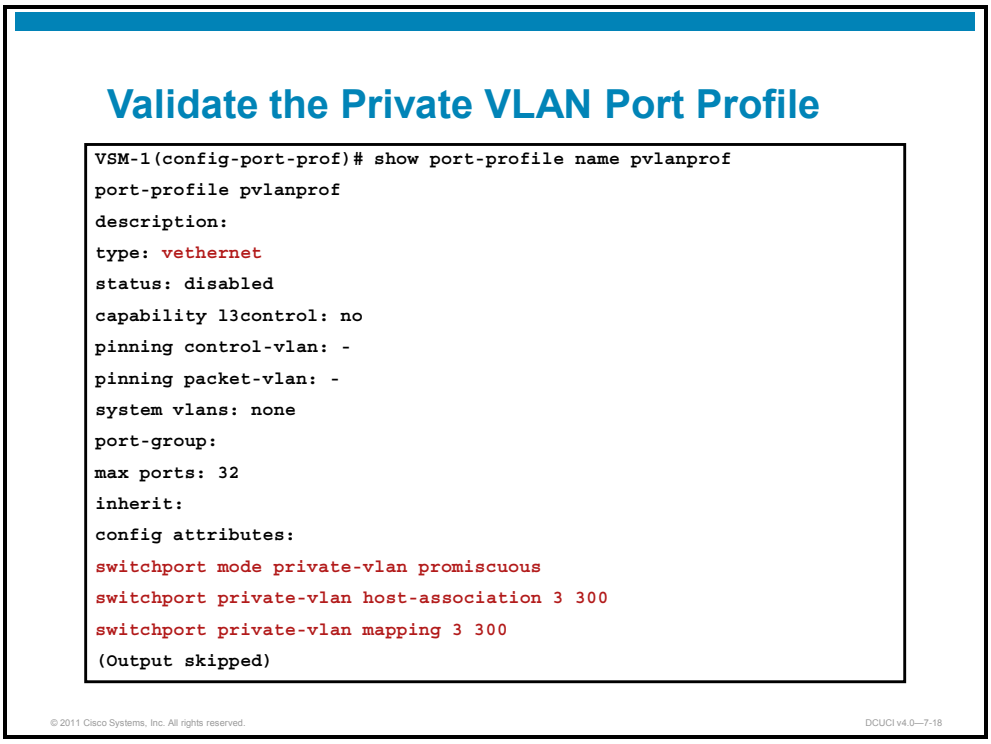

The output of the **show port-profile name pvlanprof** command can be used to validate the configuration and to assure correct secondary-to-primary VLAN association and the promiscuous port assignment.

The output indicates that the pvlanprof port profile is the primary (promiscuous) VLAN number 3. VLANs 300 and 301 are secondary VLANs. Depending on the port profile that is created for the secondary VLANs, they can be isolated or community ports.

Creating Uplink Profiles

This topic discusses how to create uplink profiles.

Create a System Uplink

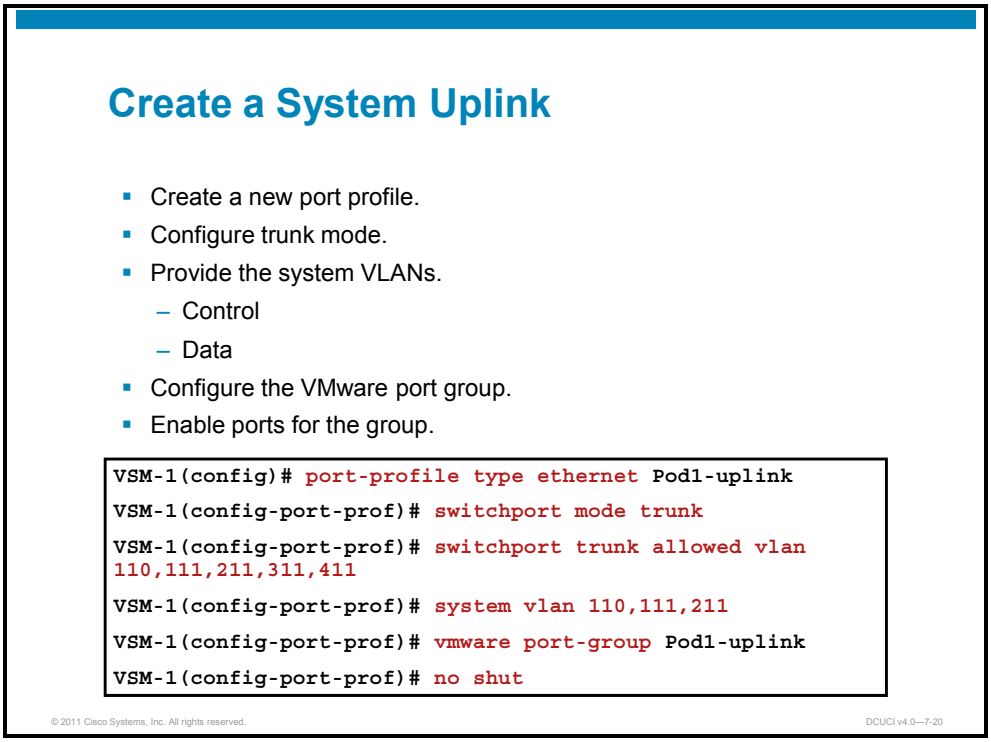

The figure shows how to create a port profile for uplink ports.

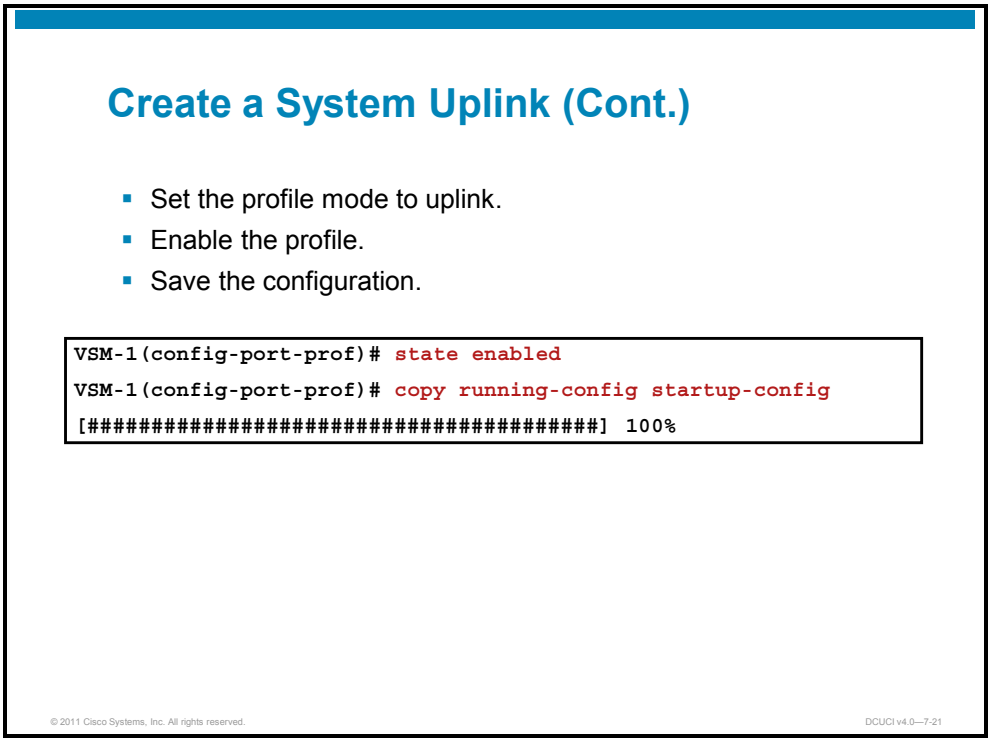

The figure shows the next steps in creating a port profile for uplink ports.

Verify the System Uplink

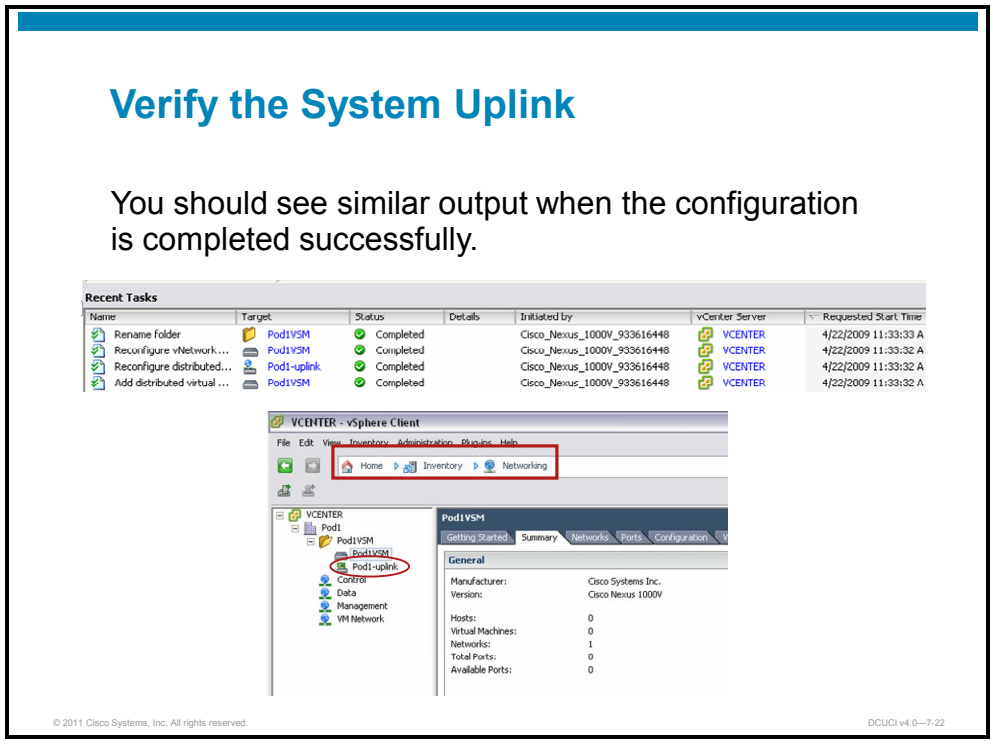

The figure shows how to verify the port profile configuration within vSphere.

Create Uplink Profiles for VM Data

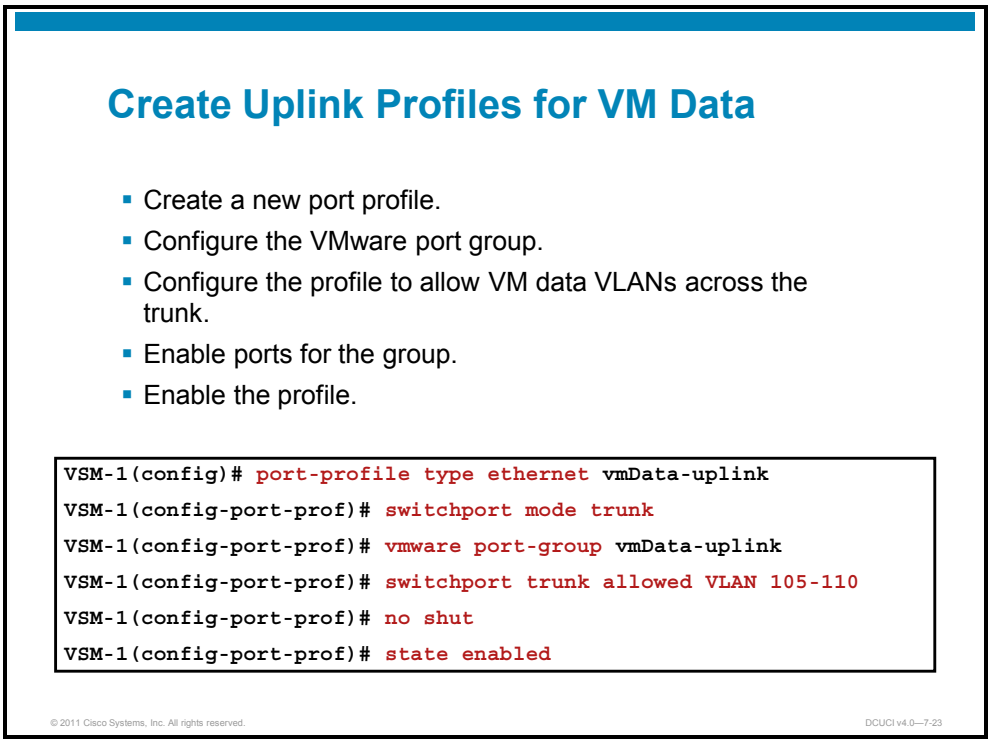

The figure describes how to create a VM data uplink profile.

Creating vEthernet Port Profiles

This topic discusses how to create vEthernet port profiles.

Create a VM Data Port Profile

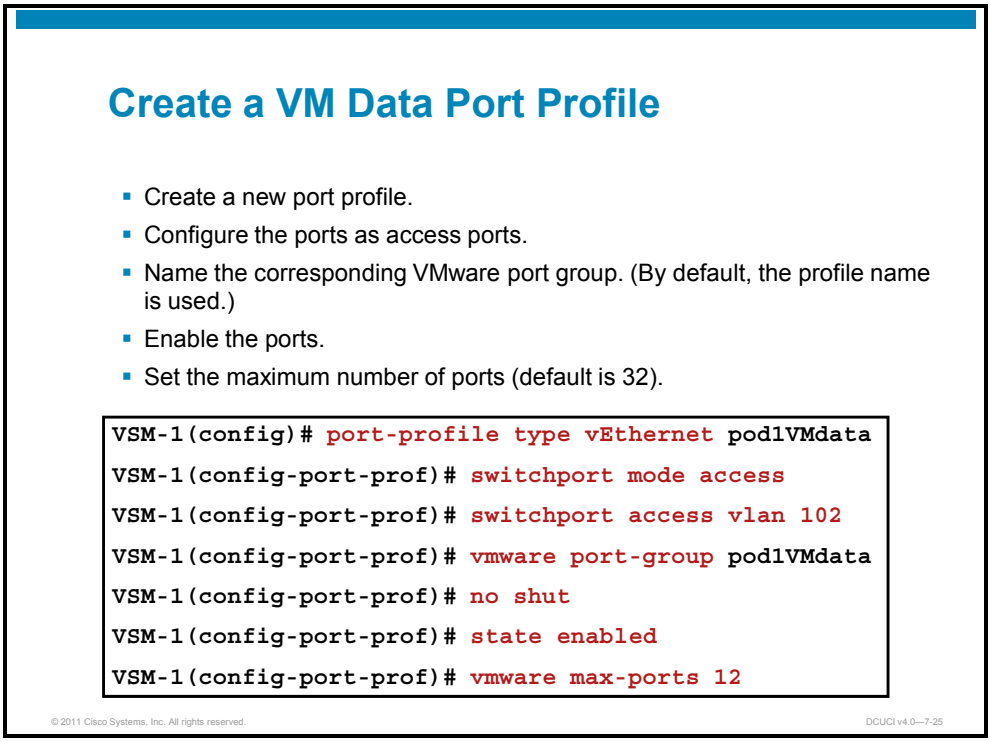

The figure describes the process of creating a VMware data port profile.

