Scalability

Module 10

You Are Here

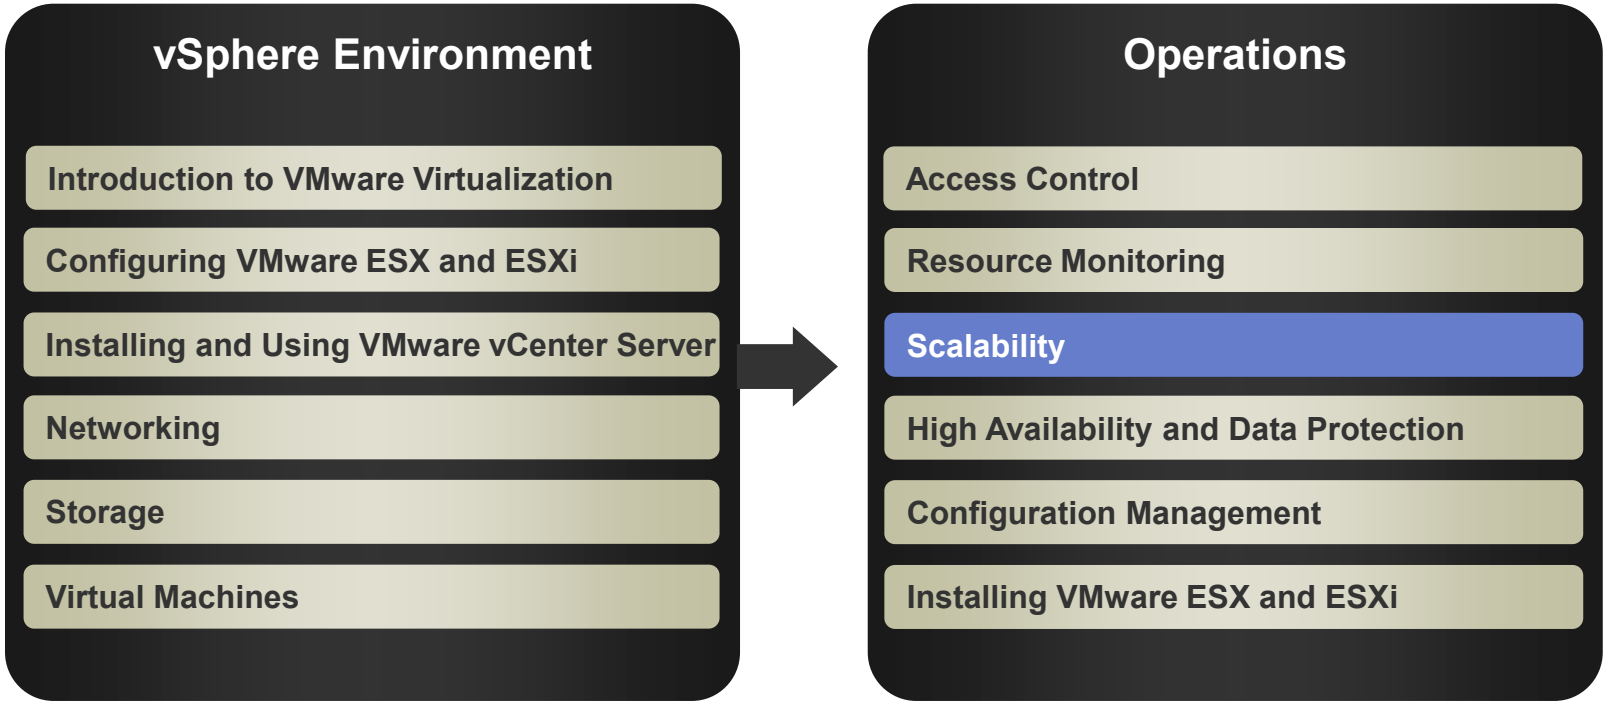

Importance

- Resource pools allow CPU and memory resources to be hierarchically assigned. Clusters enabled for VMware® Distributed Resource Scheduler (DRS) provide automated resource management for multiple VMware ESX™/ESXi hosts.

Module Lessons

Lesson 1: Scaling CPU and Memory Management

Lesson 2: Scaling Storage and Network Management

Lesson 3: VMware VMotion Migration

Lesson 4: VMware Distributed Resource Scheduler

Lesson 1: Scaling CPU and Memory Management

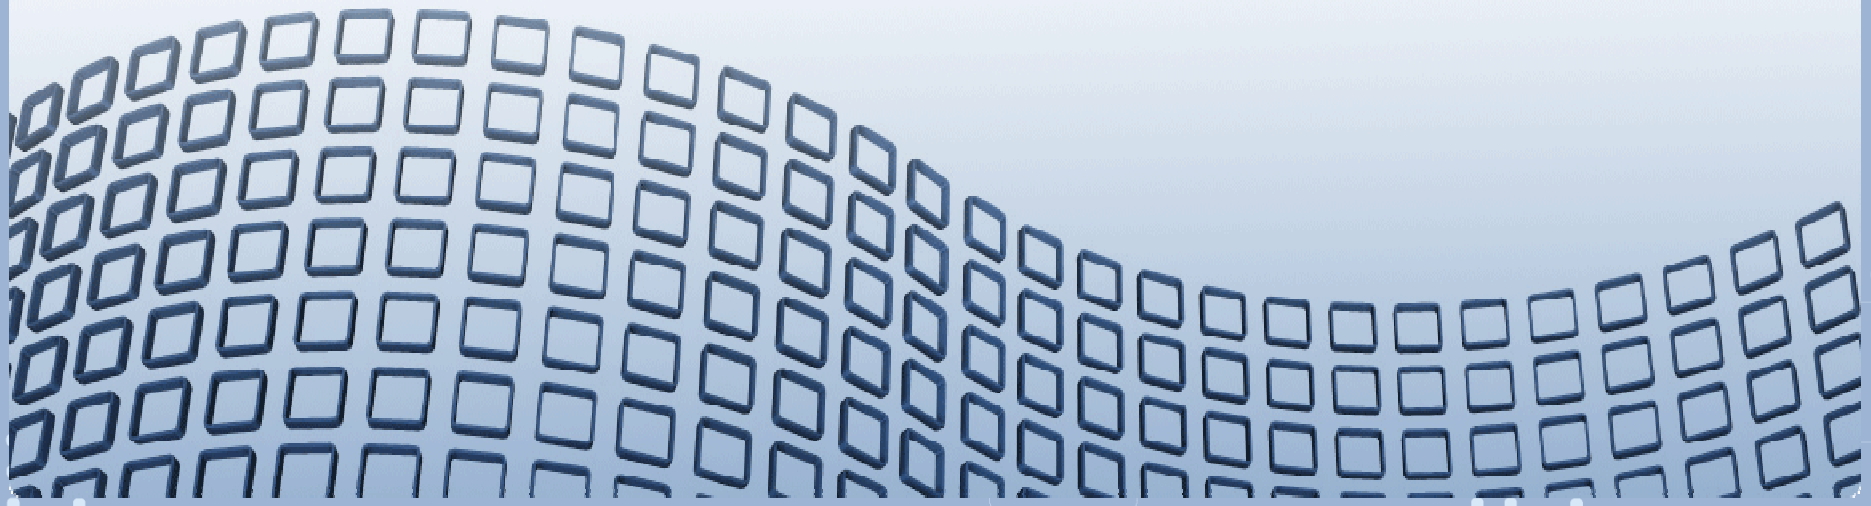

Lesson Objectives

- Describe the CPU and memory resource allocation settings
- Describe a resource pool
- Create a resource pool
- View resource allocation

Resource Management

Resource management is the allocation of resources from providers (hosts and clusters) to consumers (virtual machines).

- Resources include CPU, memory, storage, and network.

Resource management:

- Resolves resource overcommitment
- Prevents virtual machines from monopolizing resources
- Exploits undercommitted resources
- Controls the relative importance of virtual machines

Resource allocation settings – shares, reservation, and limit – are used to determine the amount of CPU and memory resources provided for a virtual machine.

Virtual Machine CPU Resource Settings

Limit

- This value is a cap on the consumption of CPU time by this virtual machine, measured in MHz.

Reservation

- This value is a certain number of CPU cycles reserved for this virtual machine, measured in MHz.
- The VMkernel chooses which CPUs it can migrate.

Shares

- More shares means that this virtual machine will win competitions for CPU time more often.

All the VCPUs in a virtual machine must be simultaneously scheduled.

- Therefore, a reservation of 1,000MHz might be generous for a one-VCPU virtual machine but not for a four-VCPU virtual machine.

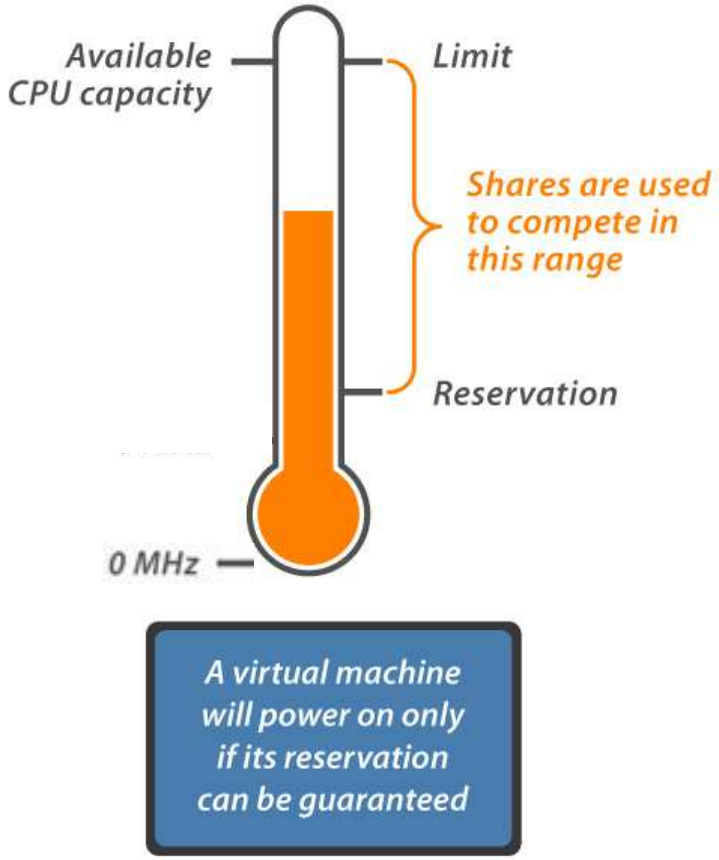

Virtual Machine Memory Resource Settings

Available memory

- This value is the memory size defined when the virtual machine was created.

Limit

- This value is a cap on the consumption of physical memory by this virtual machine, measured in MB.

Reservation

- This value is a certain amount of physical memory reserved for this virtual machine, measured in MB.

Shares

- More shares means that this virtual machine will win competitions for physical memory more often.

VMkernel allocates a per–virtual machine swap file to cover each virtual machine's range between available memory and reservation.