Verify the New Port Group

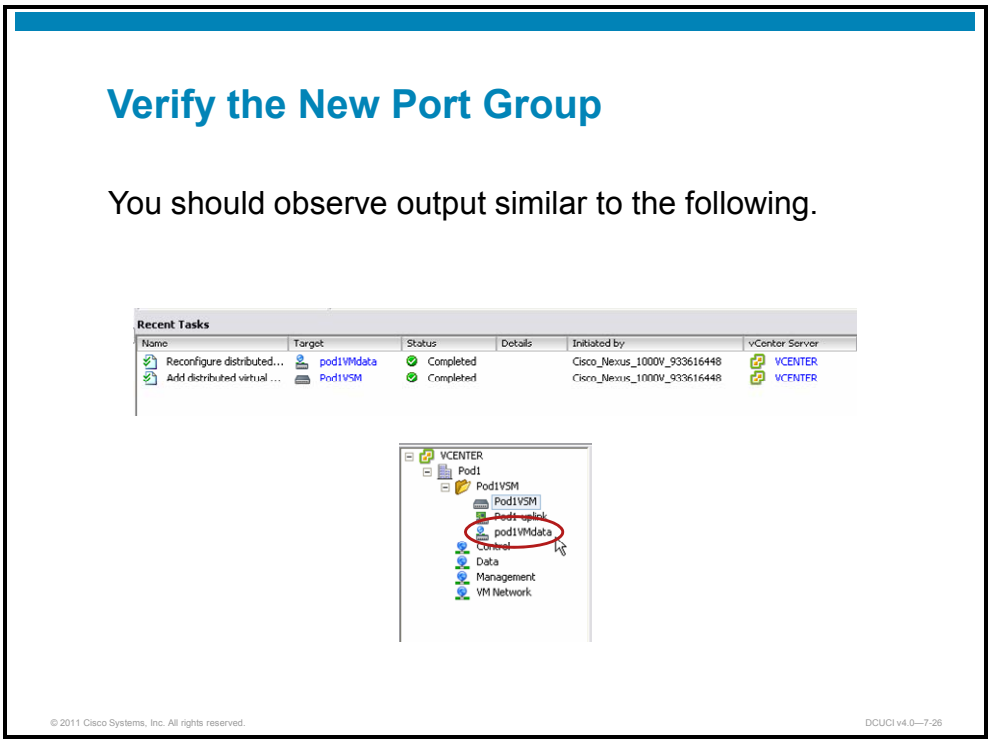

The figure shows the process of verifying VMware data port profile configuration. From within vSphere, choose **Inventory** > **Networking** within the navigation pane. The network inventory objects appear, including the newly created port profile named pod1VMdata. The vSphere Recent Tasks window shows that the creation of the new port profile has been completed successfully.

Configuring Cisco Nexus 1000V Port Channels

This topic discusses how to configure Cisco Nexus 1000V port channels.

Understanding Port Channels

Understanding Port Channels

Port channels are used to create link bundles.

- Port channels act as a single link with bandwidth equal to the combined bandwidth of the member links.
- Load balancing is performed on available links within a port channel.

(2) 1 Gigabit Ethernet Links =
(1) 2 Gigabit Ethernet Port Channel

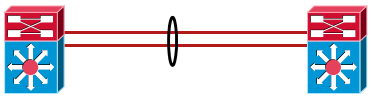

Port channels are used to bundle links, to increase availability. The Cisco Nexus 1000V uses VMNICs to create port channels to upstream switches. These port channels allow VM data that is destined for the physical network to be load-balanced actively across available links. Without port channels, NIC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

Port Channel Load Balancing

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 by 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

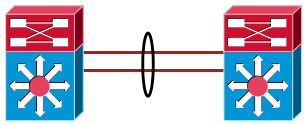

Port channels create a single logical link that comprises multiple physical links. When a port channel has been selected as an egress interface, the switch must choose a physical link, within the port channel, on which to transmit data. The process of selecting a physical link is based on load-balancing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

Port Channel Guidelines

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 these conditions:
  - The uplinks from the host are going to the 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 by using a port profile.

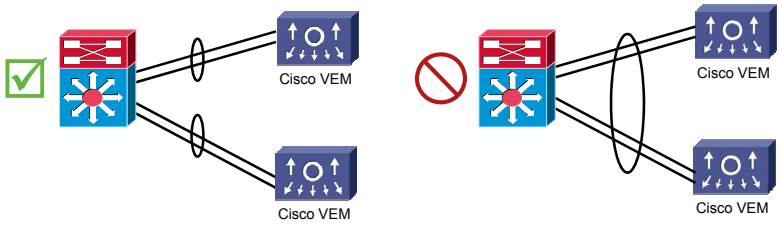

The Cisco Nexus 1000V does not support port channels across multiple hosts. Each VEM on each host requires its own port channel. These port channels can exist between the VEM and a single upstream switch or between the VEM and two upstream switches. Multiple port channels are supported on a single VEM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

Port Channel Types

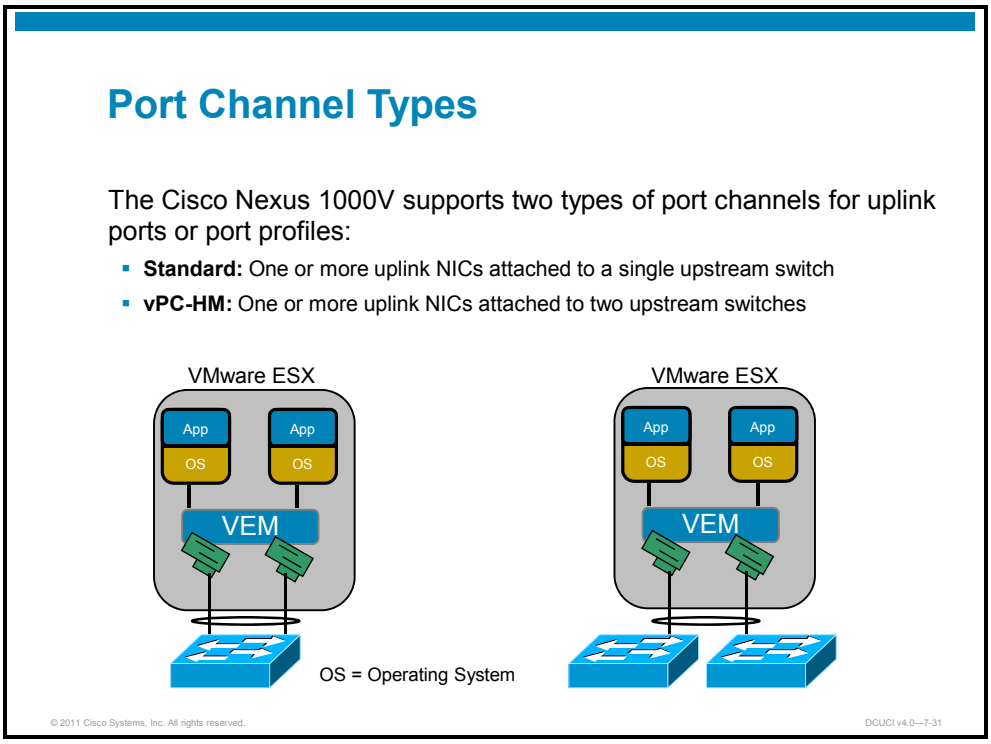

The Cisco Nexus 1000V supports two types of port channels: standard, and virtual port channel host mode (vPC-HM).

Standard 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.

In vPC-HM, VMNIC port channel members can be cabled to two upstream switches, for redundancy. This mode has the following requirements:

- The uplinks from the host go to multiple upstream switches, but the load-balancing algorithm is source MAC-based.
- The upstream switch can and does enable Cisco Discovery Protocol.
- The Cisco Nexus 1000V port channel is configured as enabled for Cisco Discovery Protocol.

Cisco Discovery Protocol information allows the Cisco Nexus 1000V to address issues that typically occur in asymmetric configurations:

- **MAC address flapping:** The Cisco Nexus 1000V prevents MAC address flapping, by using MAC-based load balancing.
- **Duplicate packets that are 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 accepts only matching packets.
- **Packets that are 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.

vPC Host Mode

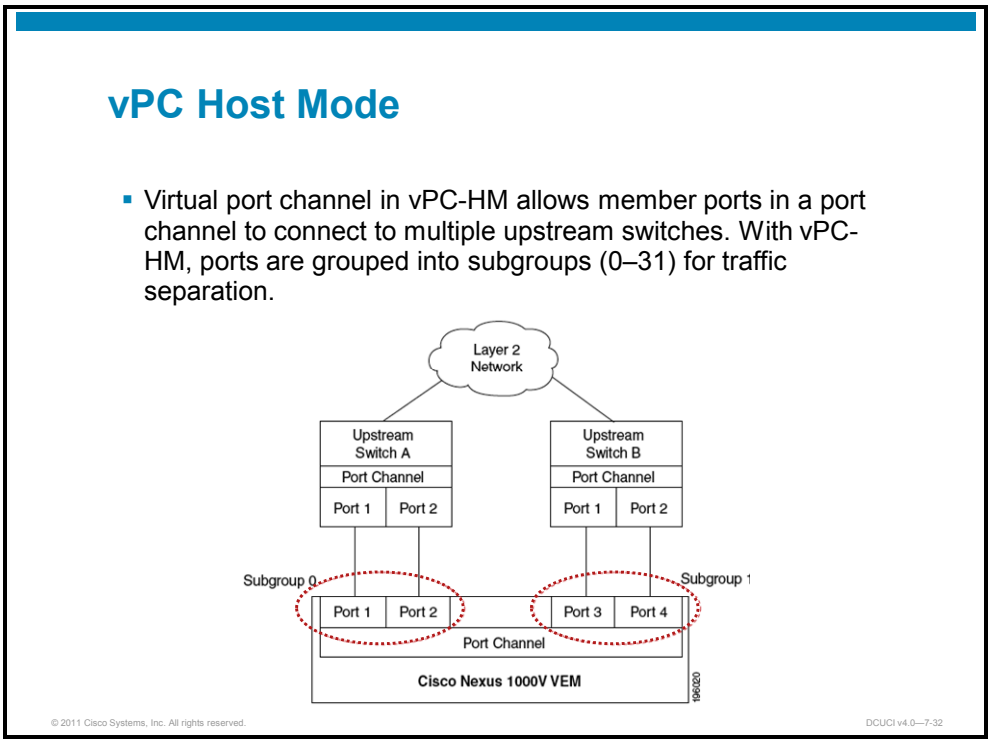

Before beginning this procedure, you must confirm or complete these tasks:

- You are logged in to the CLI in EXEC mode.
- In vPC-HM, the port channel member ports connect to multiple upstream switches, and the traffic must be managed in separate subgroups.
- When you create a port channel, an associated channel group is automatically created.
- vPC-HM is supported only in port channels that are configured in the on mode. vPC-HM is not supported for Link Aggregation Control Protocol (LACP) channels that use the active and passive modes.
- You need to know whether Cisco Discovery Protocol is configured in the upstream switches. If configured, then Cisco Discovery Protocol creates a subgroup for each upstream switch to manage its traffic separately. If Cisco Discovery Protocol is not configured, then you must manually configure subgroups to manage the traffic flow on the separate switches.
- If you are using Cisco Discovery Protocol with the default Cisco Discovery Protocol timer (60 seconds), then links that advertise that they are in service and then out of service in quick succession can take as much as 60 seconds to be returned to service.
- If any subgroup has more than one member port, then you must configure a port channel for the member ports of each subgroup on the upstream switch.
- If vPC-HM is not configured when port channels connect to multiple upstream switches, then the VMs behind the Cisco Nexus 1000V receive duplicate packets from the network for unknown unicast floods, multicast floods, and broadcasts.
- The subgroup command that is used in this procedure overrides any subgroup configuration that is specified in the port profile that the port channel interface inherits.

Standard Port Channel Configuration

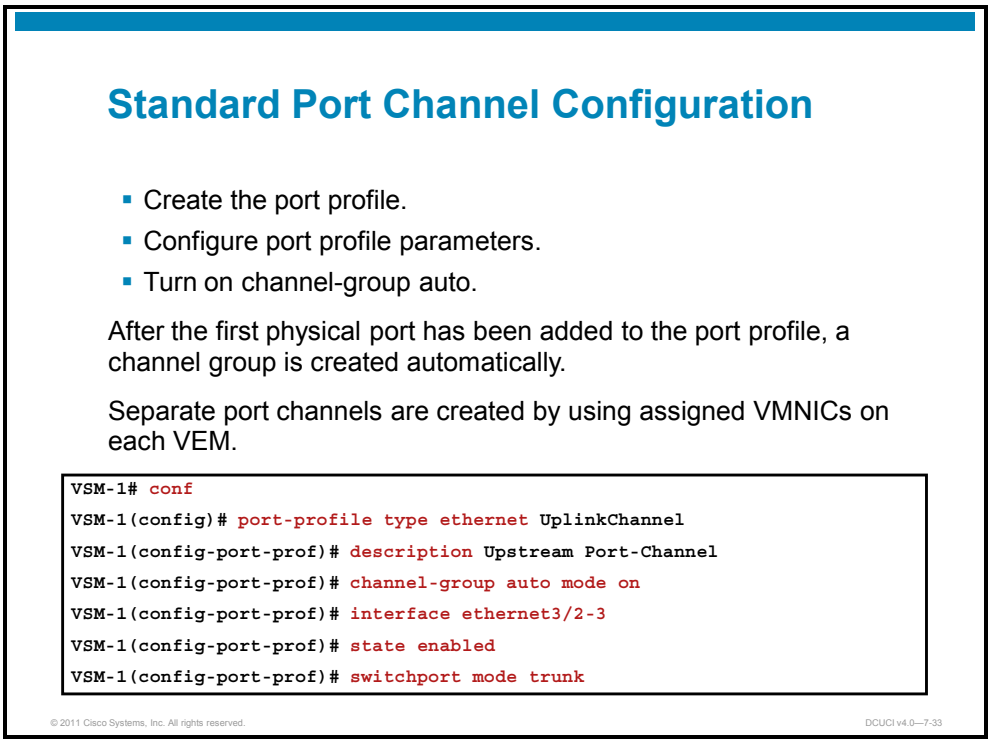

Port channels should be configured by using port profiles rather than at the individual port level. This method ensures consistent configuration across hosts and compatibility for advanced functionality, such as VMware vMotion, high availability, and Distributed Resource Scheduler (DRS).

For standard port channels, enter the **channel-group auto mode on** command under the port profile configuration. When VMNICs (physical NICs on ESX hosts) are added to this profile, port channels are automatically created on each host, and additional VMNICs are added to the channel. A port channel cannot span multiple hosts, so an individual port channel is created from the VEM of each ESX host.

Standard Port Channel Verification

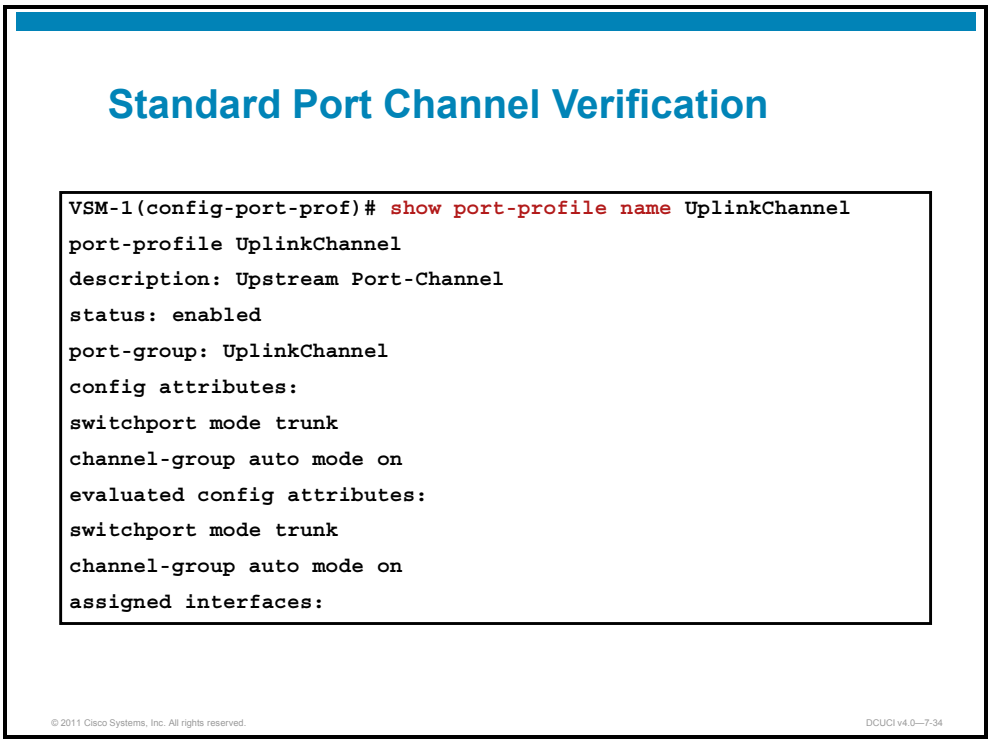

The figure illustrates the verification of the uplink port profile standard port channel configuration.

vPC-HM Port Channel Verification

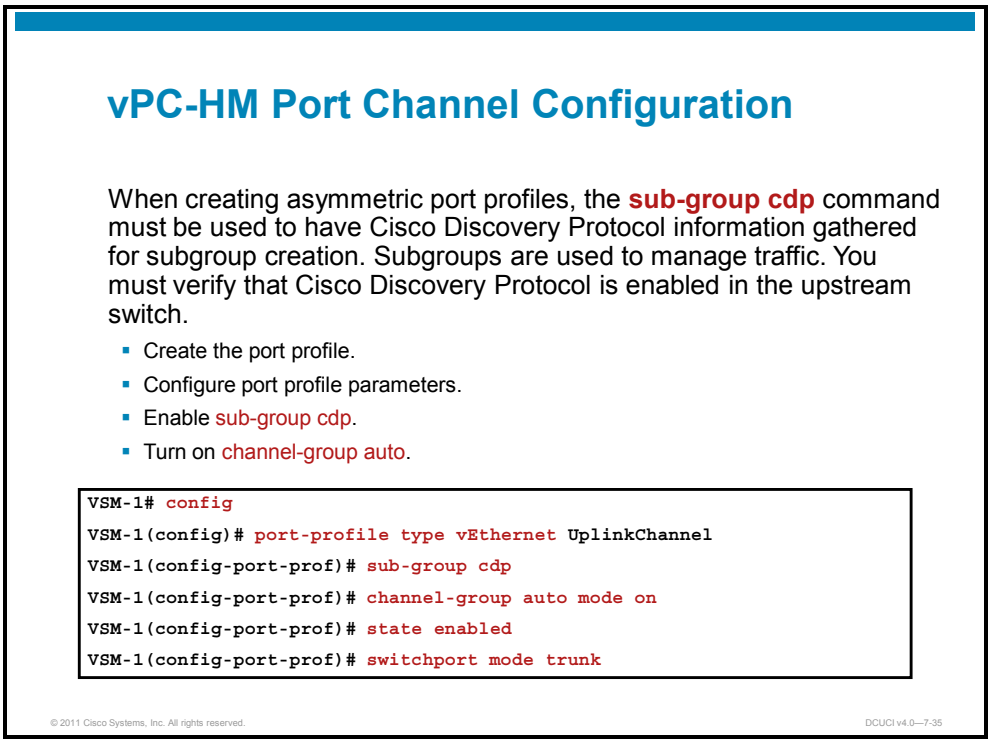

vPC-HM port channels rely on source MAC address load balancing and subgroups, to prevent MAC address instability and multiple frame copies on upstream switches. Loops are prevented automatically because frames that are received on uplink ports are not forwarded to other uplink ports.

- **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 always uses the same uplink port in asymmetric configurations. This method prevents upstream switches from having issues that can be caused when a source MAC is seen on multiple switches and from having to make constant MAC table changes.
- **Subgroups:** Using Cisco Discovery Protocol, the Cisco Nexus 1000V can analyze traffic and create subgroups for the separate upstream switches in a vPC-HM port channel. This method allows the Cisco Nexus 1000V to prevent multiple copies of the same frame that are received on separate ports from being forwarded to a VM.

Adding VEMs to the VSM

This topic discusses adding VEMs to the VSM.

Add a Host

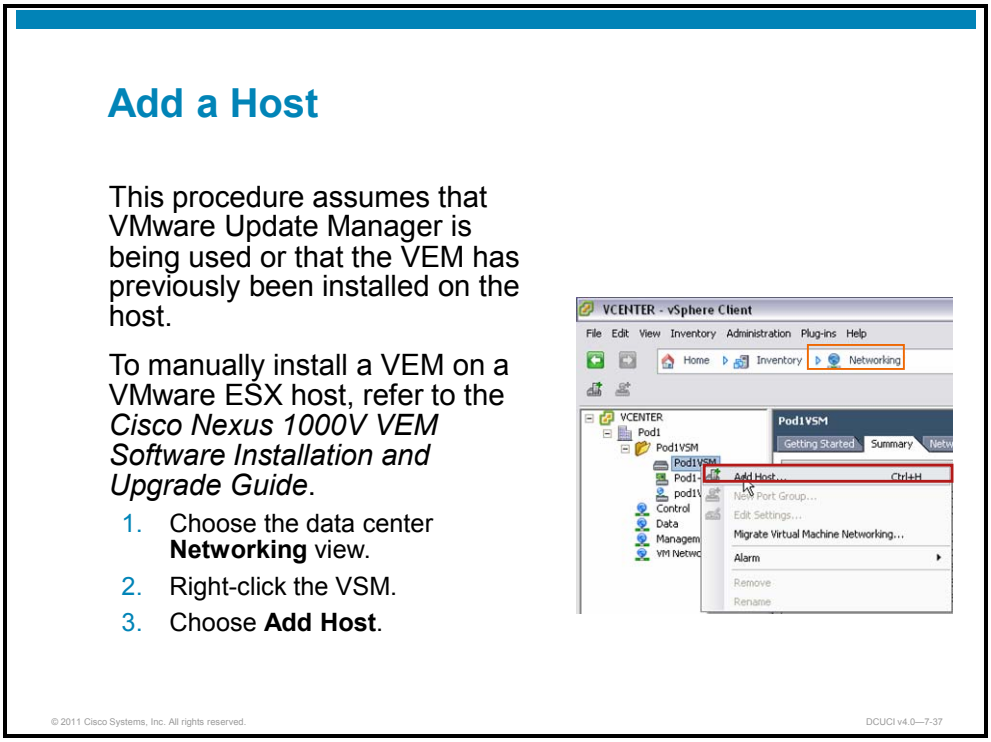

The VMware administrator adds hosts to a VSM. The administrator uses the data-center Networking view to assign hosts and their VMNICs to a Cisco Nexus 1000V DVS.

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

If you need to manually install a VEM on an ESX host, refer to the *Cisco Nexus 1000V* VEM *Software Installation and Upgrade Guide*.

Add a Host (Cont.)

4. Choose the host to add.
5. Choose the uplink NIC or NICs to add.
6. Click **Next**.

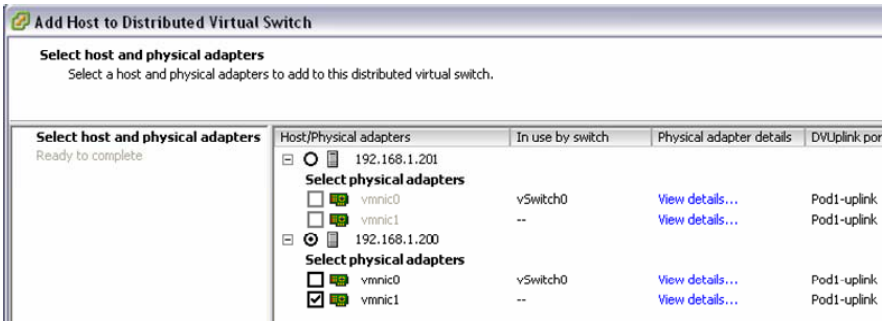

On the Add Host to Distributed Virtual Switch screen, follow these steps:

- Choose the host that you want to add.
- Choose the uplink NICs that you want to add.
- Click **Next**.

Add a Host (Cont.)

7. Verify the configuration.
8. Click **Finish**.
9. Repeat as necessary for all hosts.

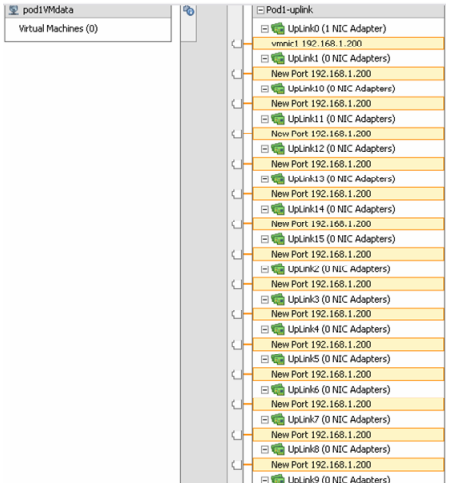

After you have added the host, verify the configuration, then click **Finish**. Repeat these steps for all the hosts that need to be added.

Verify the New VEMs from the VSM

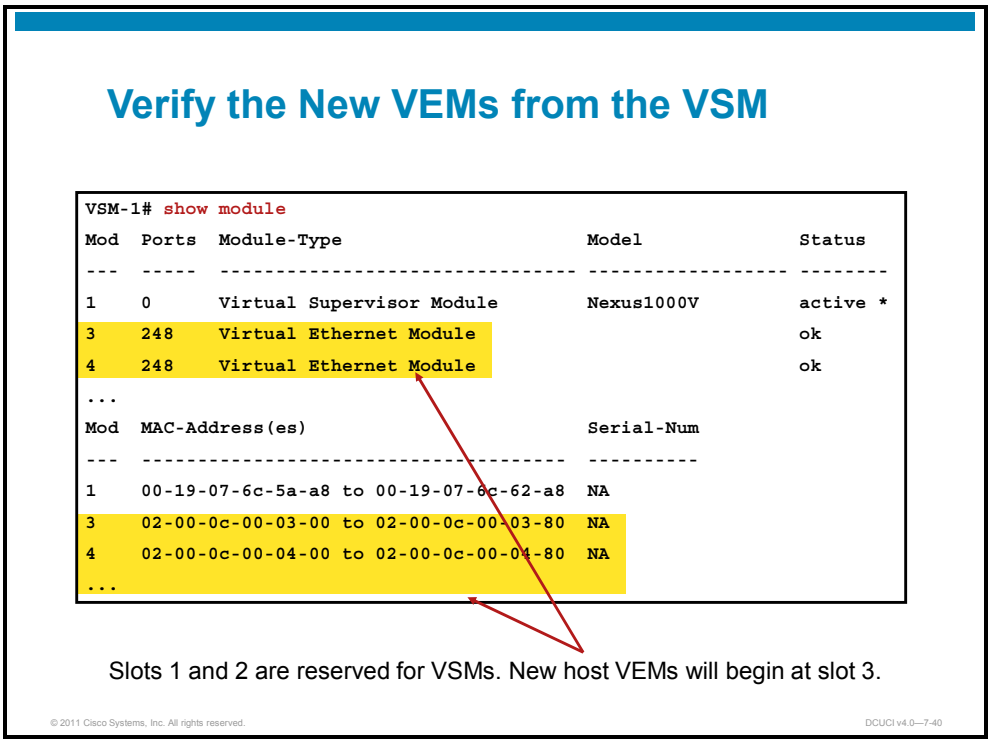

After a host has been added and the VEM has been installed successfully, the VEM appears as a module on the VSM CLI, like modules that are added to a physical chassis.

Note	Slots 1 and 2 are reserved for the VSMs. New host VEMs start from slot 3.

Verify the New VEMs from the VSM (Cont.)

The **show module vem map** command shows the status of all VEMs, as well as the UUID of the host on which the VEM runs.

```
VSM-1# show module vem map
Mod      Status           UUID
---      -----------      ----
3       powered-up      34343937-3638-3355-5630-393037415833
4       powered-up      34343937-3638-3355-5630-393037415834
```

This is the UUID of the host on which the VEM resides.

The **show module vem map** command shows the status of all VEMs, as well as the Universally Unique Identifier (UUID) of the host on which the VEM runs. This command can be used for verification purposes.

Add a VM to a VSM Port Group

Add a VM to a VSM Port Group

1. Right-click the desired VM.
2. Choose **Edit Settings**.

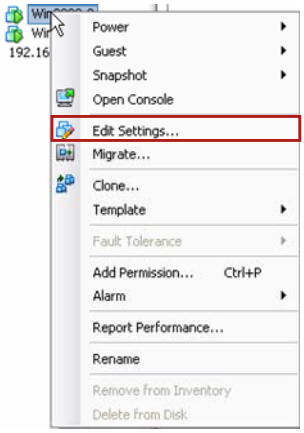

To add a VM to a VSM port group, right-click the VM, then choose **Edit Settings**.

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

3. Select the NIC to configure.
4. Change the Network Label to the desired VSM port profile.

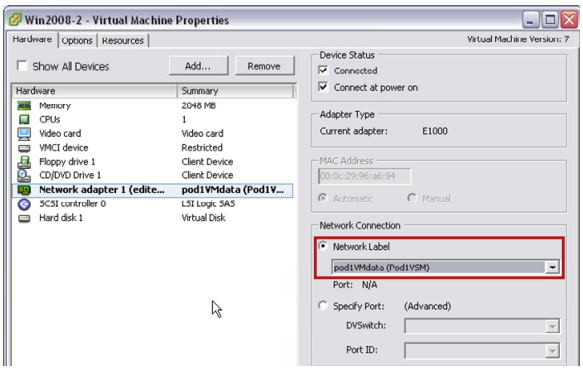

The VMware administrator adds VMs to the Cisco Nexus 1000V DVSs. The port group becomes available for selection as a network label within the VM configuration, after the port profile has been created and the corresponding port group has been pushed to vCenter.

Backing up a VSM Configuration

This topic discusses how to back up a VSM configuration.

Managing VSM Files and Configurations

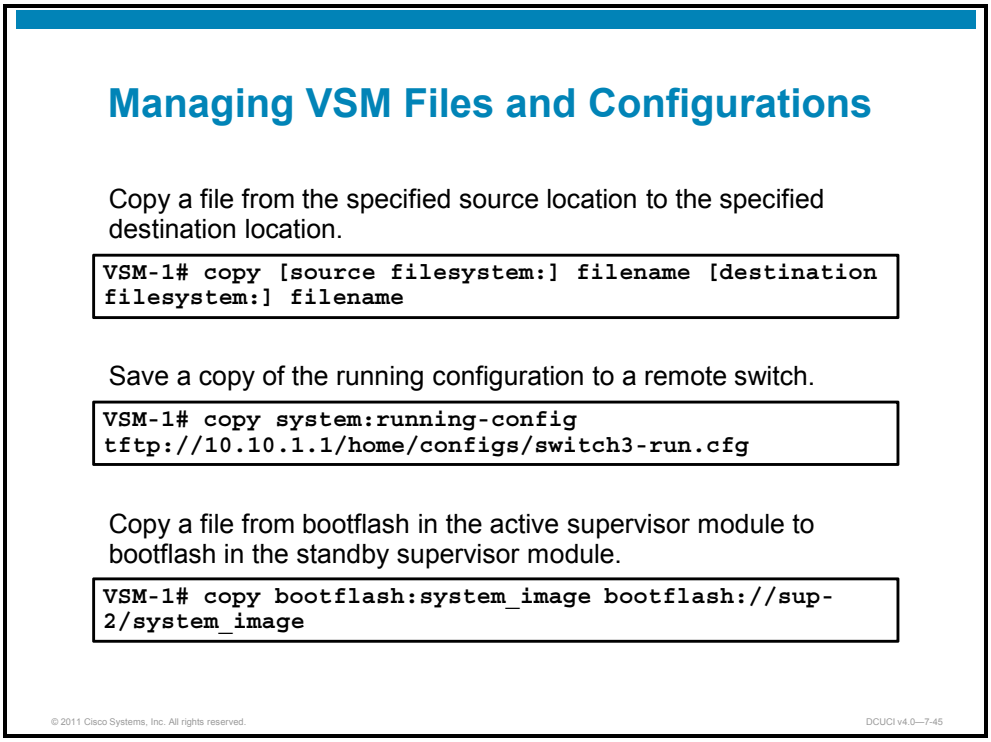

The figure shows the various **copy** commands that can be used within the Cisco Nexus 1000V CLI to manage configuration file images and Cisco NX-OS kickstart and system images. These commands can be issued from EXEC mode.