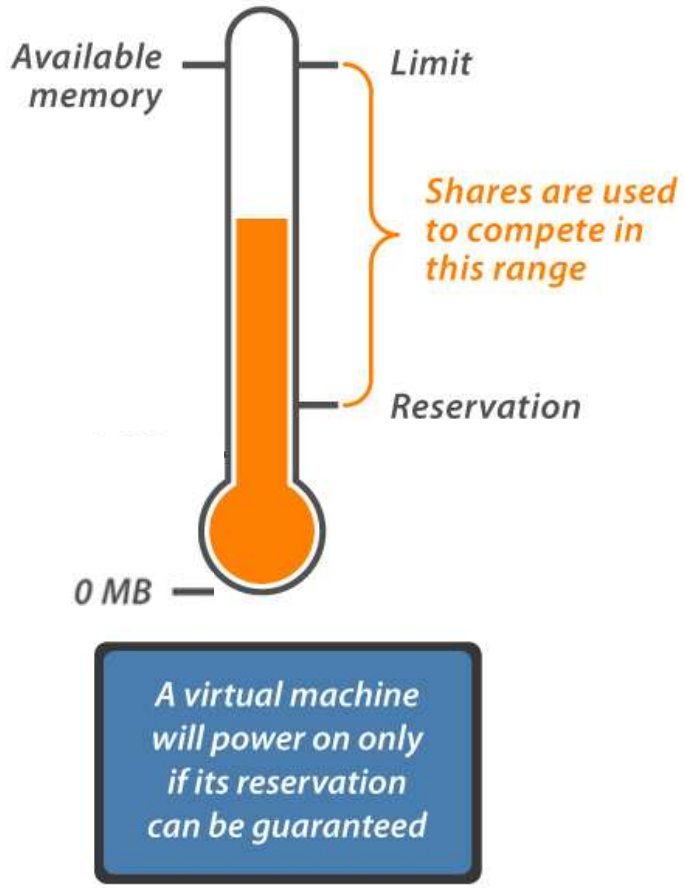

How Virtual Machines Compete for Resources

Proportional-share system for relative resource management

- Applied during resource contention
- Prevents virtual machines from monopolizing resources
- Guarantees predictable resource shares

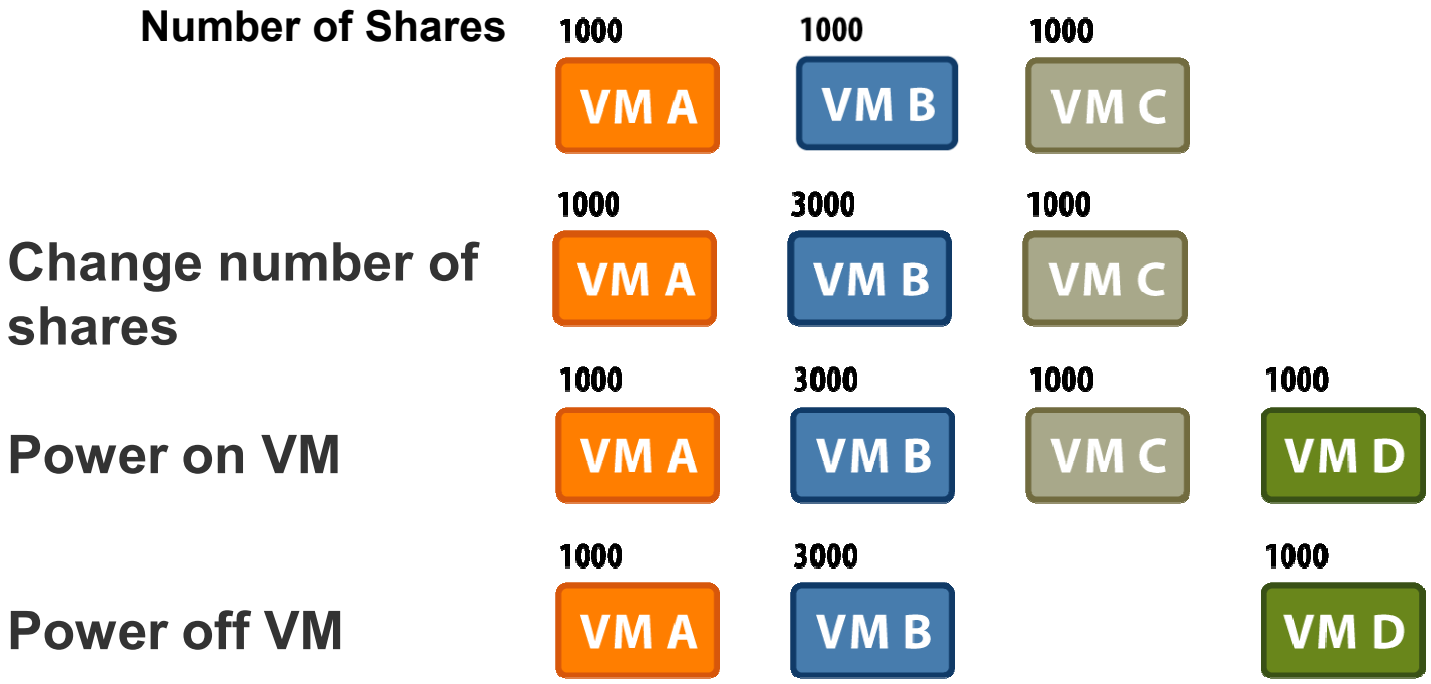

What Is a Resource Pool?

- A logical abstraction for hierarchically managing CPU and memory resources
- Used on standalone hosts or DRS-enabled clusters
- Provides resources for virtual machines and child pools

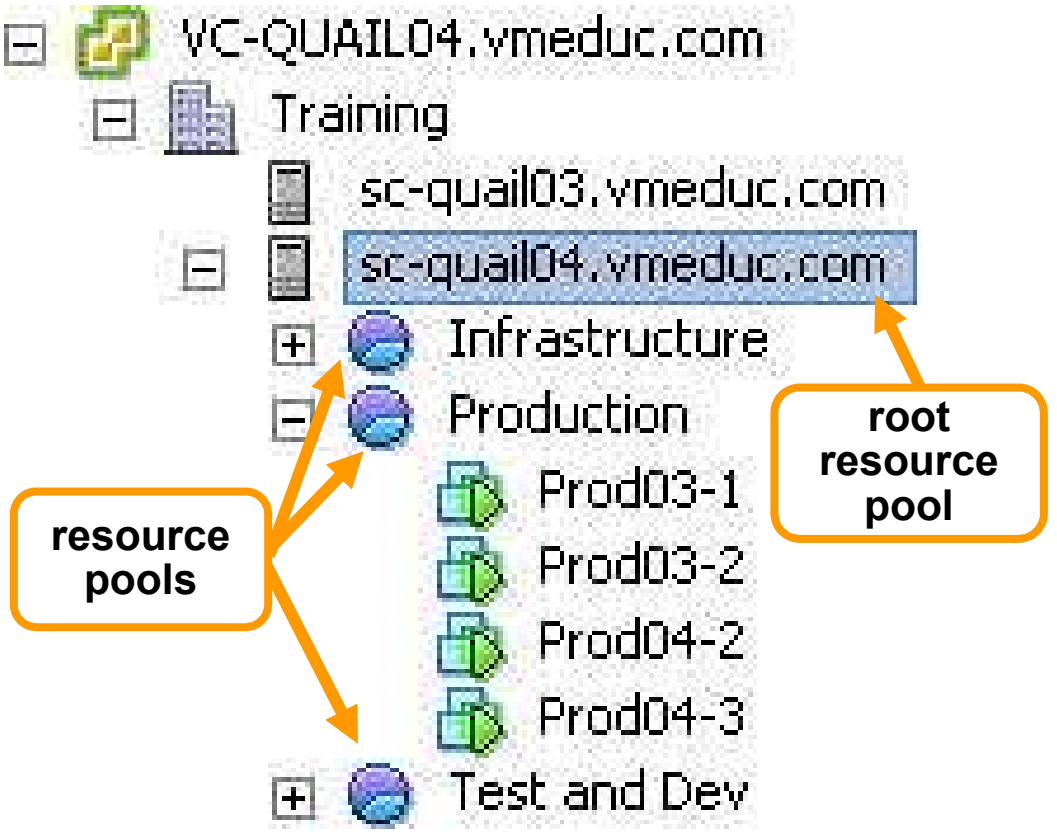

Why Use Resource Pools?

Using resource pools can result in these benefits:

- Flexible hierarchical organization
- Isolation between pools, sharing within pools
- Access control and delegation
- Separation of resources from hardware
- Management of sets of virtual machines running a multitier service

Resource Pool Attributes

Resource pools have the following attributes:

- Shares
  - Low, Normal, High, Custom
- Reservations, in MHz and MB
- Limits, in MHz and MB
  - Unlimited access, by default (up to maximum amount of resource accessible)
- Expandable reservation?
  - Yes – Virtual machines and subpools can draw from this pool's parent.
  - No – Virtual machines and subpools can draw only from this pool, even if its parent has free resources.

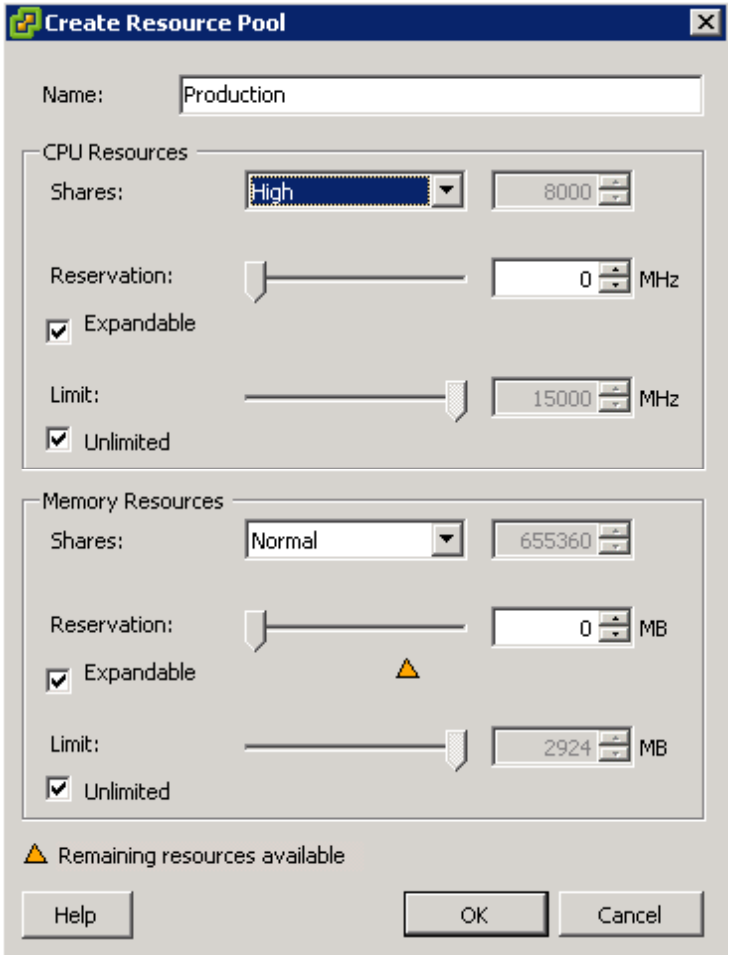

Resource Pool Scenario

Company X's IT department has two internal customers.

- The finance department supplies two-thirds of the budget.
- The engineering department supplies one-third of the budget.

Each internal customer has both production and test/dev virtual machines.

We must cap the test/dev virtual machines' resource consumption.