The first command illustrates the command structure to copy a file from a source to a destination location. The second command saves a copy of the active running configuration to a remote switch. The third command copies a file in the bootflash of the active supervisor to the bootflash of the standby supervisor.

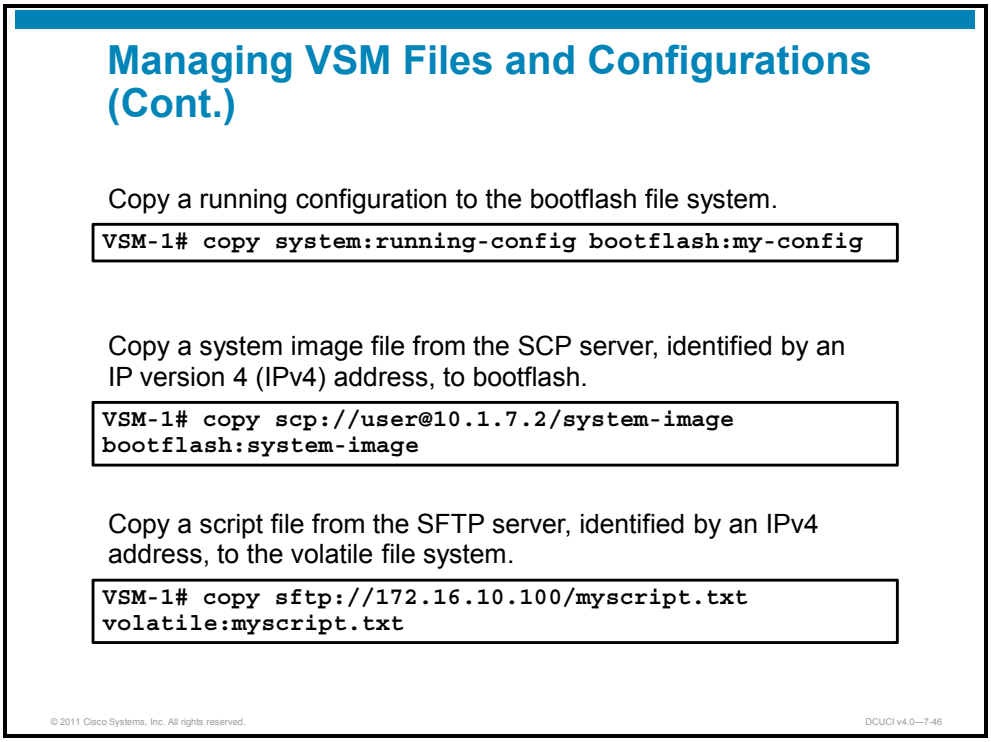

The first command in this figure can be used to copy the active running configuration to the bootflash. The saved file will survive a system reset (write erase and reload command sequence). The second command downloads a copy of the system image file from a Secure Copy Protocol (SCP) server to the bootflash. The third command copies a script file from a Secure File Transfer Protocol (SFTP) server to the volatile file system.

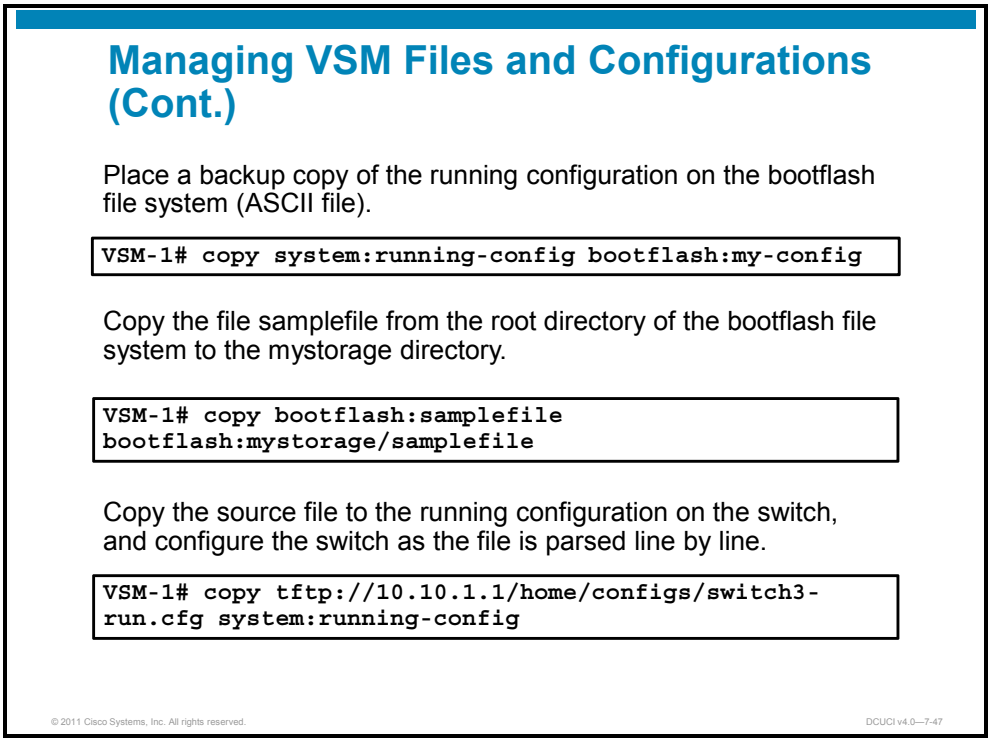

The first command in this figure places a backup copy of the active running configuration on the bootflash. The second command saves a file on the bootflash root directory to a different bootflash directory. The third command copies the source file to the active running configuration and configures the switch, line by line, as the file commands are compiled.

vMotion and Port Profile Mobility

This topic discusses vMotion and port profile mobility.

Policy-Based VM Connectivity

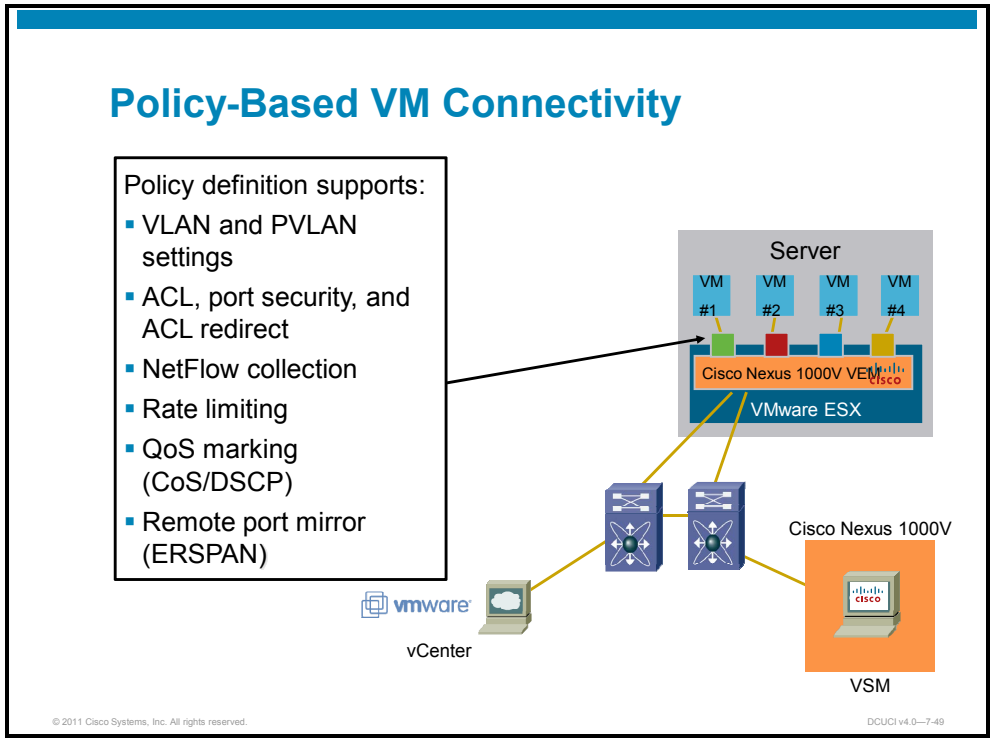

Port profiles are created within the Cisco Nexus 1000V CLI by the network administrator and are pushed out to vCenter. On vCenter, the profiles appear in the vCenter inventory as a port group.

These port profiles represent policies that include VLANs and PVLANs, ACLs, Layer 2 port security, NetFlow collection, rate limiting, quality of service (QoS) marking that uses either the differentiated services code point (DSCP) or class of service (CoS) values, and Encapsulated Remote Shared Port Analyzer (ERSPAN).

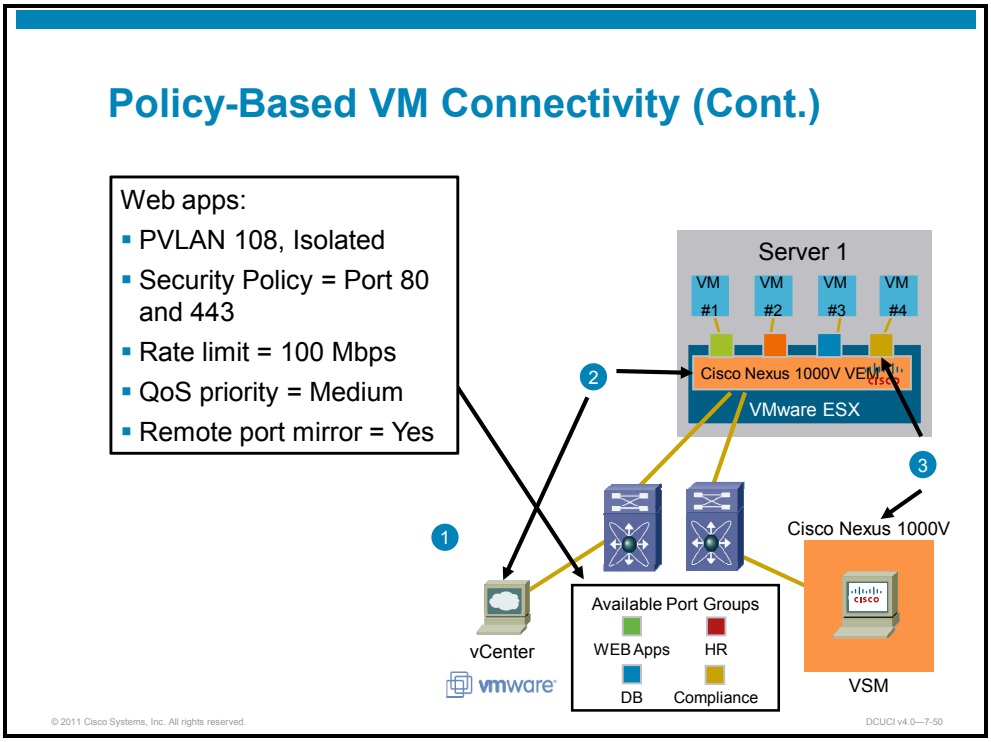

In this configuration, a port profile named Web Apps has been created and assigns a security policy for TCP ports 80 and 443 to the isolated secondary PVLAN 108. The profile rate-limits the port to 100 Mb/s with a QoS medium priority level and enables remote port mirroring.

This policy has been state enabled and pushed out to vCenter, where it appears within the vCenter inventory. The policy has been assigned to a vNIC within VM 1 on server 1.

Mobility of Security and Network Properties

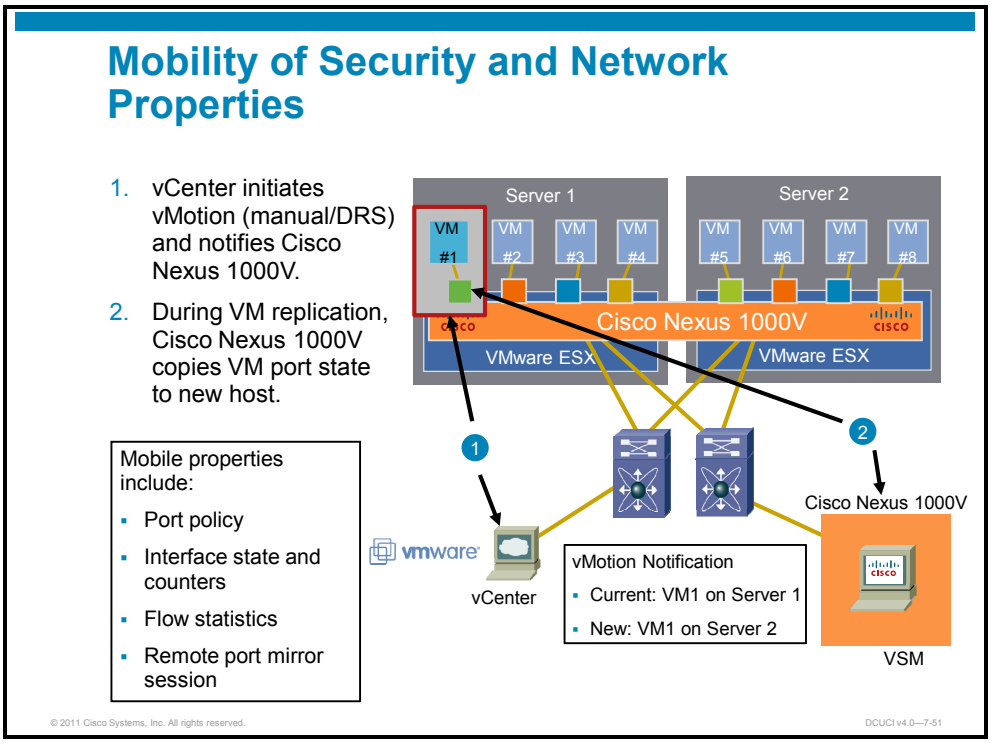

If a manual or automatic vMotion event occurs, the Cisco Nexus 1000V is notified of the process. During VM replication on the secondary host, the primary VSM copies the port state to the new host.

The mobility properties that are retained and copied to the new host include the port policy, interface state and counters, flow statistics, and remote Switched Port Analyzer (SPAN) session.

Mobility of Security and Network Properties (Cont.)

1. vCenter initiates vMotion (manual/DRS) and notifies Cisco Nexus 1000V.
2. During VM replication, Cisco Nexus 1000V copies VM port state to new host.
3. When VMotion completes, port on new ESX host is brought up and VM MAC address is announced to the network.

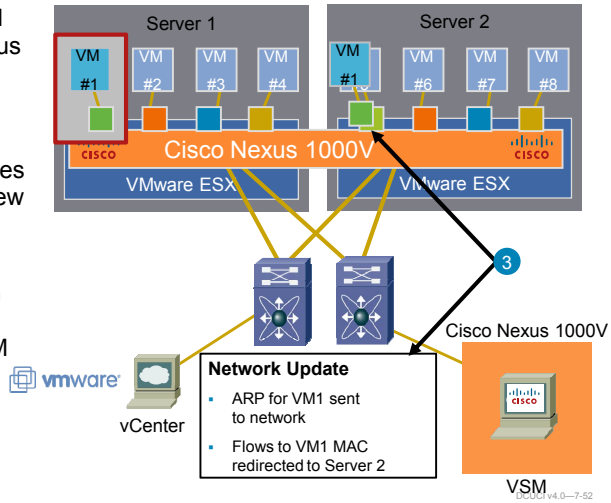

After the server administrator, operating within vSphere, assigns these port profiles to either uplink or vNIC ports, these policies survive both manual and automatic vMotion events.