Resource Pool Example

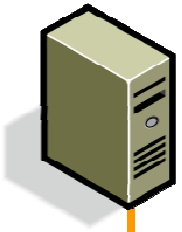

standalone host – Svr001
(root resource pool)
CPU: 12,000MHz
Memory: 4GB

Engineering (Resource Pool)

CPU Shares: 1,000
Reservation: 1,000MHz
Limit: 4,000MHz
Expandable Reservation: Yes

Eng-Test (VM)

CPU Shares: 1,000
Reservation: 0MHz
Limit: 4,000MHz

Eng-Prod (VM)

CPU Shares: 2,000
Reservation: 250MHz
Limit: 4,000MHz

Resource Pools Example: CPU Shares

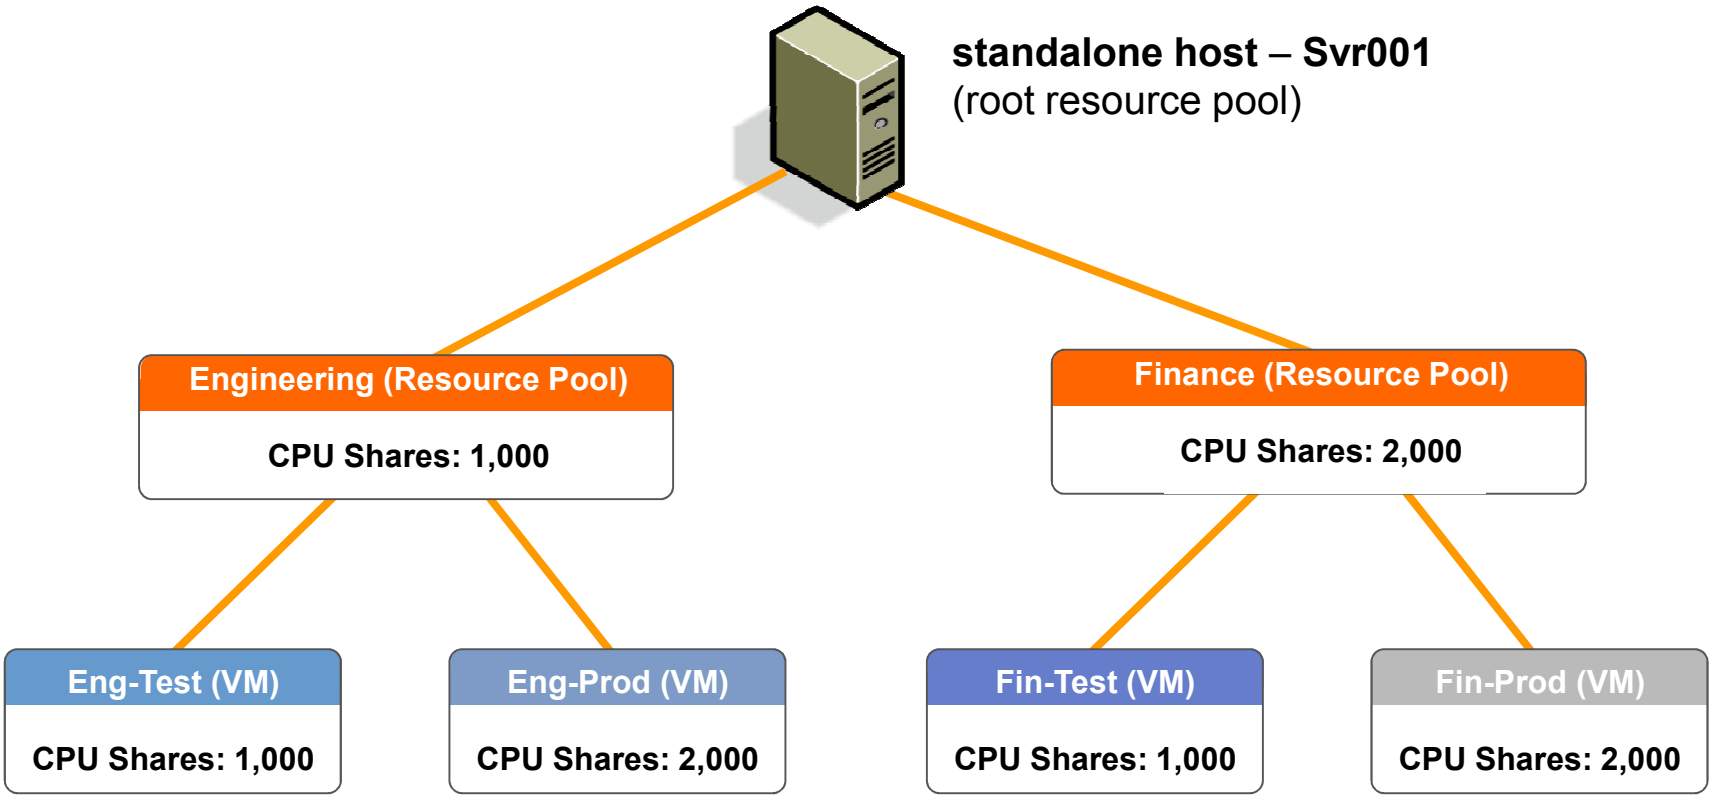

Resource Pools Example: CPU Contention

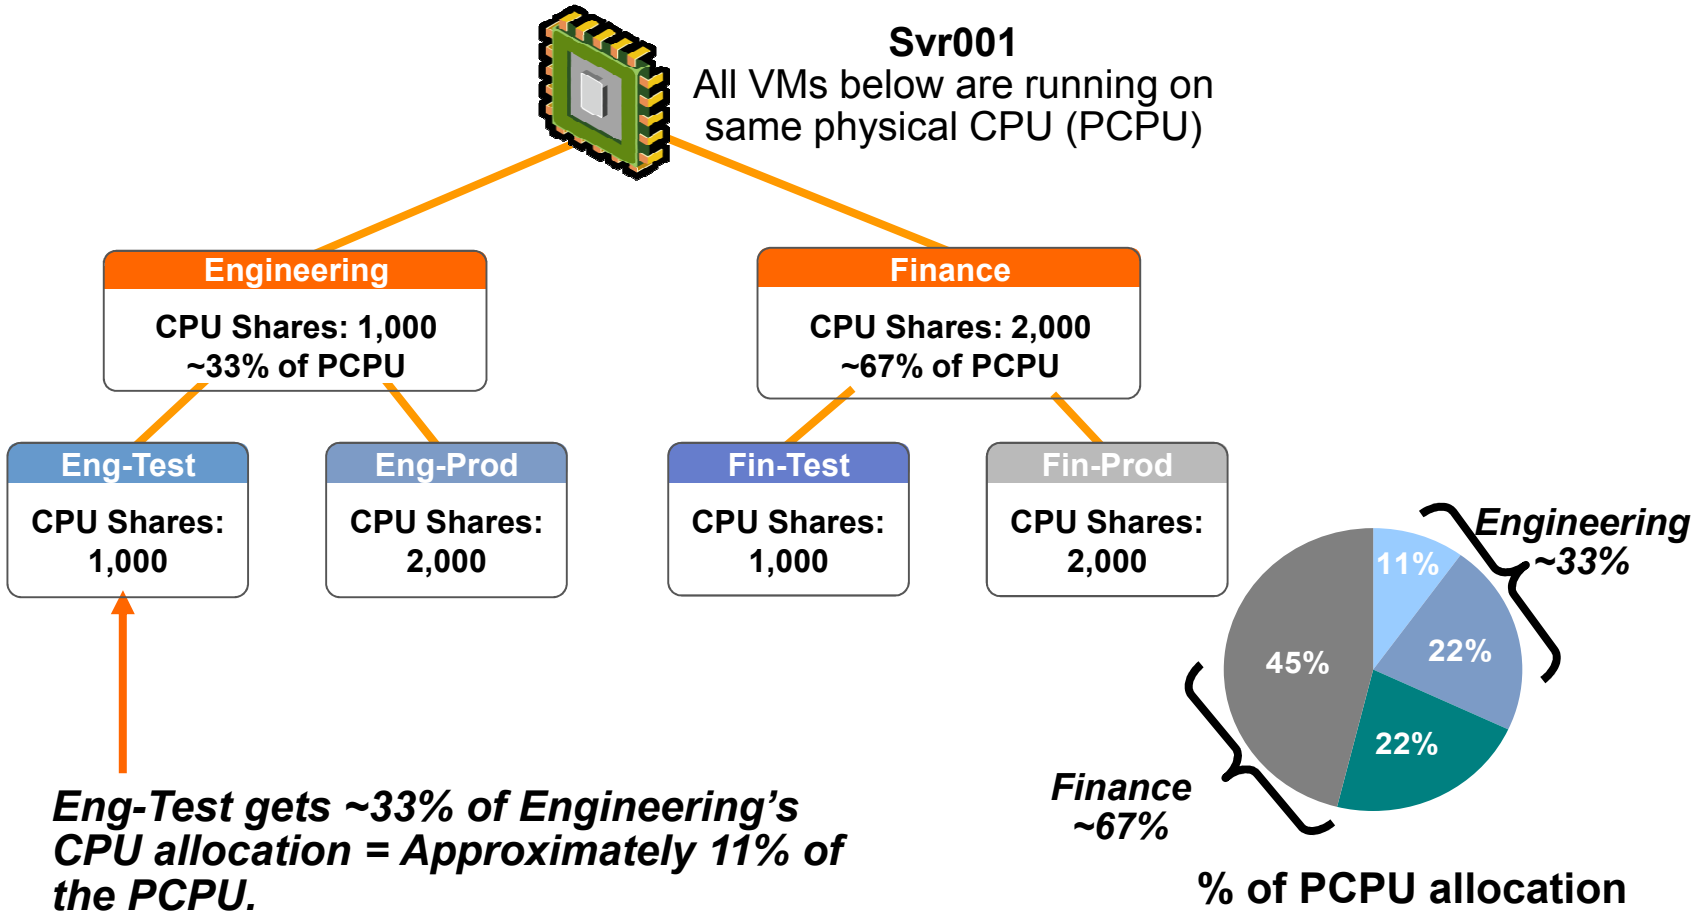

Expandable Reservation

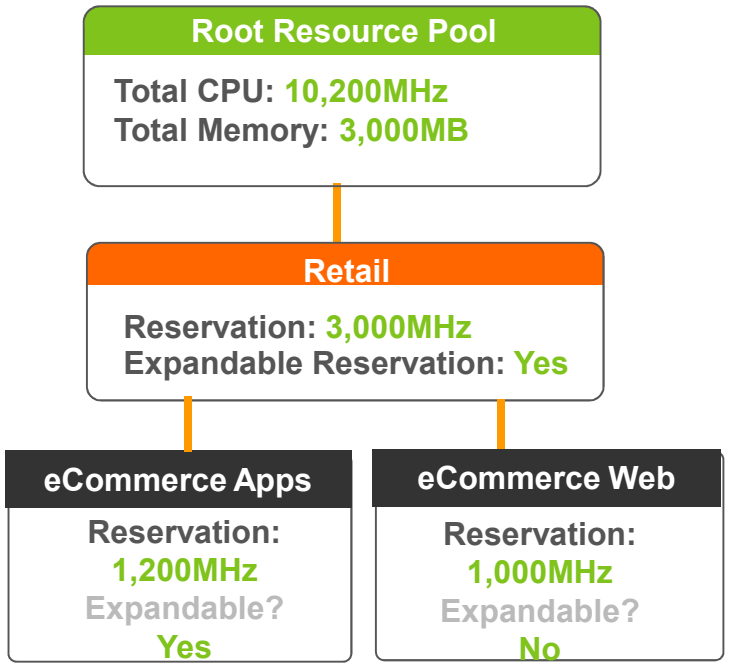

Borrowing resources occurs recursively from the ancestors of the current resource pool.

- As long as **Expandable Reservation** is selected
- Offers more flexibility but less protection

Expanded reservations are not released until the virtual machine that caused the expansion is shut down or its reservation is reduced.

An expandable reservation could allow a rogue administrator to claim all unreserved capacity in the environment.

Example of Expandable Reservation (1)

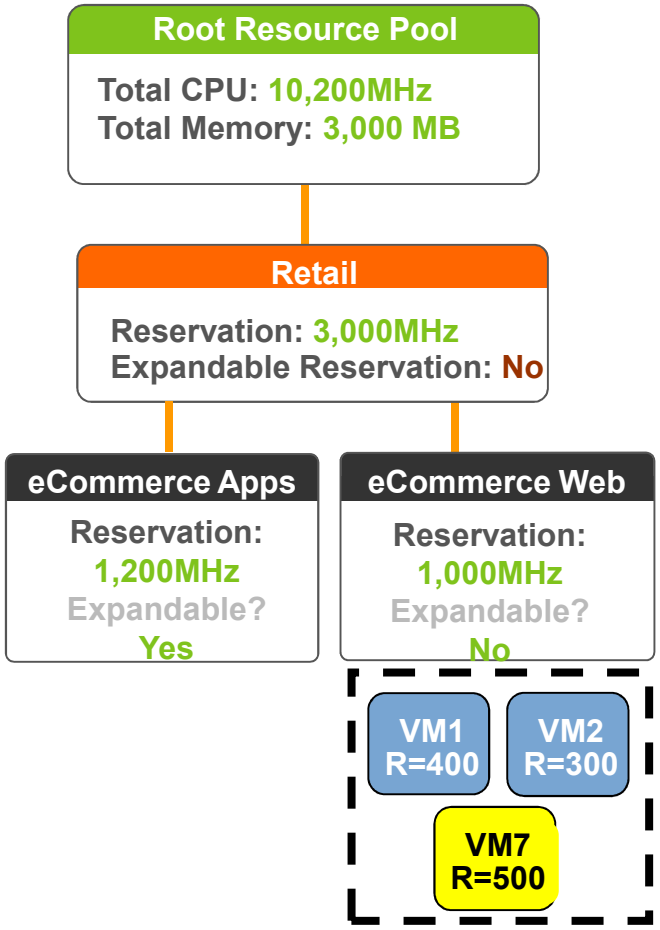

eCommerce resource pools reserve 2,200MHz of 3,000MHz the Retail resource pool has reserved.

Power on virtual machines in the eCommerce Web resource pool.

With expandable reservation disabled on the eCommerce Web resource pool, it is not possible to start VM7 with a reservation of 500MHz.

- Lower the virtual machine reservation.
- Select **Expandable Reservation**.
- Increase eCommerce Web pool's reservation.

Example of Expandable Reservation (2)

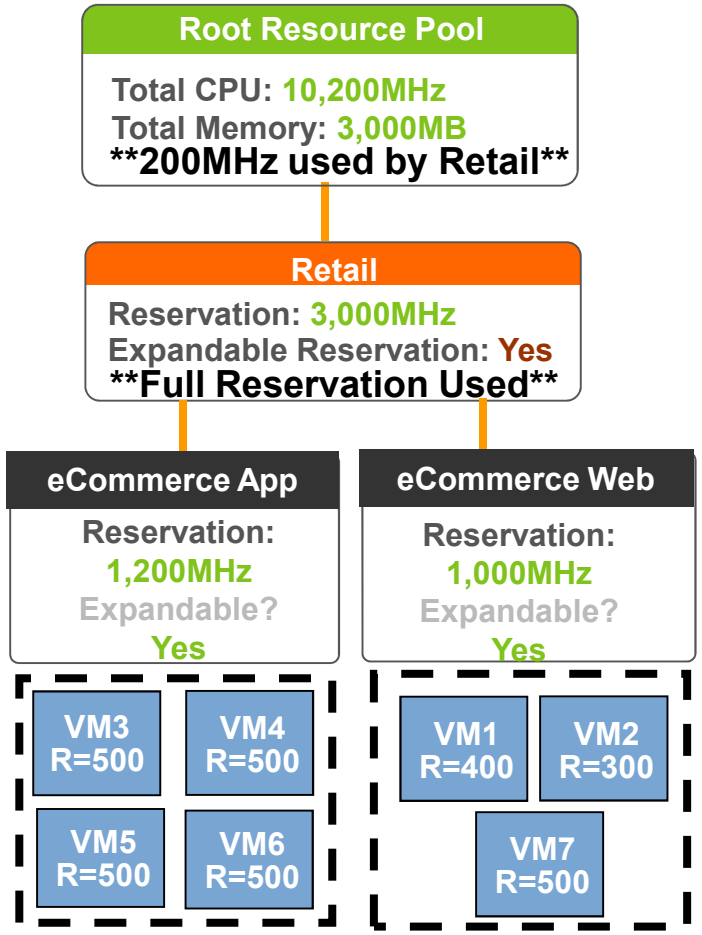

Enable expandable reservation on the eCommerce Web resource pool.

The system considers the resources available in the child resource pool and its direct parent resource pool.

The virtual machine's reservation is charged against the reservation for eCommerce Web.

eCommerce Web's reservation is charged against the reservation for Retail.

Creating a Resource Pool

Right-click the host, then choose New Resource Pool**.**

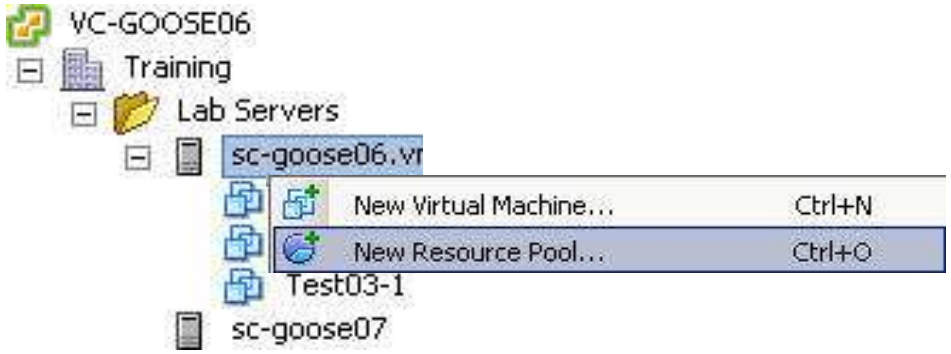

Drag virtual machines into the resource pool.