When the vMotion process is completed, the port on the new ESX host is made active and the VM MAC address is announced to the network.

Summary

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

Summary

- Port profiles are used to configure multiple ports globally.
- Port profiles reduce administrative time and minimize opportunities for error.
- After a port profile has been created, it is passed to vCenter; new port groups will be available for VMs.
- VLANs and PVLANs can be configured within the CLI and associated with a port profile.
- Port channels are used for link aggregation, redundancy, and load balancing.
- The Cisco Nexus 1000V supports two types of port channels: symmetric, from the Cisco Nexus 1000V to a single upstream switch, and asymmetric, from the Cisco Nexus 1000V to two upstream switches.

Summary (Cont.)

- All port configuration for VM and uplink ports should be performed by using port profiles.
- Cisco Nexus 1000V image and configuration files can be managed from the VSM CLI.
- Port profile assignment to a VM assures policy mobility during a manual or automatic vMotion event.

Lesson 6

Configuring Cisco UCS Manager for VMware PTS

Overview

The Cisco M81KR represents a quantum leap forward in performance and management in VMware vSphere environments. Implementers must be able to configure Cisco UCS Manager to integrate the MK81KR into VMware vSphere deployments. This lesson presents the installation and configuration of the Cisco Virtual Interface Card (VIC) and its connectivity to VMware vCenter. The lesson also describes the creation of static and dynamic virtual network interface cards (vNICs) and static virtual host bus adapters (vHBAs) and the creation and assignment of port profiles.

Objectives

Upon completing this lesson, you will be able to describe the installation and configuration of the Cisco VIC and the creation of vNICs, vHBAs, and port profiles. This ability includes being able to meet these objectives:

- Describe how to install the Cisco UCS Manager extension in vCenter
- Describe how to configure Cisco UCS Manager to connect to vCenter
- Describe how to configure uplink and vEthernet profiles
- Describe how to configure service profiles that contain dynamic NICs
- Describe how to configure vMotion hosts and port profile mobility within Cisco M81KR

Install Cisco UCS Manager Extension in vCenter

This topic discusses how to install the Cisco Unified Computing System (UCS) Manager extension in vCenter.

Preparing the Environment

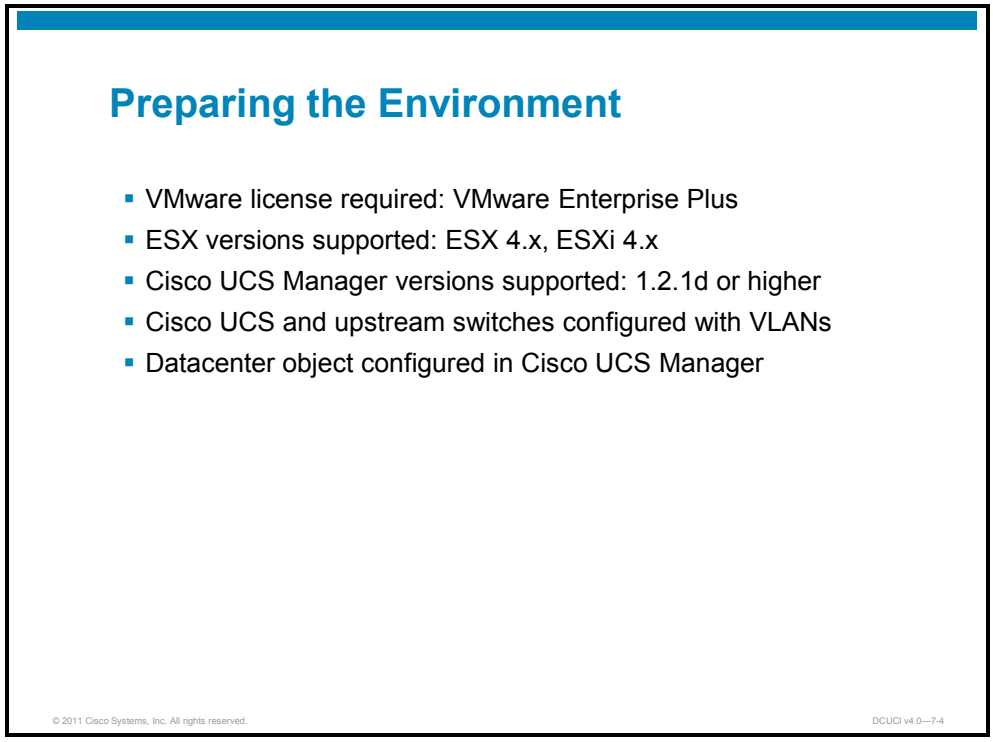

Before integrating Cisco UCS Manager with vCenter, you should verify that a few requirements are in place. VMware vCenter requires an Enterprise Plus license to create a distributed virtual switch (DVS), which is how Cisco Virtual Network Link (VN-Link) in hardware appears in vCenter.

Support for DVSs in vCenter starts in VMware ESX 4.0, and support for the Pass-Through Switch (PTS) DVS starts with ESX 4.0 Update 1. Cisco UCS Manager supports the VIC PTS, starting in Cisco UCS Manager Release 1.2(1d).

Other requirements include configuring your upstream network switches with the proper VLANs that you want to expose to your virtual machines (VMs). You must also configure a datacenter object in vCenter to contain the PTS DVS.

vCenter Integration Methods

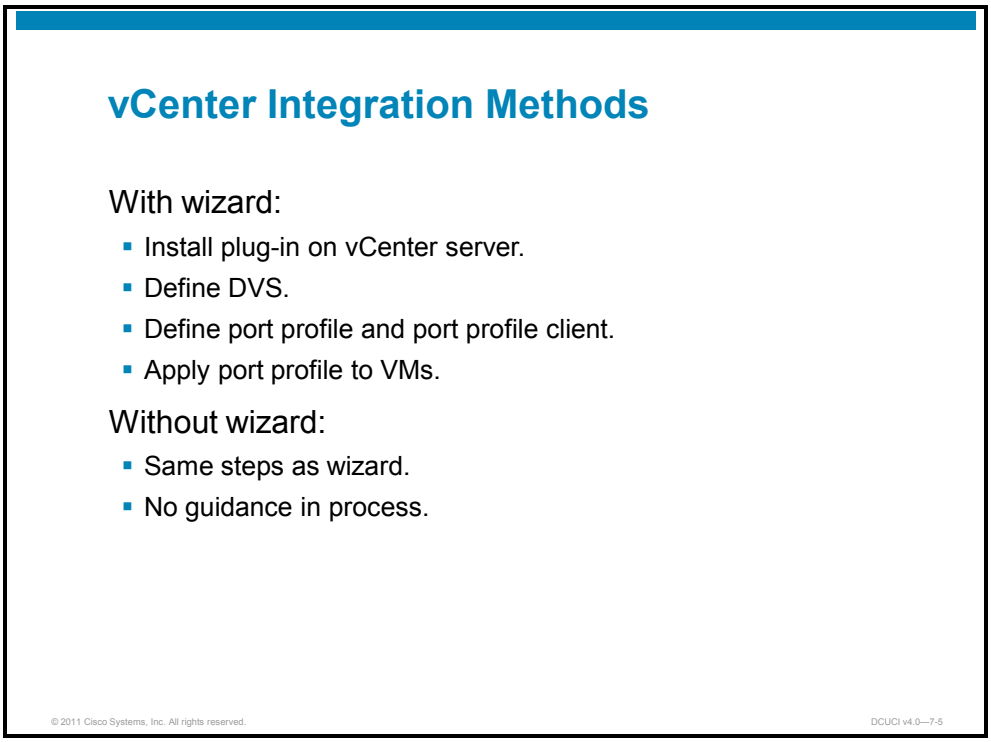

You can use two methods to integrate vCenter with Cisco UCS Manager: using a wizard or manual integration. The wizard guides you through all the required steps. Manual integration requires that you know which steps to perform and in which order.

vCenter Integration Wizard

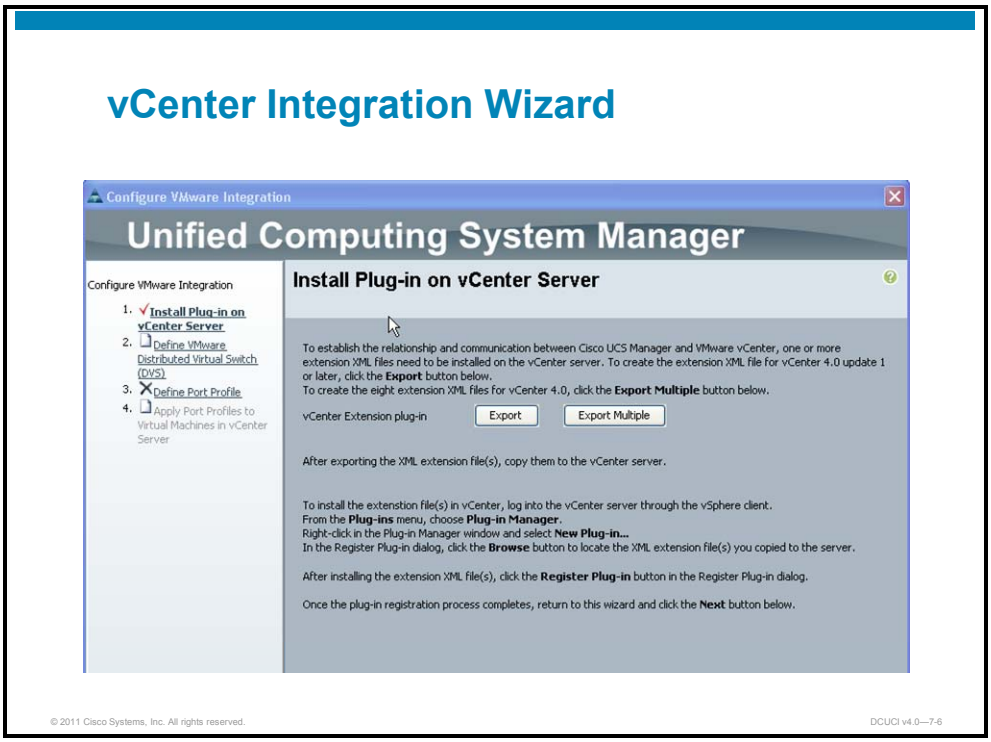

The figure depicts the vCenter integration wizard. The left-hand side of the window shows the four steps that the wizard guides you through:

1. Install the plug-in.
2. Configure the DVS.
3. Produce port profiles in Cisco UCS Manager.
4. Consume port profiles in vCenter.

Downloading the Security Plug-In in Cisco UCS Manager

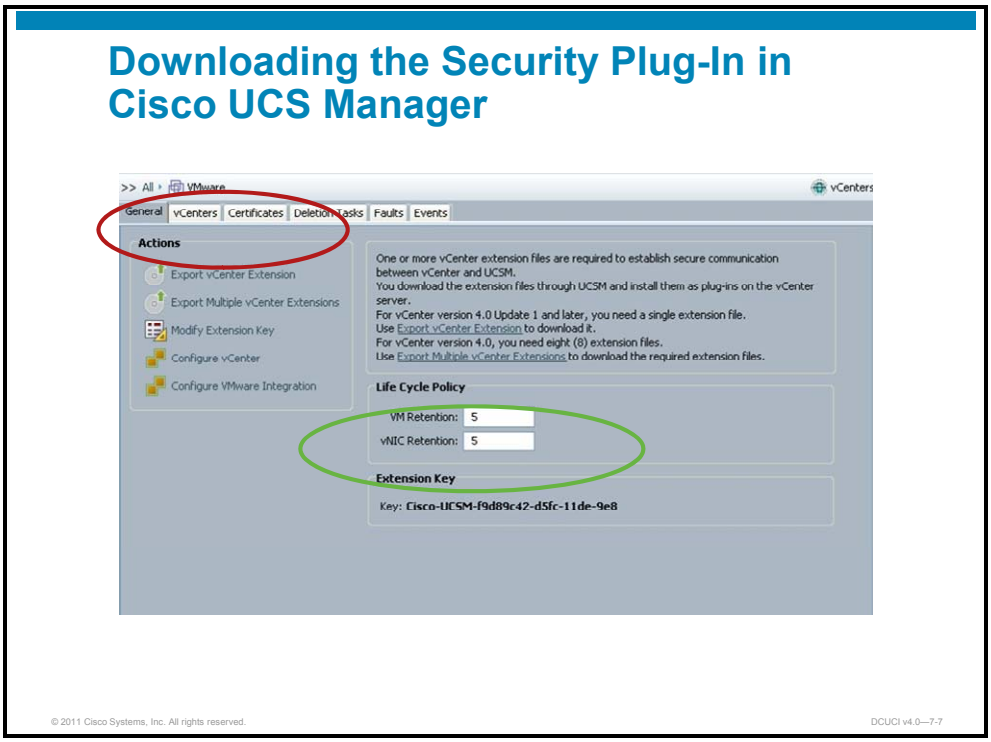

The security plug-in is a Secure Sockets Layer (SSL) certificate that Cisco UCS Manager generates. You can use the default key that Cisco UCS Manager creates (as shown in the figure), or you can manually modify the key. Either way, you must download the certificate in a file to your client machine so that you can then upload it into vCenter.

In a lab environment, set the lifecycle times to a range of 1 to 5 minutes. If you need to remove all port profiles, Cisco UCS Manager locks the configurations for a default period of 15 minutes.

Installing the Security Plug-In in vCenter

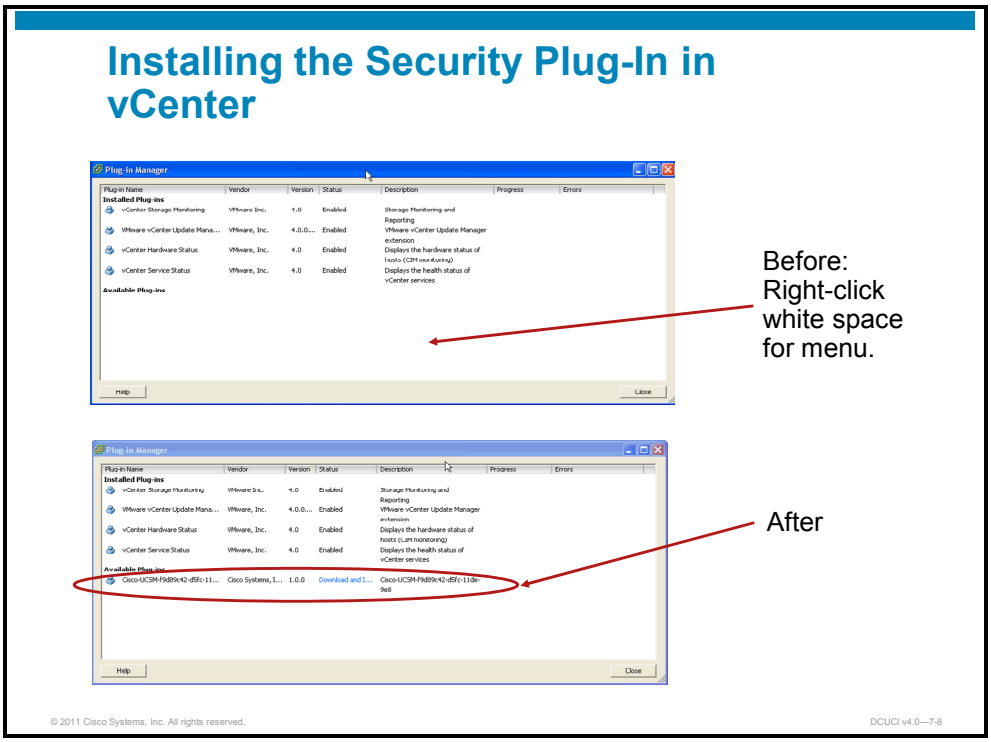

To install the certificate into vCenter, navigate to the plug-in manager. Right-click the available plug-ins area, and choose **Install Plug-in** from the menu.