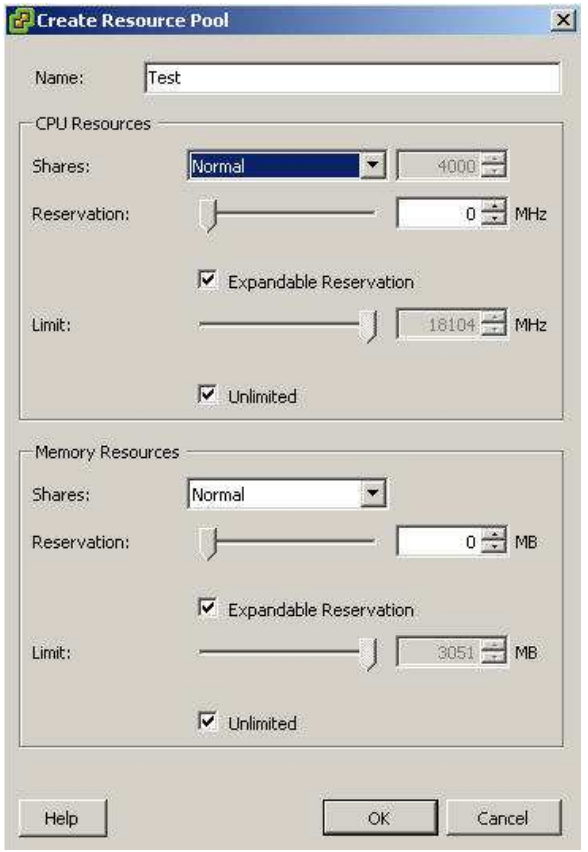

Admission Control for CPU & Memory Reservations

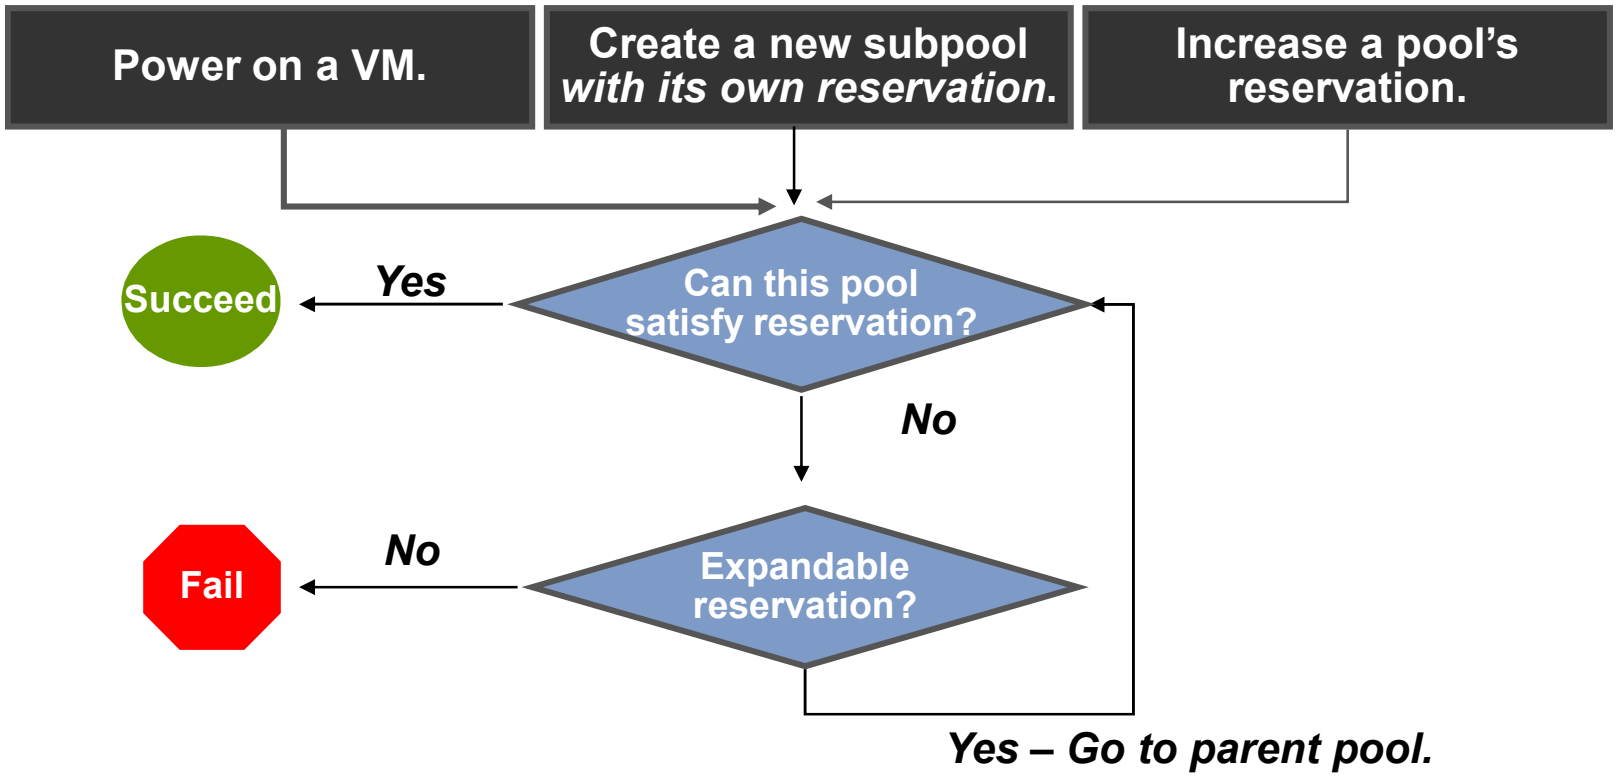

Resource Pool Summary Tab

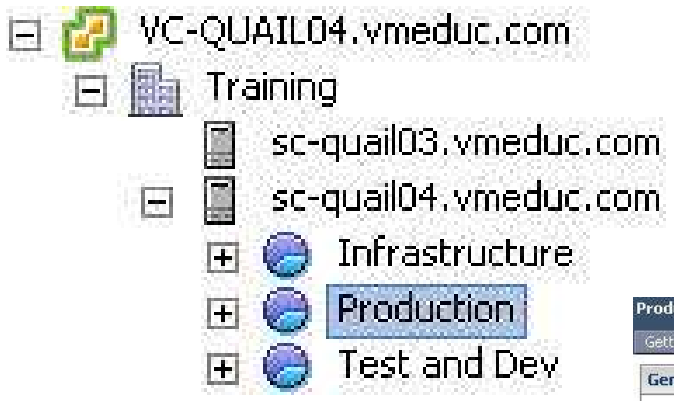

Display the resource pool's Summary **tab.**

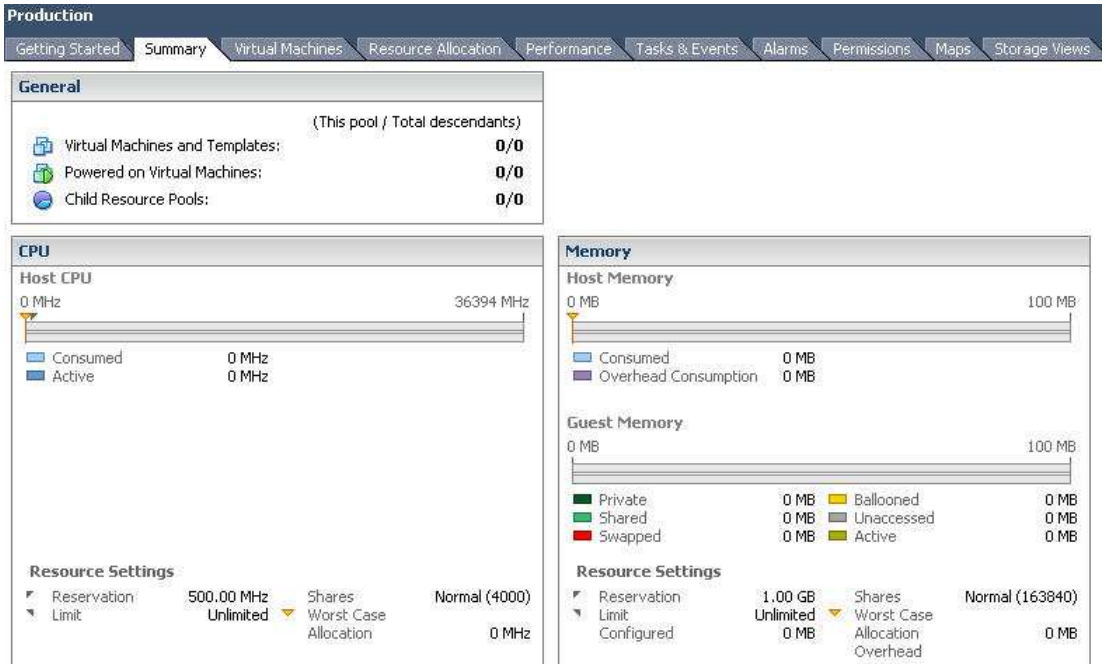

Resource Allocation Tab

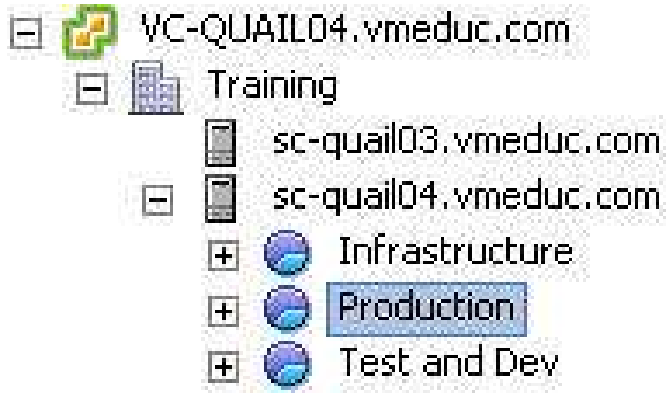

Display the resource pool's Resource Allocation **tab.**

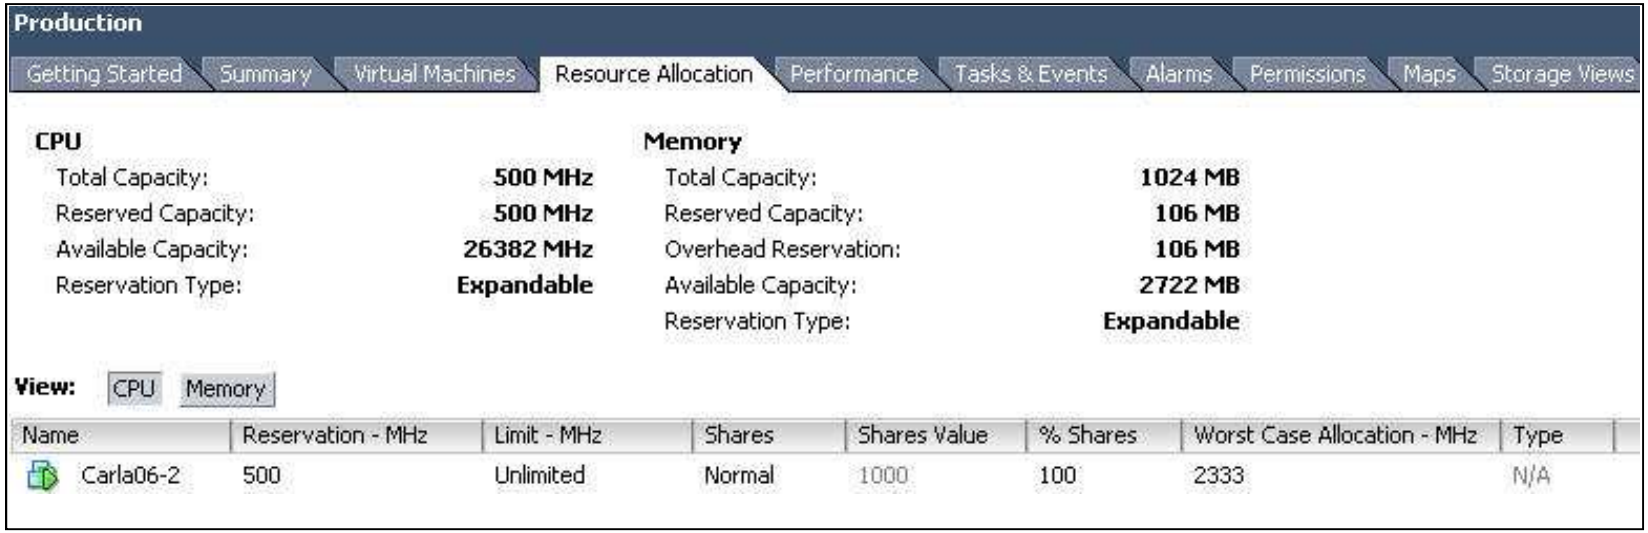

Scheduling Changes to Resource Settings

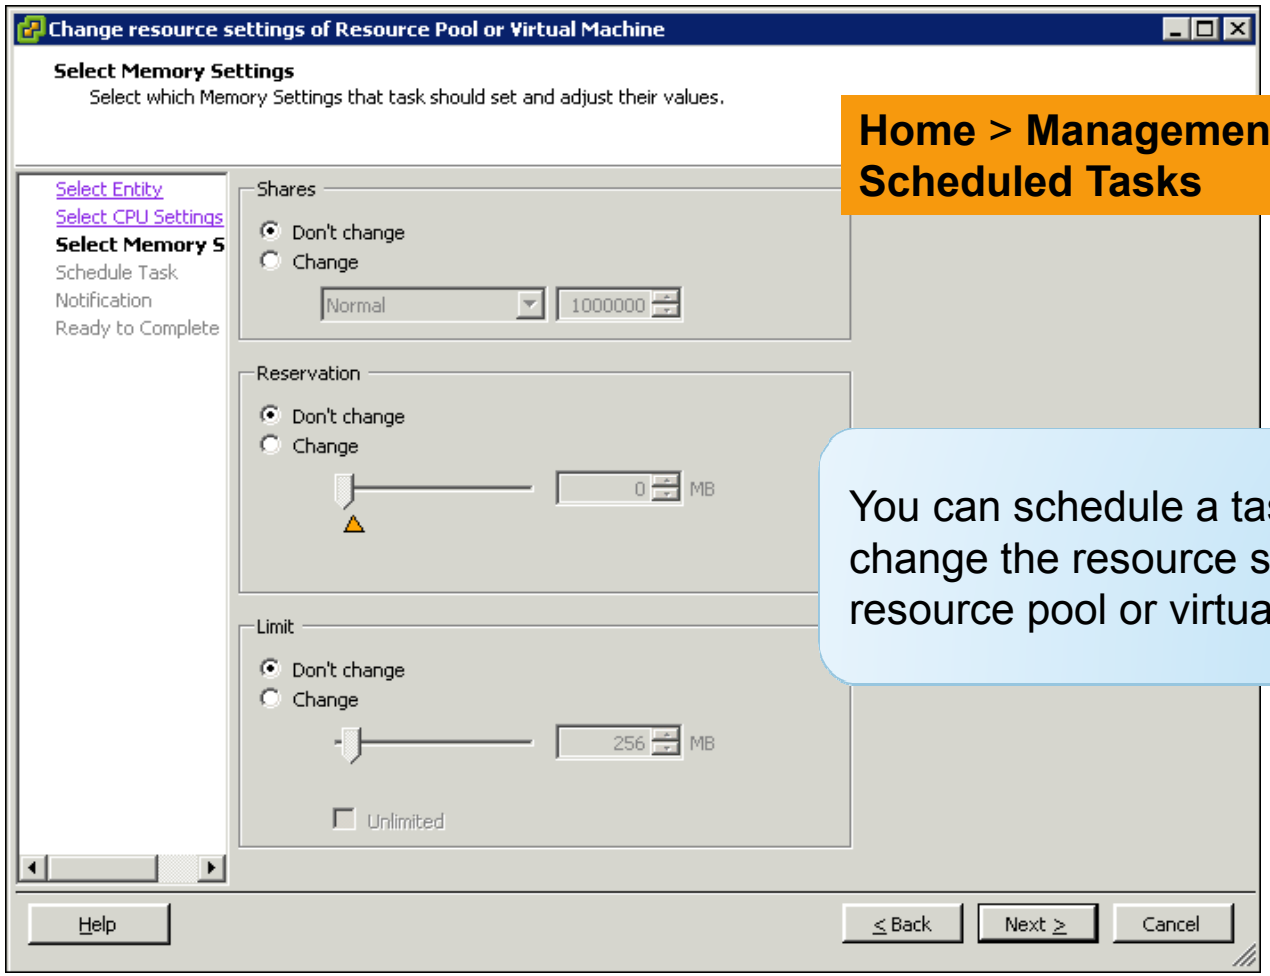

Home > Management > Scheduled Tasks

You can schedule a task to change the resource settings of a resource pool or virtual machine.

Lab 17

In this lab, you will create and use resource pools on an ESX host.

1. Create resource pools.
2. Verify resource pool functionality.

Lesson Summary

- A virtual machine's CPU and memory allocation can be controlled using a combination of limits, reservations, and shares.
- Use resource pools for a flexible hierarchical organization of CPU and memory resources.
- Expandable reservation is a resource pool attribute that allows virtual machines and subpools to use resources from the pool's parent, if necessary.

Lesson 2: Scaling Storage and Network Management

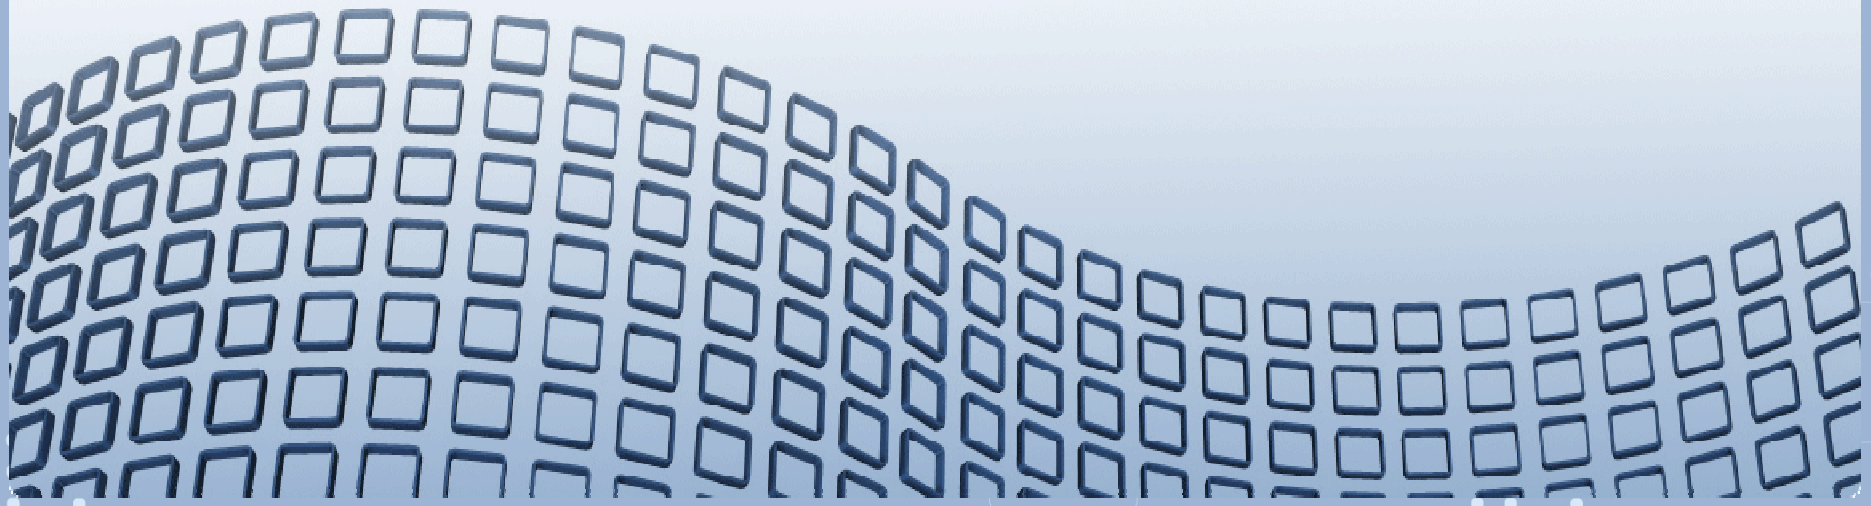

Lesson Objectives

- Describe the methods for scaling storage management
  - Multipathing
    - Load balancing
- Describe the Pluggable Storage Architecture (PSA)
- Describe the methods for scaling network management
  - NIC teaming
  - Multipathing for iSCSI storage

Storage Multipathing

Multipathing allows continued access to SAN LUNs in the event of hardware failure. It also provides load balancing.

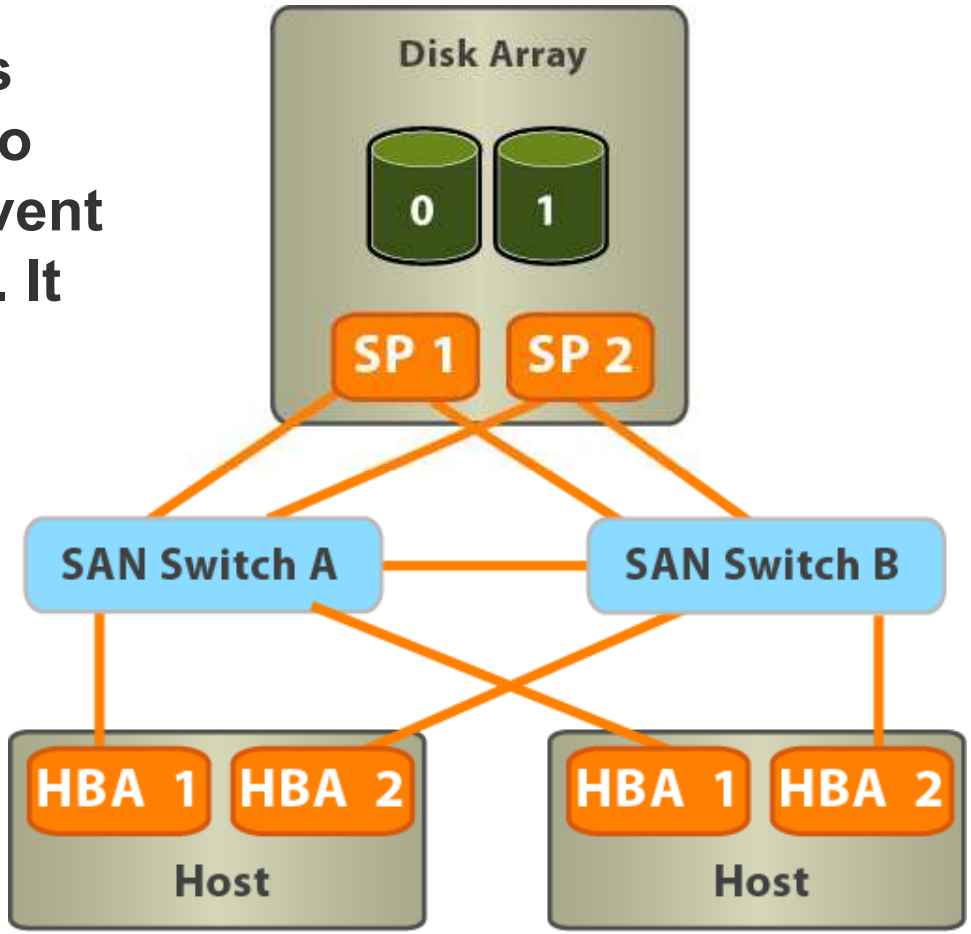

Managing Multiple Storage Paths

Multiple paths can exist to a datastore on the host.

To modify storage path information, click the datastore's Properties **link.**

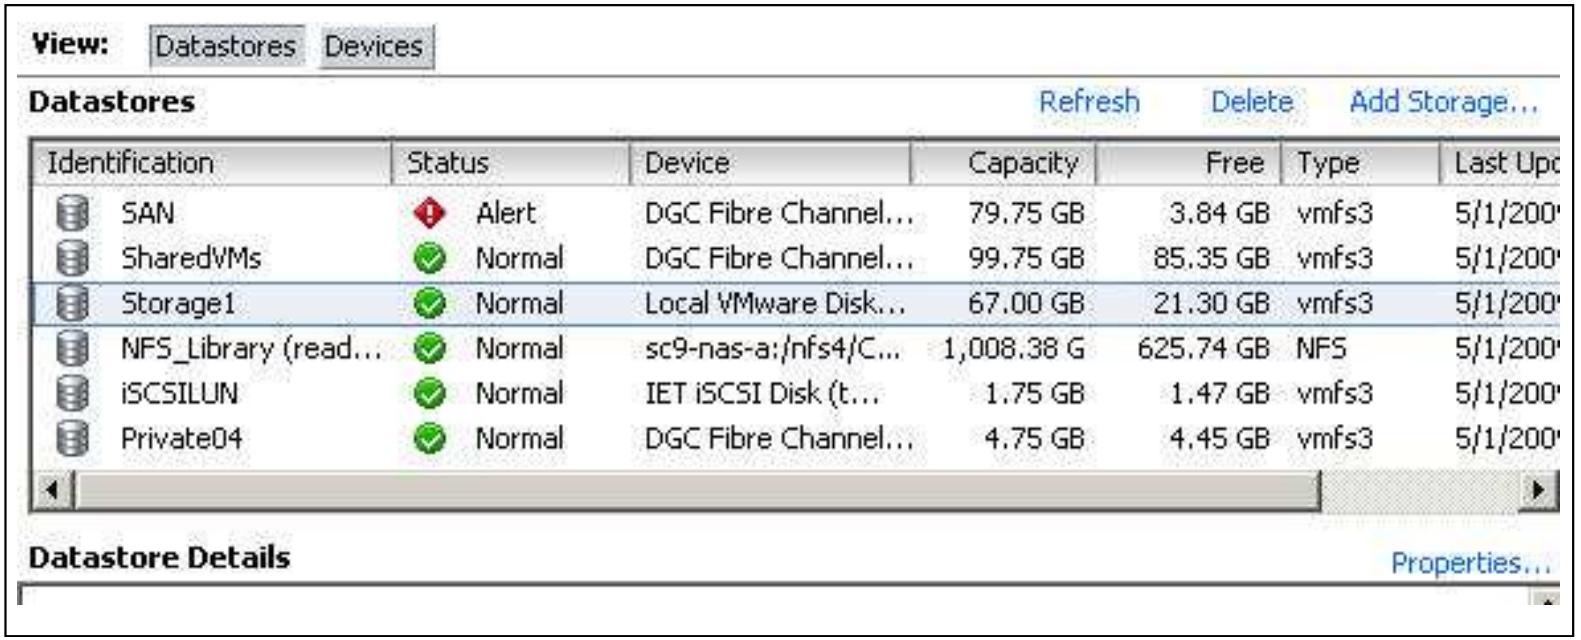

Configuring Storage Load Balancing

Path selection policies exist for:

- Scalability
  - Round Robin – a multipathing policy that performs load balancing across paths.
- Availability
  - MRU and Fixed (discussed in a later module)