You are prompted to navigate your client machine file system to the location to which you downloaded the plug-in from Cisco UCS Manager.

Configure Cisco UCS Manager to Connect to vCenter

This topic discusses how to configure Cisco UCS Manager to connect to vCenter.

Using Folders

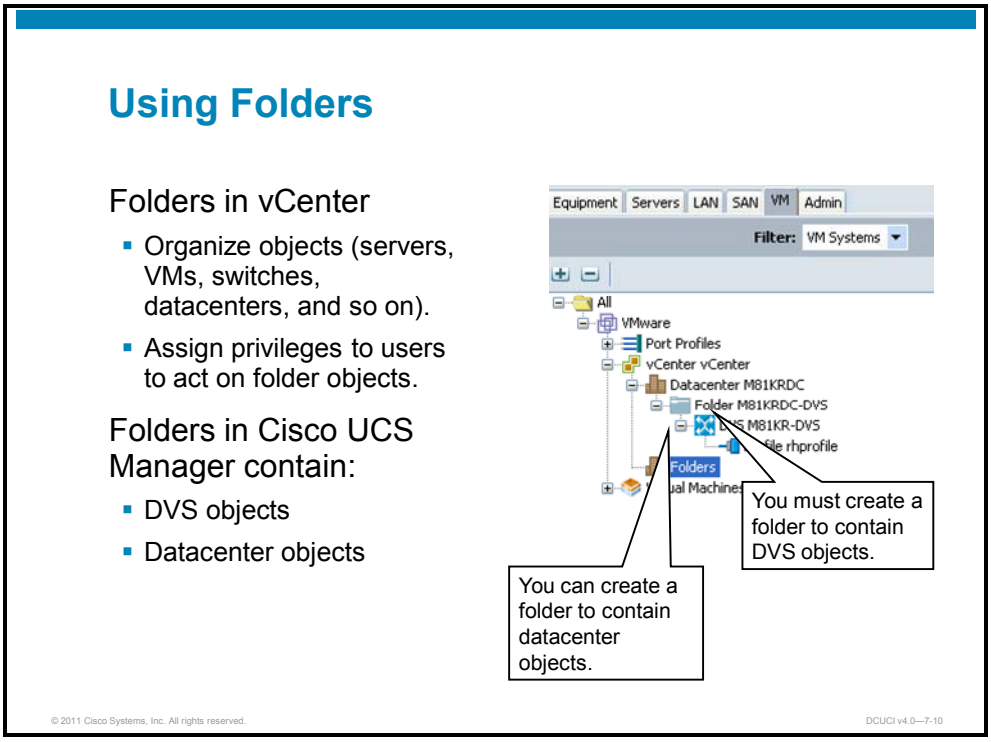

Folders are container objects in vCenter. Folders are used to organize other vCenter objects, such as servers, VMs, switches, and data centers. You can also use folders to assign privileges to users so that they can act on the objects in the folder.

When you integrate vCenter with Cisco UCS Manager, the folders are used to contain DVS and datacenter objects only. You must configure a folder to contain the DVS objects, whereas you may decide whether or not to create folders to contain datacenter objects.

Note	Although your folder names do not need to match those used in vCenter, you should use the same names to simplify correlation.

Configure Uplink and vEthernet Profiles

This topic discusses configuring uplink and vEthernet profiles.

QoS Policy

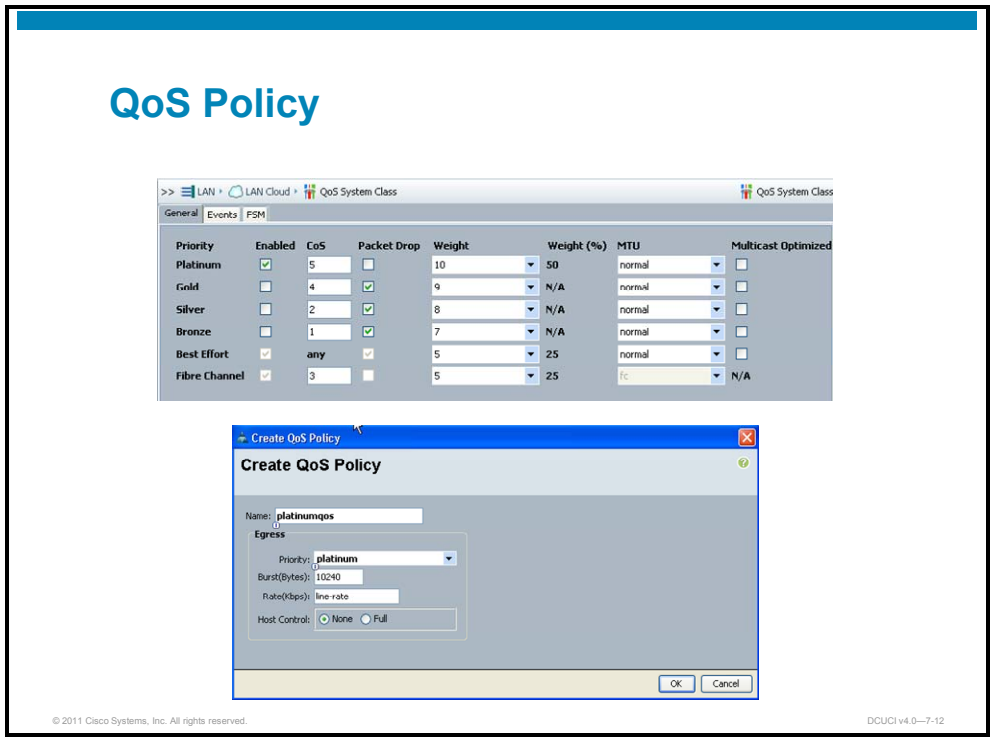

To associate a quality of service (QoS) policy with a port profile, you first must enable the corresponding QoS system class and create a QoS policy with that class. By default, all Ethernet traffic is delivered with best effort priority. You can enable and use bronze, silver, gold, and platinum classes, as shown in the figure.

For each of these classes, you can define the class of service (CoS) value, whether to tolerate packet drop, associated weight (for platinum class only, relative to best effort and Fibre Channel traffic), maximum transmission unit (MTU), and whether to optimize multicast traffic.

If you enable jumbo frames for any interface, you further limit the total number of PTS interfaces that you can configure, as described earlier in this lesson.

Creating a Port Profile

Creating a Port Profile

Port profile contains:

- QoS policy
- Network control policy
- Max ports
- Pin group
- VLANs

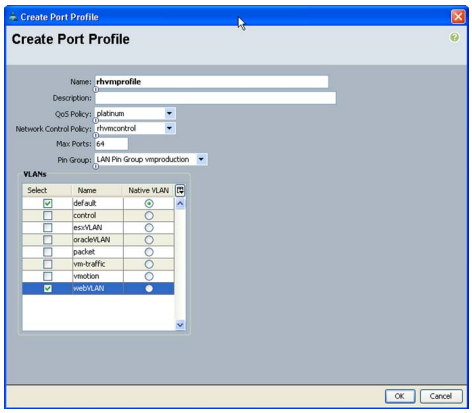

A port profile contains configuration information that ultimately is associated with a VM network interface. The configuration includes QoS, network control, maximum number of ports, a pin group, and one or more VLANs.

Network Control Policy

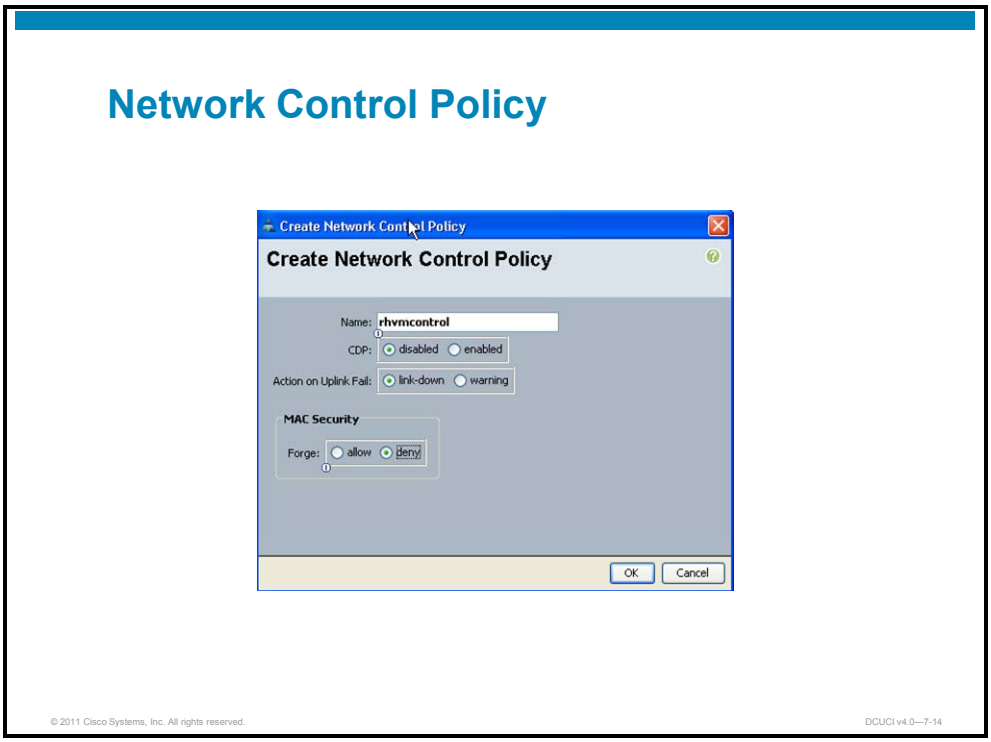

The network control policy allows you to enable or disable Cisco Discovery Protocol membership, define which action Cisco UCS Manager takes on link failures, and decide whether to allow MAC forging.

LAN Pin Group

LAN Pin Group

Dynamic pinning

- Cisco UCS Manager assigns vNICs to uplinks
- Round robin
- Automatic repinning on uplink failure

Static pinning

- Manually assign vNICs to uplinks
- Admin responsible for load
- No repinning on uplink failure

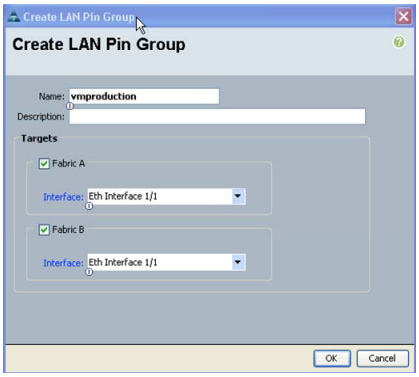

DCUCI v4.0—7-15

The default behavior for Cisco UCS Manager is to dynamically assign (or pin) all vNICs to available uplinks, by using a round-robin algorithm. The benefit of this approach is that Cisco UCS Manager automatically repins a vNIC when the uplink to which it is pinned has failed. The disadvantage of this approach is that the round-robin algorithm might not yield optimal overall traffic distribution across uplinks.

Alternatively, you can statically pin vNICs to uplinks. The benefit of this approach is that you can define the overall traffic distribution for all the vNICs that you statically pin. The disadvantage of this approach is that Cisco UCS Manager does not automatically repin such a vNIC upon uplink failure.

Creating a Port Profile Client

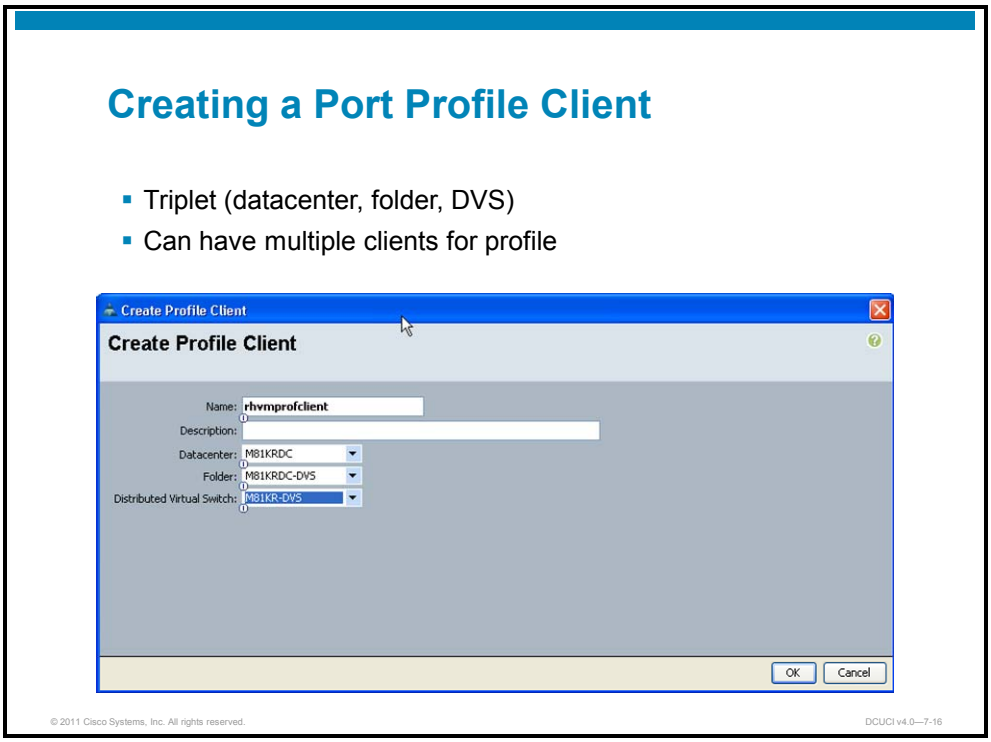

A port profile client defines the consumer of a port profile. The consumer is defined as a DVS in a folder that is associated with a datacenter object in vCenter. You can associate the same port profile with multiple port profile clients.

Configure Service Profiles with Dynamic NICs

This topic discusses configuring service profiles with dynamic NICs.

Calculating the Deployment Limits with the M81KR VIC

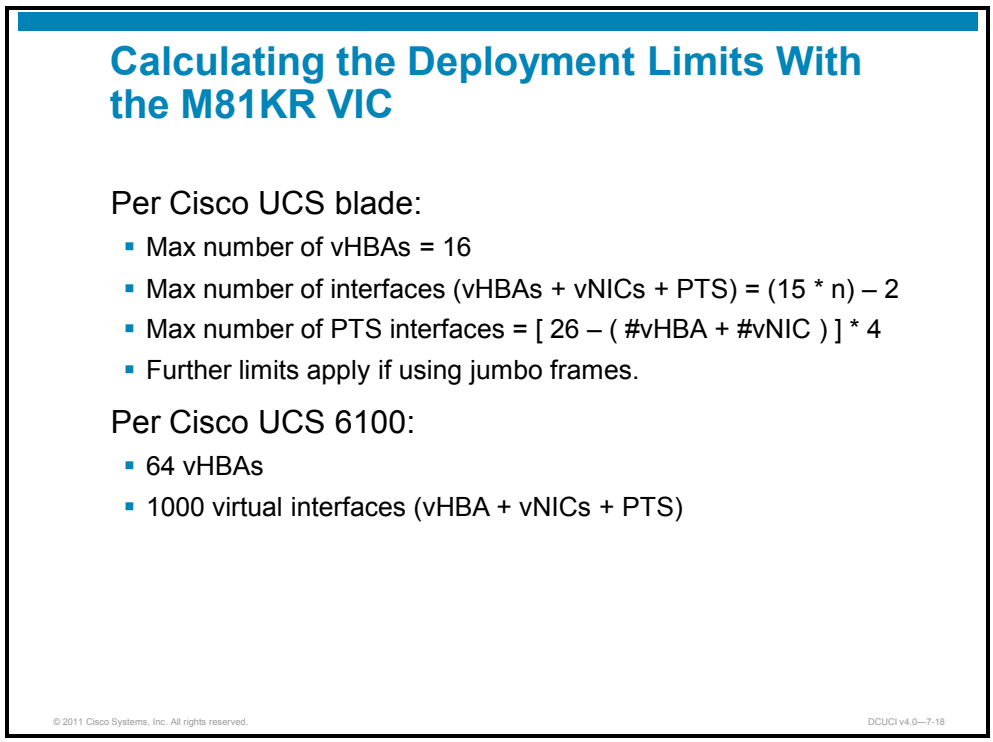

There are several limits to consider when you are trying to figure out how many static and dynamic interfaces you can configure on an M81KR VIC. The M81KR can support a maximum of 128 interfaces, but this number is further reduced by the limits that are defined in the figure.