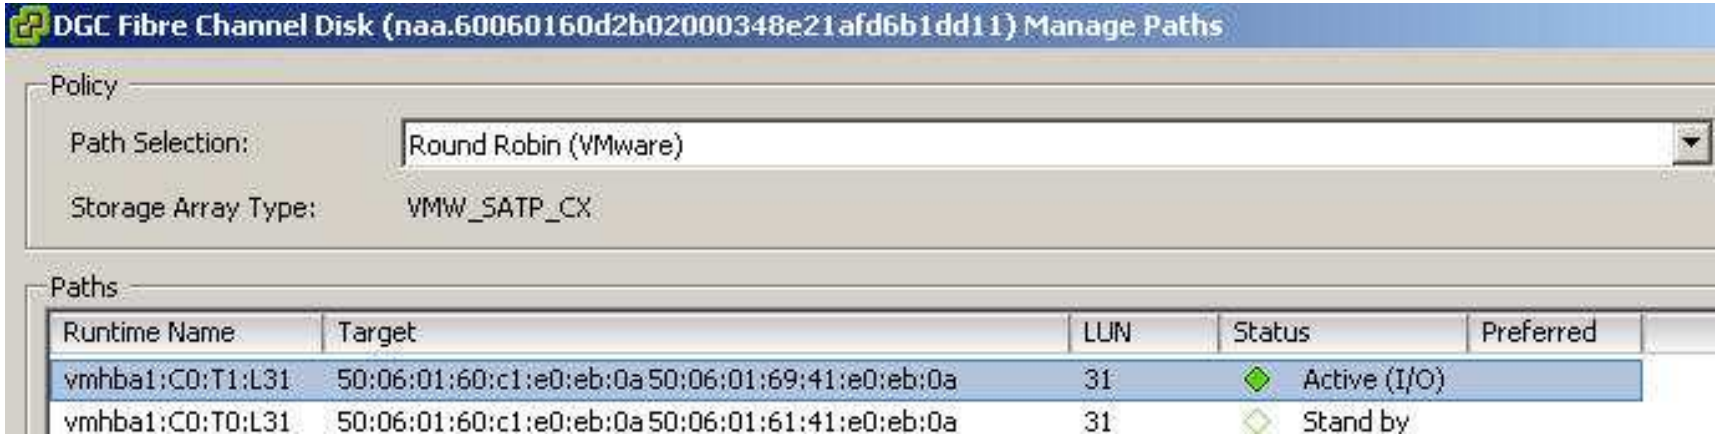

Pluggable Storage Architecture

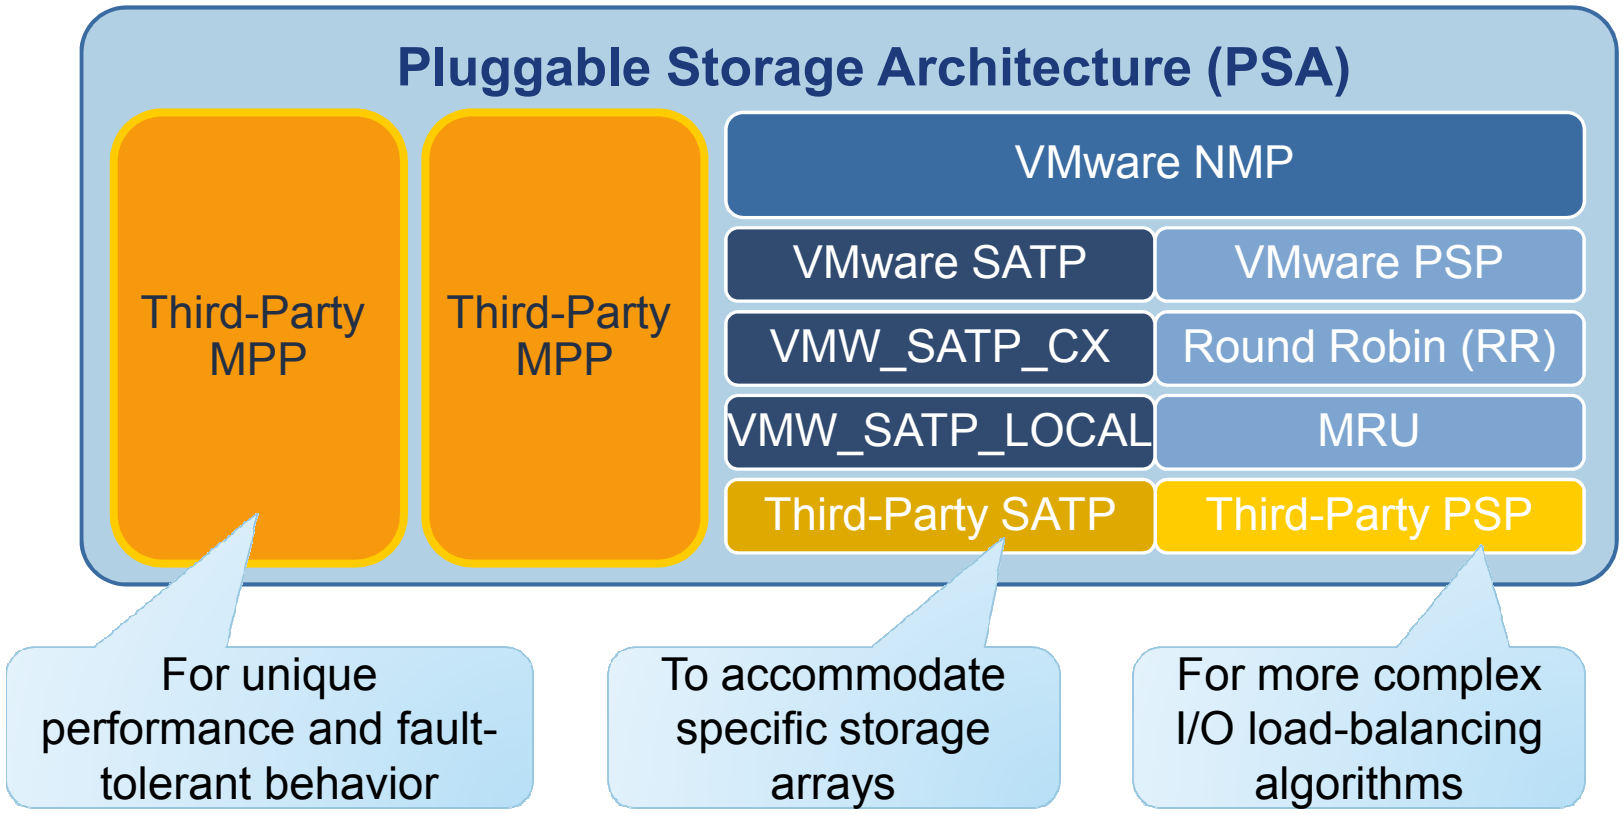

Configuring NIC Teaming

NIC teaming

- Provides multipathing for ESX/ESXi networks
- Occurs when multiple uplinks are associated with a single vSwitch
- Can share the network traffic load or provide NIC failover

To configure, add network adapters to the virtual switch.

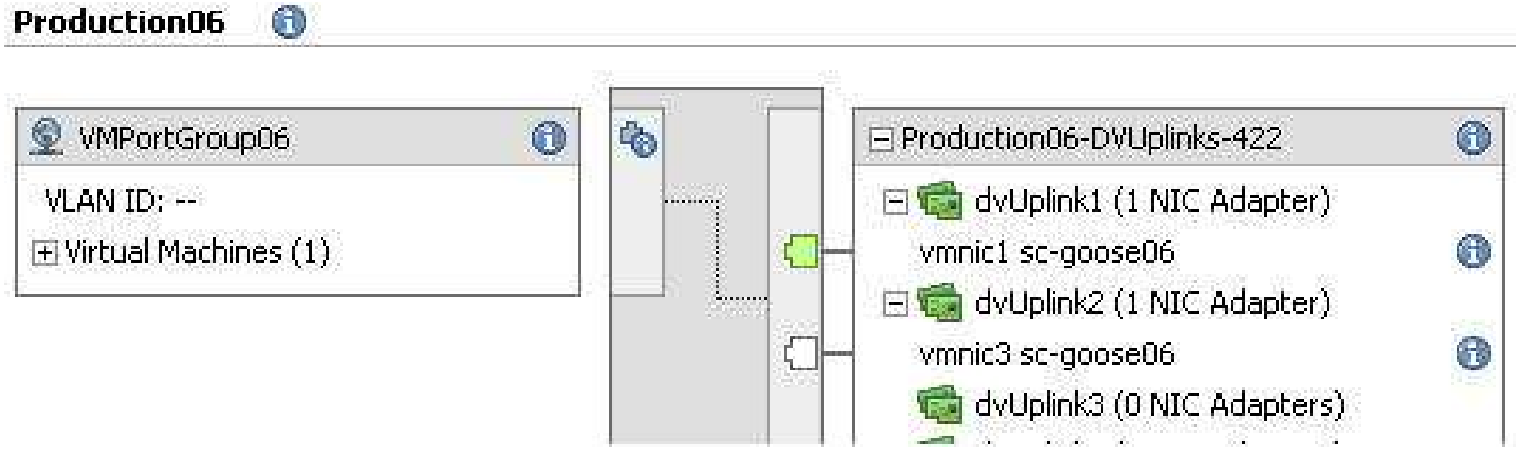

Configuring Network Load Balancing

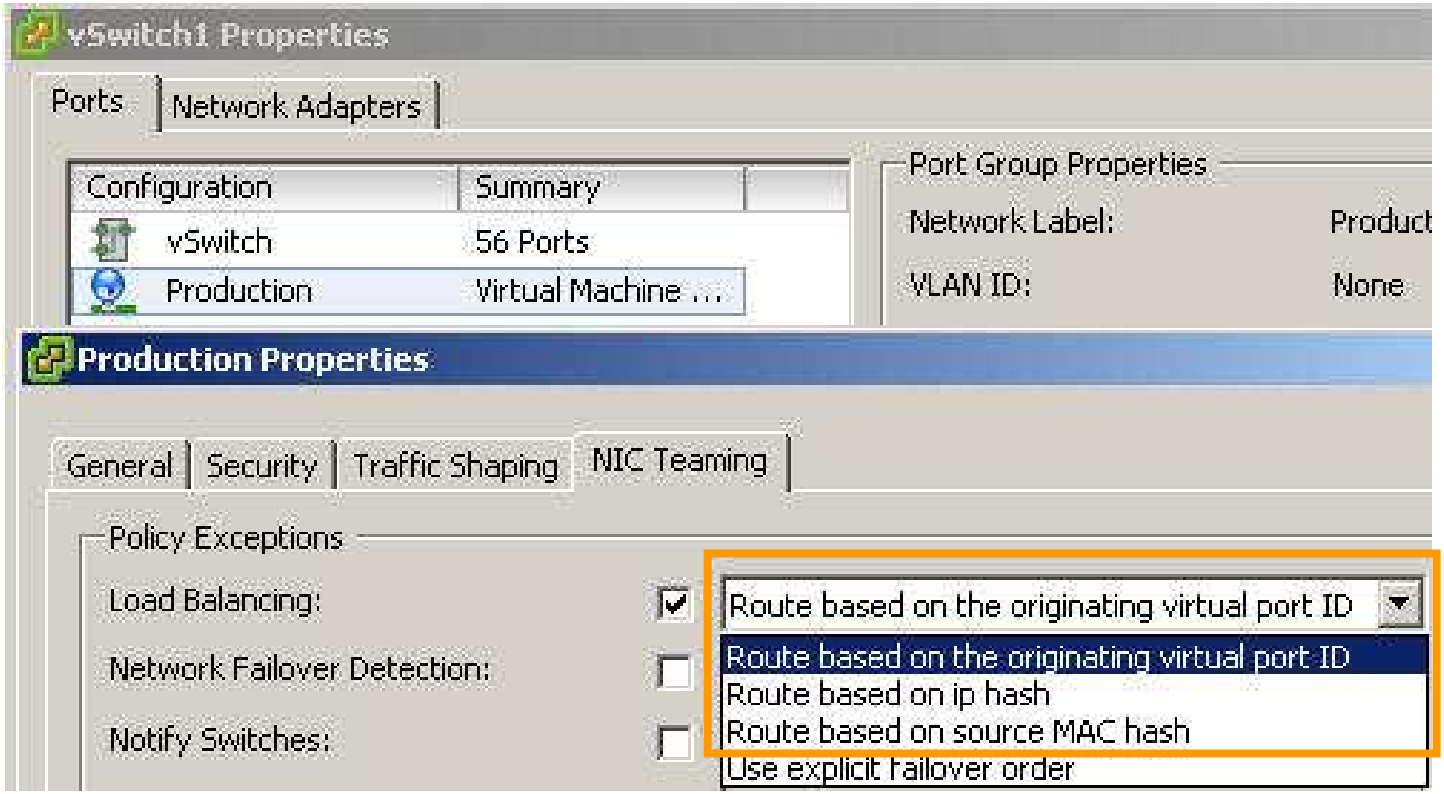

The NIC Teaming configuration allows you to set a load-balancing policy.