The maximum number of interfaces that you can configure is based on a physical limit that is defined in the Cisco UCS 6100 Series Fabric Interconnects. The formula is defined as (15 * *n*) – 2, where *n* is the number of server links that connect the I/O module (IOM) to the fabric interconnect. For example, with a single link, you can configure (15 * 1) – 2, or 13, total interfaces. These interfaces include the sum of vHBAs, static vNICs, and dynamic (or PTS) vNICs.

The maximum number of PTS interfaces that you can configure is given as [26 – (#vHBA + #vNIC)] * 4. For example, if you configure two vHBAs and four VNICs, then you are bound to configure no more than [26 – (2 + 4)] * 4, or 80, PTS interfaces. Recall, however, that you are also bound by the first formula as a function of the number of server links.

Configuring a Dynamic vNIC Connection Policy

Configuring a Dynamic vNIC Connection Policy

- Can configure multiple policies per Cisco Unified Computing System.
- Adapter policy for PTS is VMware PassThru.
- Interfaces are protected by default.

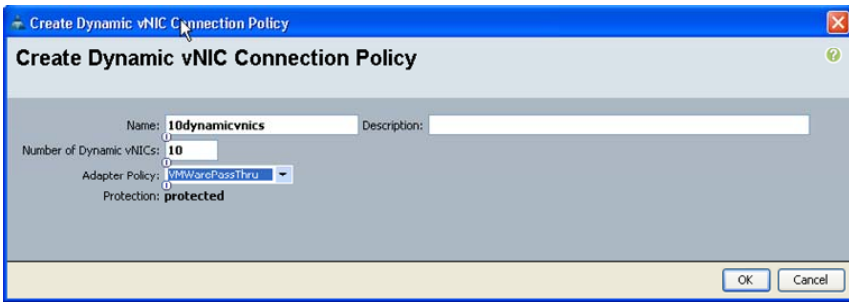

A dynamic vNIC connection policy defines the number of PTS interfaces that you want to instantiate on an M8lKR VIC. Recall the limitations that were defined earlier in this lesson.

Because you can configure varying numbers of server links to connect IOMs from different chassis to the fabric interconnects, you can feasibly have some M81KR VICs with more PTS interfaces than others. However, note that when you migrate a service profile with an associated connection policy from one chassis to another, you might run into problems. Specifically, if the number of IOM links on the target chassis is fewer than the number of IOM links on the source chassis, then fewer PTS interfaces will be available on those blades.

Associating a Dynamic vNIC Connection Policy to a Profile

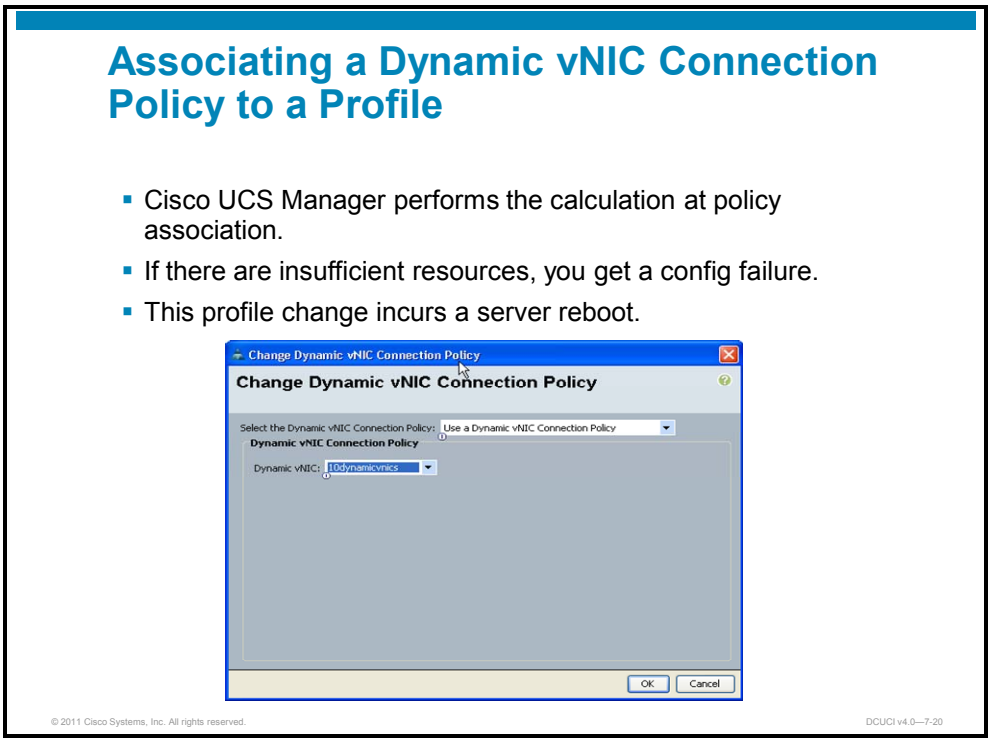

When you associate a dynamic vNIC connection policy to a service profile, Cisco UCS Manager calculates the total number of PTS interfaces that you can configure. If that number is fewer than the number that you specify in your connection policy, Cisco UCS Manager gives you a configuration failure.

Note	The server must reboot when you associate a connection policy to a service profile.

Viewing Dynamic vNICS in Cisco UCS Manager

Viewing Dynamic vNICs in Cisco UCS Manager

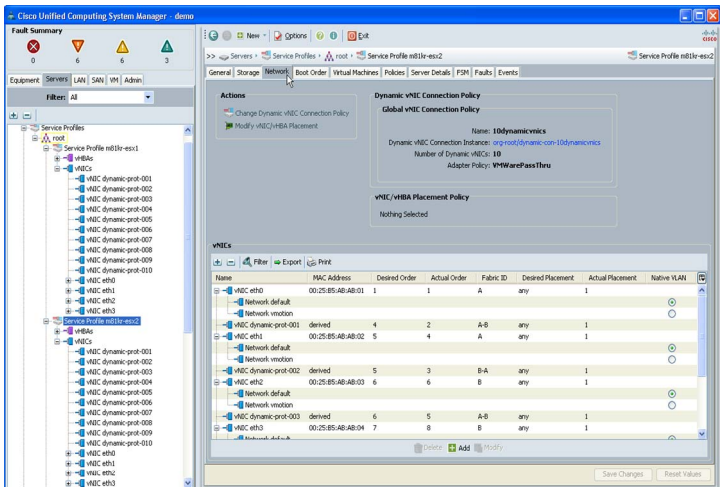

DCUCI v4.0—7-21

After successfully associating a connection policy to a service profile, you can view the instantiated PTS interfaces in Cisco UCS Manager, as shown in the figure. These interfaces are consumed dynamically as you configure network interfaces on your VMs.

Viewing Dynamic vNICS in ESX

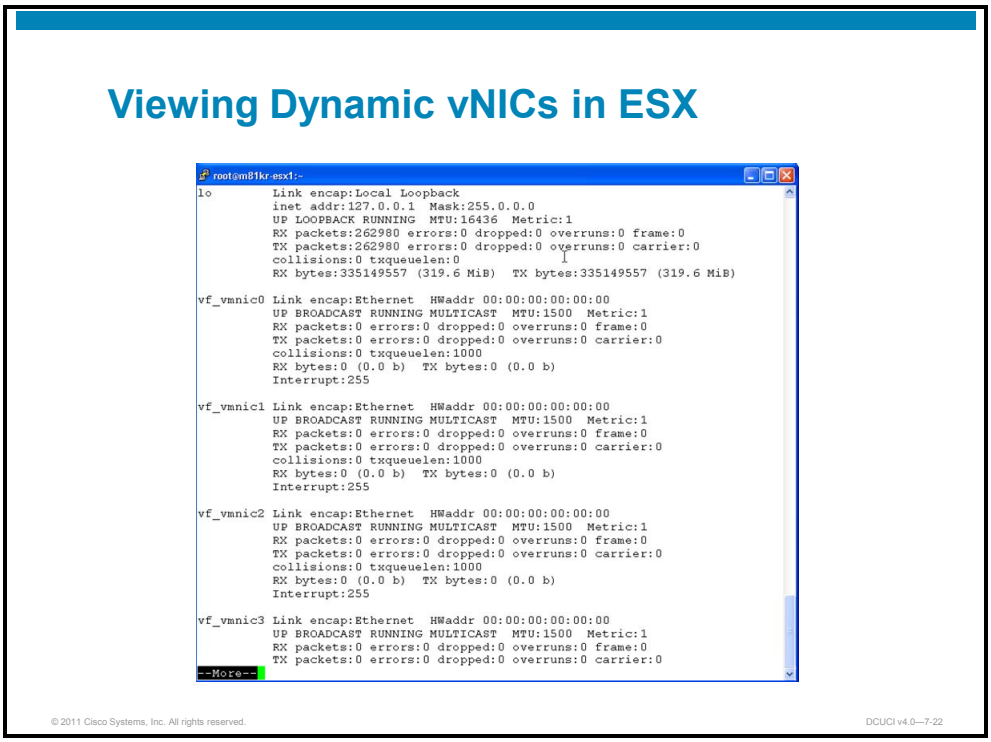

The dynamic vNICs that have been configured on the ESX server can be viewed from the ESX CLI, as shown in the figure. These dynamic vNICs are named vf_vmnicX.

Configure vMotion and M81KR Port Profile Mobility

This topic discusses configuring vMotion and M81KR port profile mobility.

Add the Hosts to the DVS in vCenter

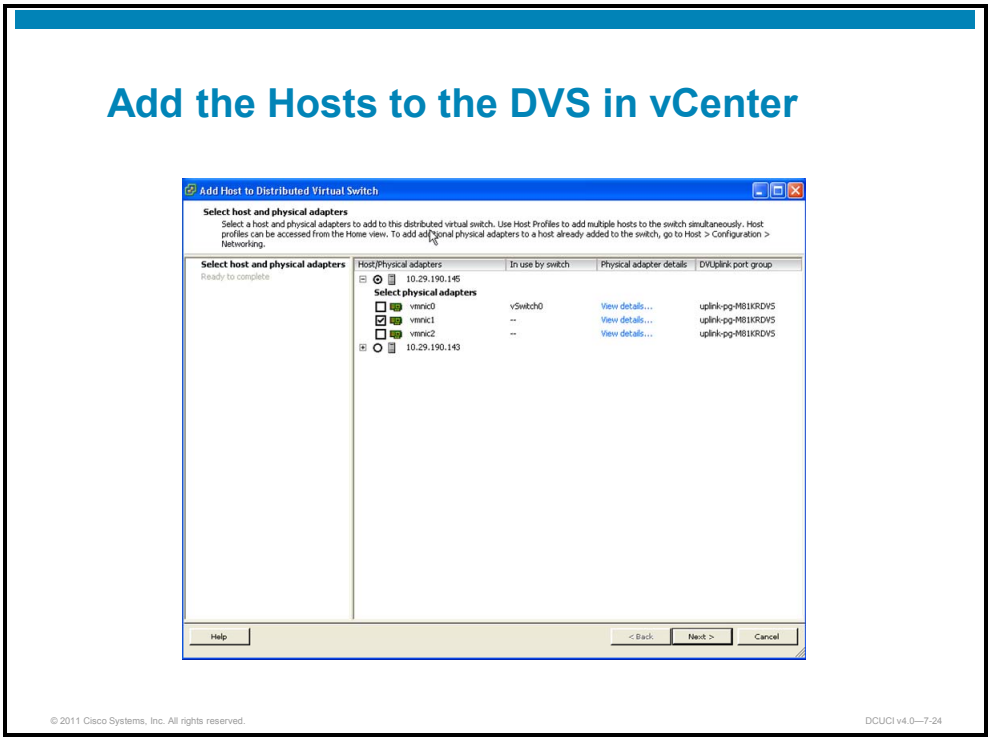

To add a host to the DVS, navigate to the DVS in the networking inventory view. Right-click **Add Host to the Distributed Virtual Switch** from the drop-down menu.

When adding a host to the DVS, you must select static interfaces that have already been configured in Cisco UCS Manager to use as uplinks to the DVS. Unlike with Cisco Nexus 1000V, you do not manually configure the uplink port groups. These ports are configured automatically for you.

Note	If you use VMware Update Manager to install the VEMs, then when adding hosts to the DVS, VMware Update Manager remediates the ESX servers to ensure that they have the correct version of the VEM installed.

View Hosts in vCenter

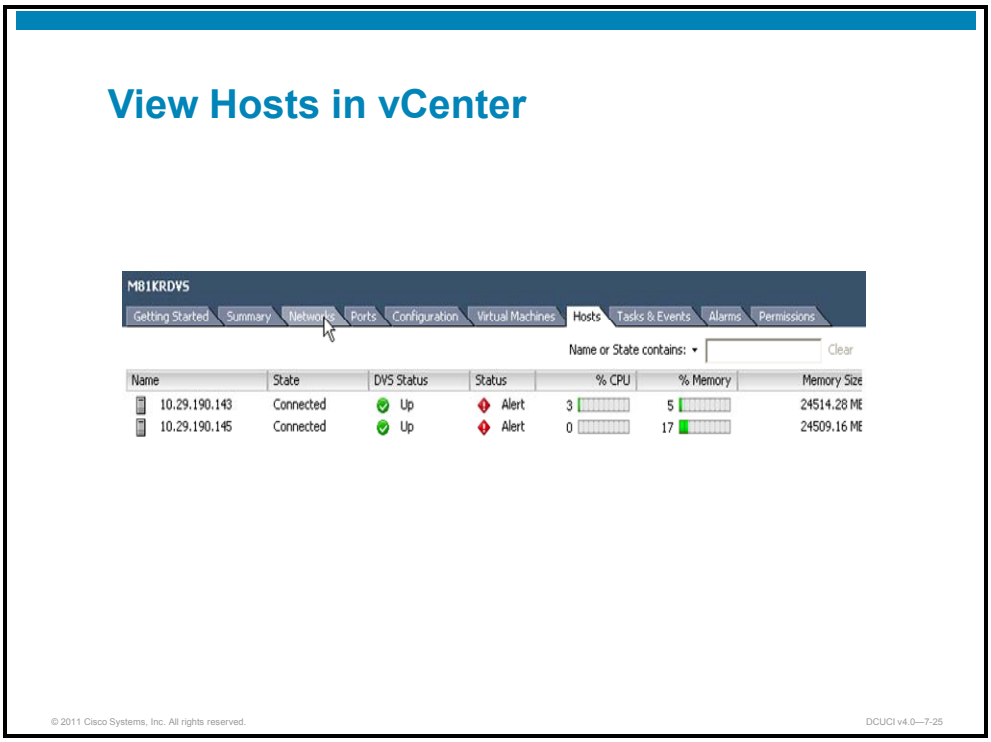

After adding all hosts to the DVS, you can view their status from the networking inventory view, as shown in the figure.

Viewing Available Port Profiles in vCenter

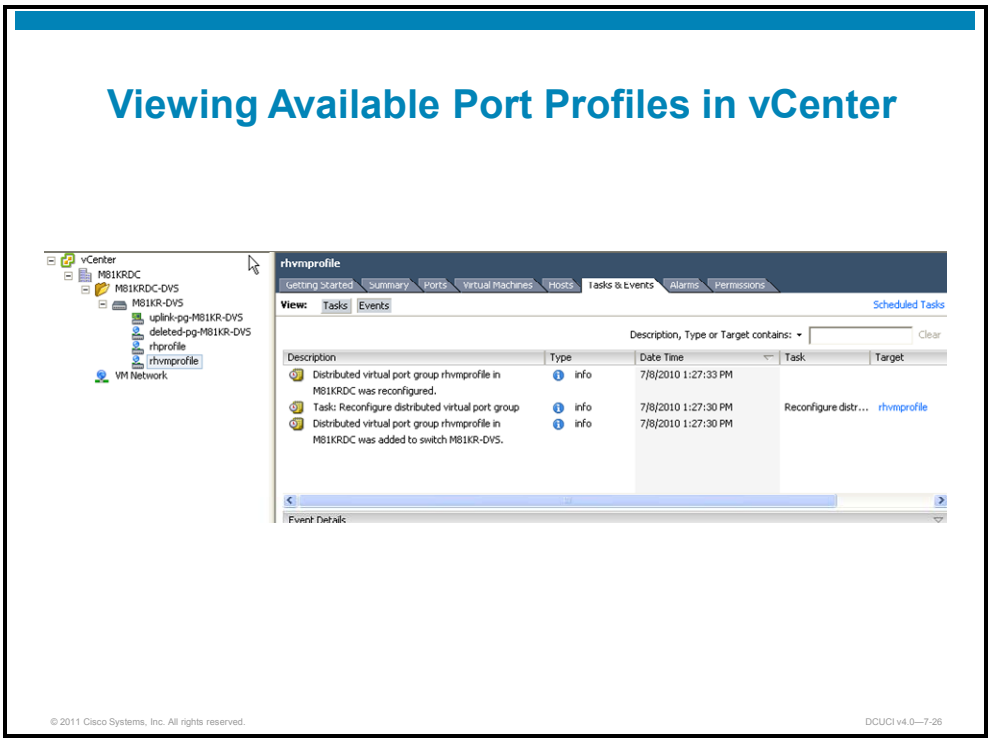

Port profiles that are defined in Cisco UCS Manager appear as port groups in vCenter. You can view all available port groups in your port profile client by navigating to the appropriate datacenter, folder, and DVS in your VMware vSphere client. In the figure, the port group named rhvmprofile is available for consumption.

Assigning Port Profile to the VM Interface

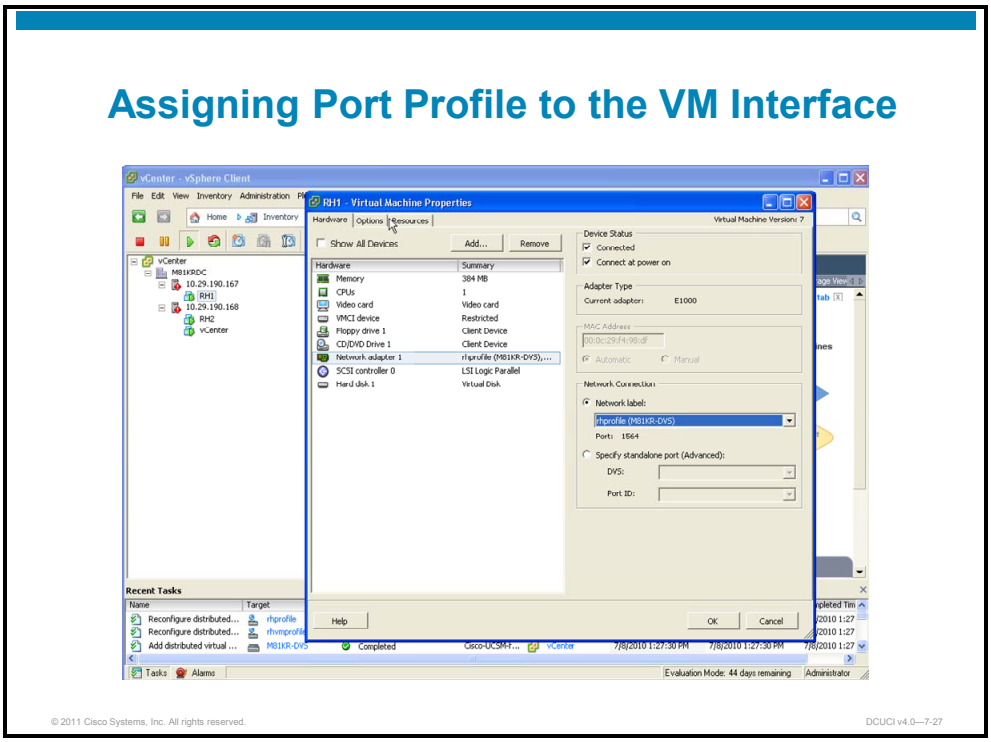

You can associate any port group available in the port profile client to any VM network interface. The figure shows the association of the port group named rhprofile to the RH1 VM network adapter.

Viewing Port Profile Consumption in Cisco UCS Manager

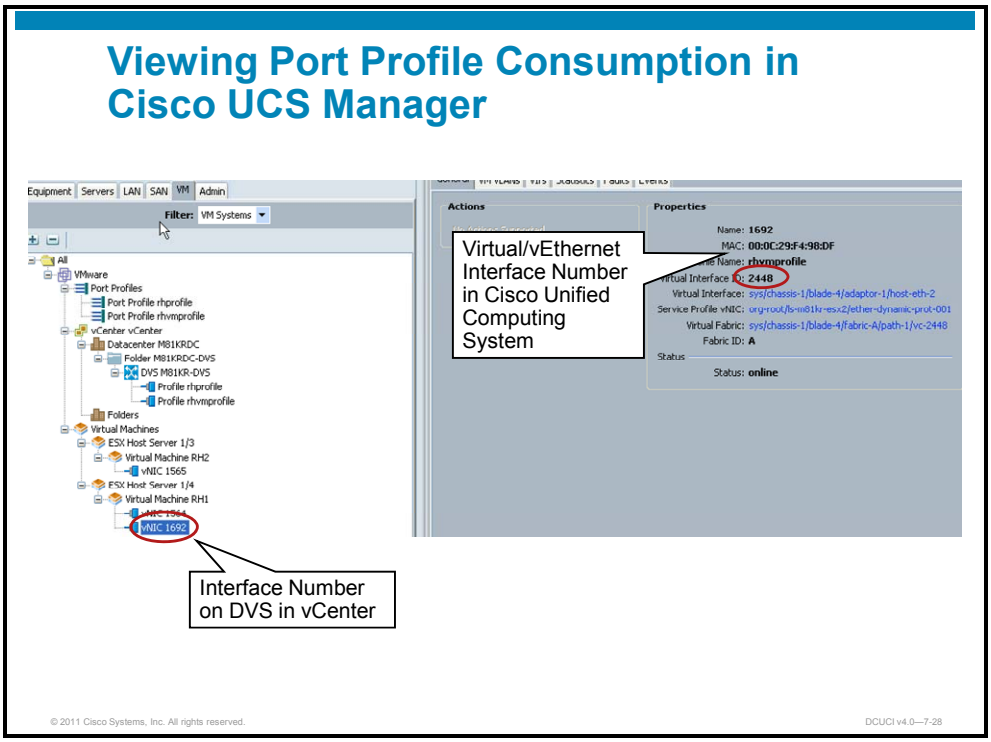

Going back to the Cisco UCS Manager interface, you can now see, under the VM tab, which port profiles have been consumed by which port profile clients, as shown in the figure. In this example, an interface vNIC 1692 was instantiated on the DVS in the port profile client. This interface is associated with a virtual interface 2448 in Cisco Unified Computing System.

Displaying Virtual Interface/vEthernet Packet Counters in the Cisco UCS Manager CLI

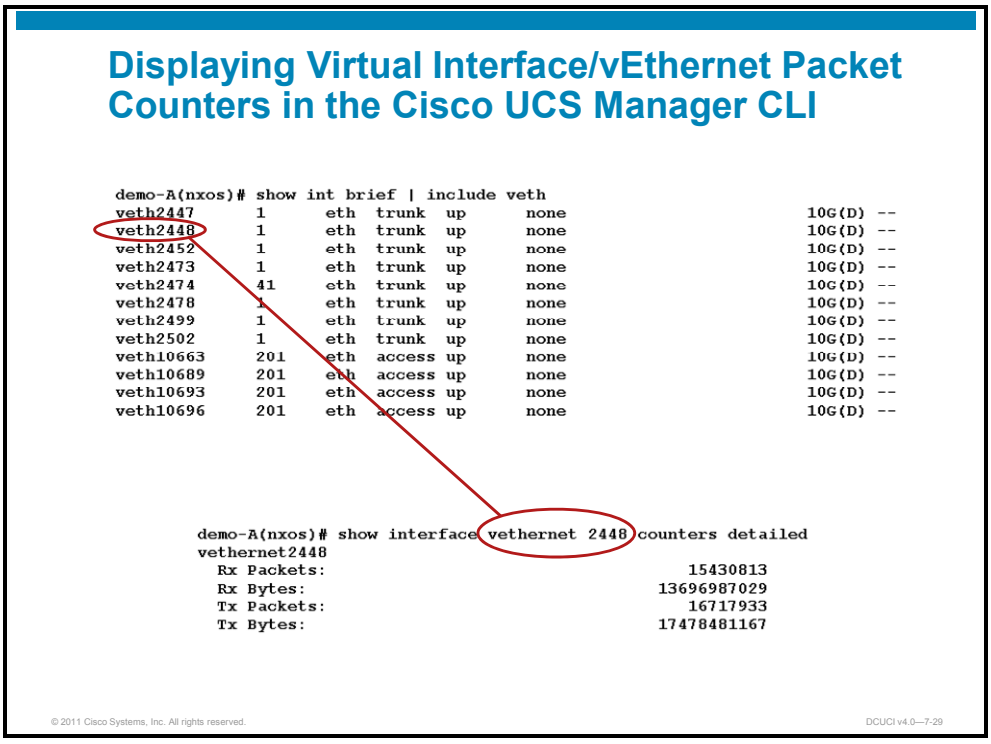

Now that you know the vEthernet interface number that is associated with your VM network interface, you can view packet counters and status for that interface in the connect nxos shell of Cisco UCS Manager.

Summary

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

Summary

- The PTS feature requires configuration on both Cisco UCS Manager and vCenter, to provide two-way communication.
- The Cisco UCS Manager extension is installed as a plug-in to vCenter.
- Datacenter, PTS, and folder objects are used to organize and assign parameters to hosts and VMs within Cisco UCS Manager.
- Uplink and vEthernet port profiles are created within Cisco UCS Manager and pushed out to vCenter.
- Dynamic NICs are created by using a policy wizard within Cisco UCS Manager, are assigned within the service profile, and appear within both the Cisco UCS Manager and vCenter inventories.

Module Summary

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

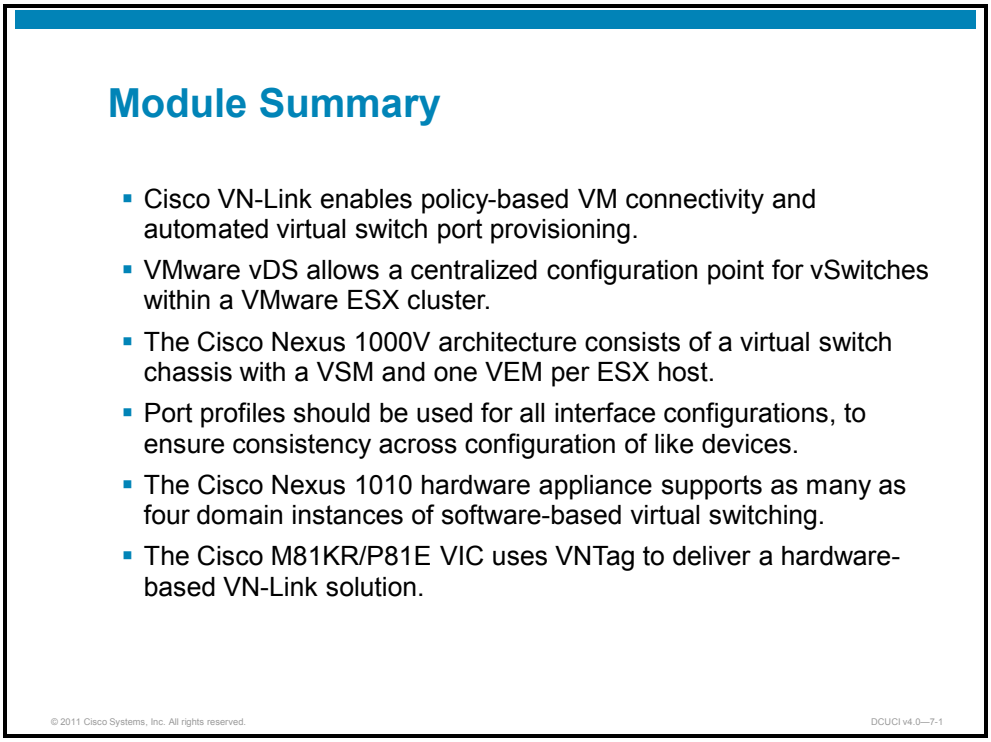

This module introduced the Cisco Nexus 1000V Distributed Virtual Switch (DVS), which is a third-party plug-in to VMware vCenter. The features and benefits of the Cisco Nexus 1000V DVS were discussed, as were detailed installation and configuration methods. In addition, a comparison of the Cisco Nexus 1000V, Cisco Nexus 1010 Virtual Services Appliance, and the Cisco M81KR/P81E Virtual Interface Card (VIC) were presented.

Module Self-Check

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

Q1) What is a feature of the Cisco Nexus 1000V Series Switch? (Source: "Evaluating the Cisco Nexus 1000V Switch")

A) port profiles
B) Fibre Channel
C) hypervisor
D) vMotion

Q2) Which VMware feature enables automated migration of VMs between physical hosts, in response to defined thresholds? (Source: "Evaluating the Cisco Nexus 1000V Switch")

A) fault tolerance
B) DRS
C) PTS
D) clustering

Q3) In traditional VMware deployments, which group configures virtual switches? (Source: "Evaluating the Cisco Nexus 1000V Switch")

A) network
B) development
C) storage
D) server

Q4) Which property can a Cisco Nexus 1000V port profile define? (Source: "Configuring Basic Cisco Nexus 1000V Networking")

A) link speed
B) VLANs
C) VM guest operating system
D) switch high availability

Q5) Which feature does the VMware vDS introduce that does not exist in traditional virtual switches? (Source: "Evaluating the Cisco Nexus 1000V Switch")

A) hypervisor
B) service console
C) port groups that span the data center
D) VMkernel

Q6) Which VN-Link implementation does the Cisco M81KR VIC use? (Source: "Configuring Cisco UCS Manager for VMware PTS")

A) software
B) virtual
C) hardware
D) integrated

Q7) Cisco Nexus 1000V VSMs can be configured for which type of high availability? (Source: “Characterizing Cisco Nexus 1000V Architecture”)

A) standalone
B) high-availability standby
C) primary
D) secondary

Q8) What are the two types of Cisco Nexus 1000V port profiles? (Choose two.) (Source: “Configuring Basic Cisco Nexus 1000V Networking”)

A) uplink
B) physical
C) vEthernet
D) server based

Module Self-Check Answer Key

Q1) A

Q2) B

Q3) D

Q4) B

Q5) C

Q6) C

Q7) B

Q8) A, C