Load-Balancing Method: Port-ID Based

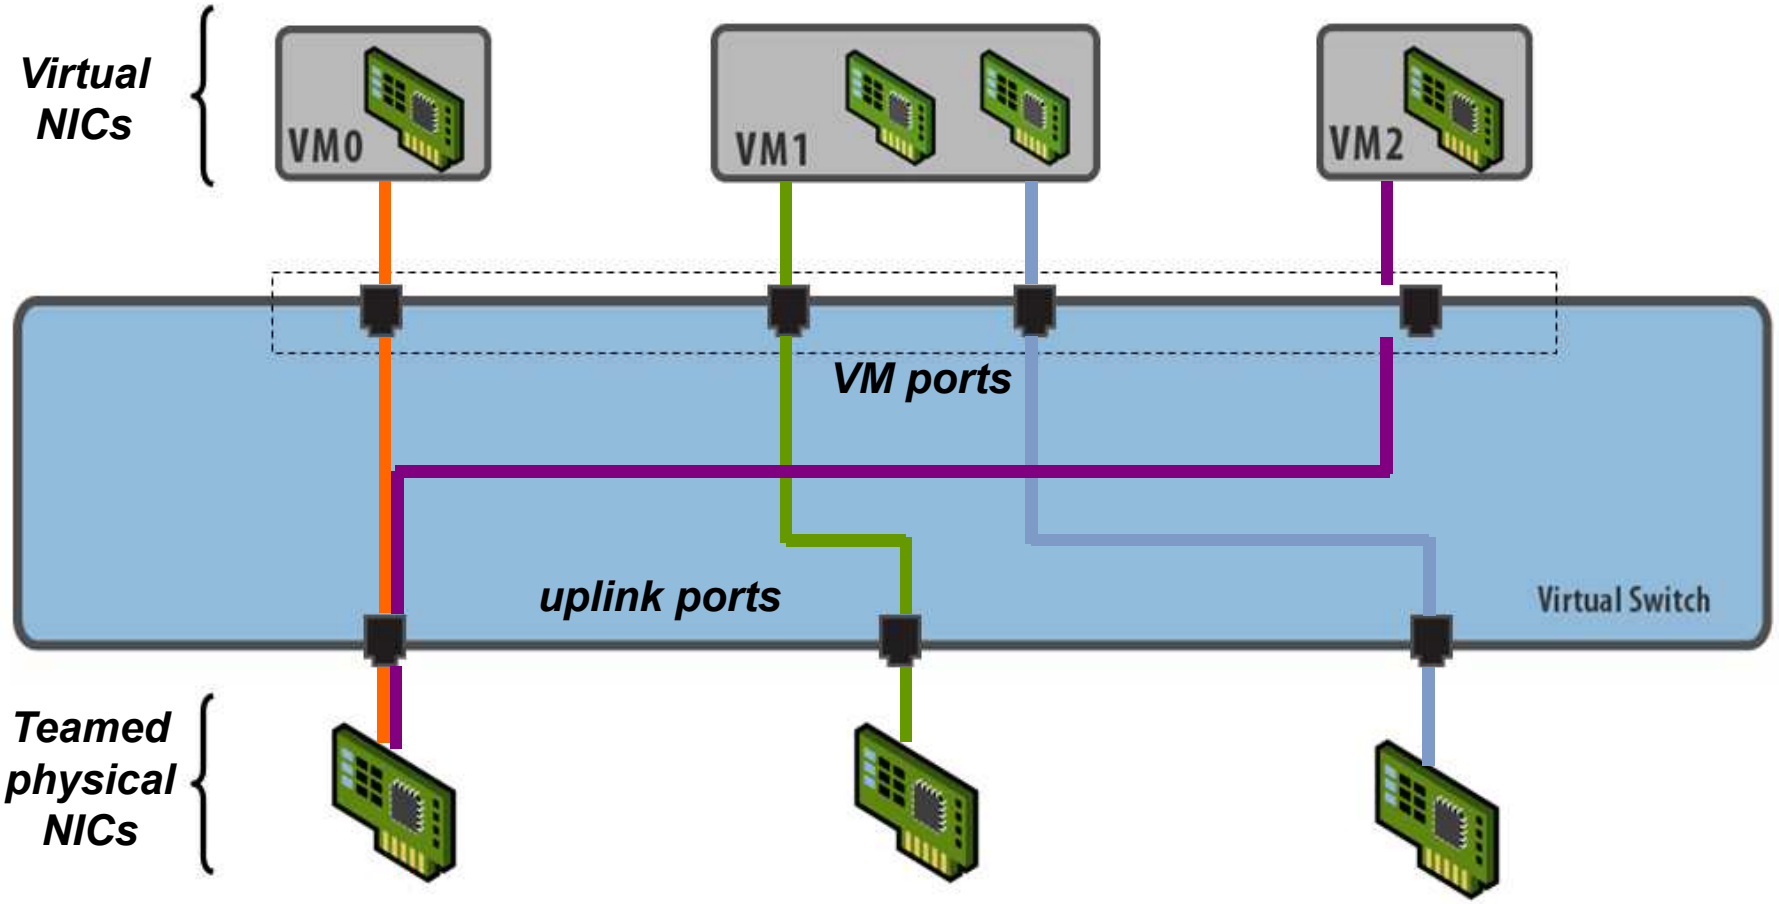

Load-Balancing Method: Source MAC-Based

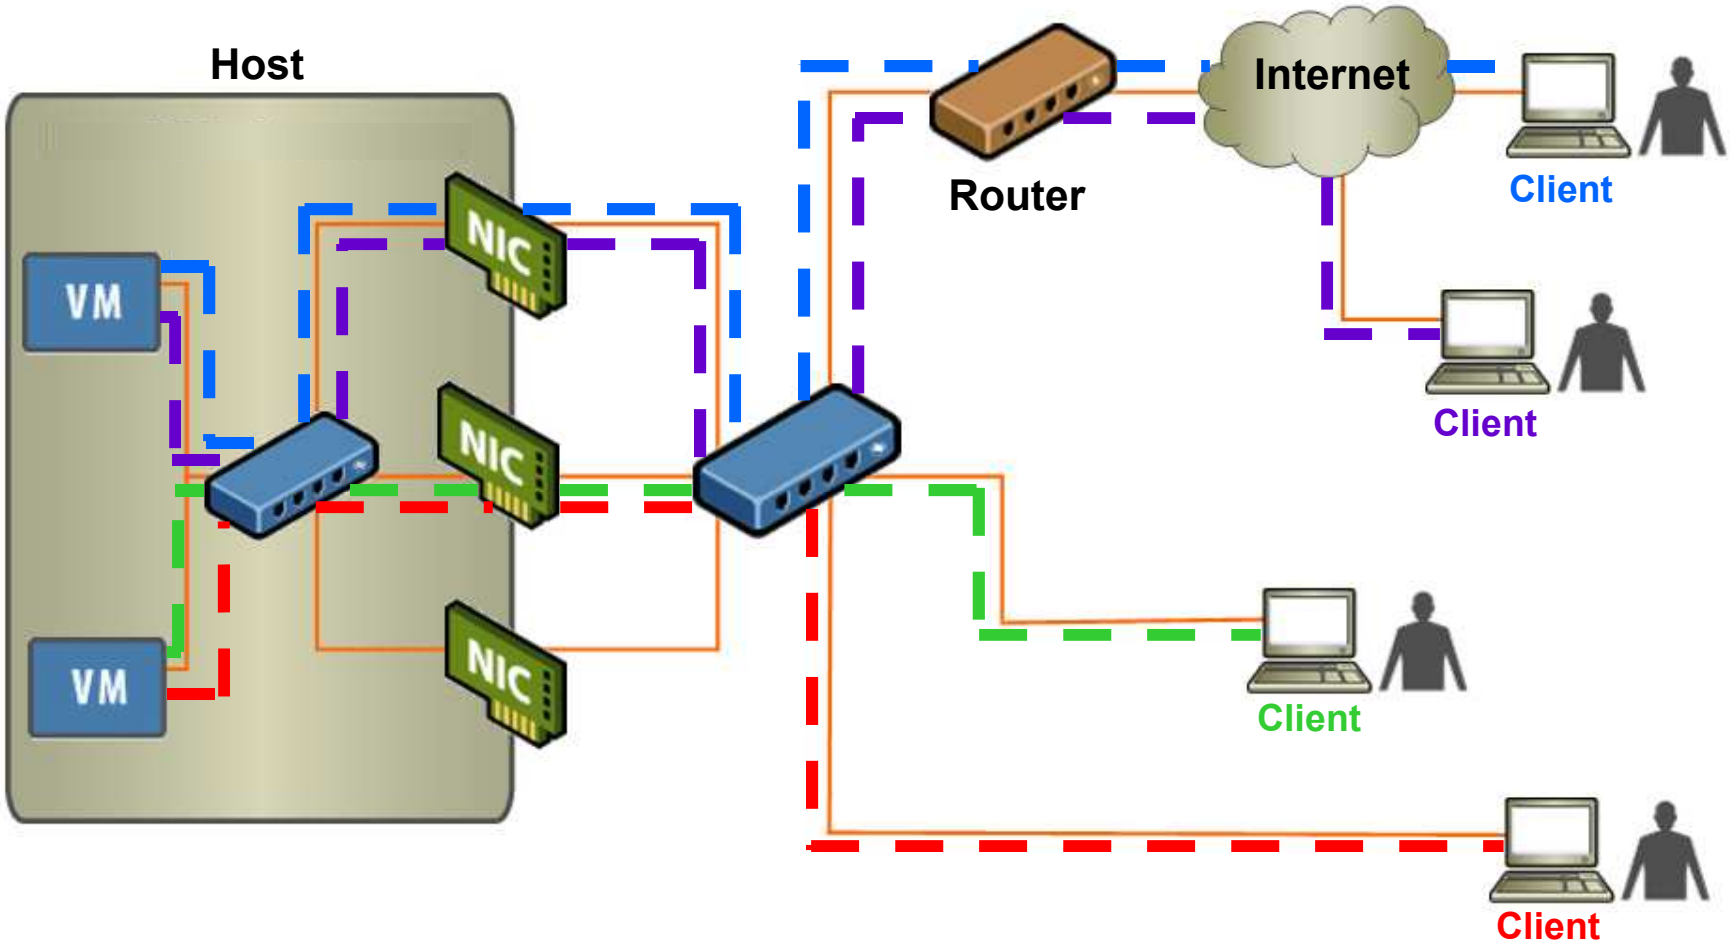

Load-Balancing Method: IP-Based

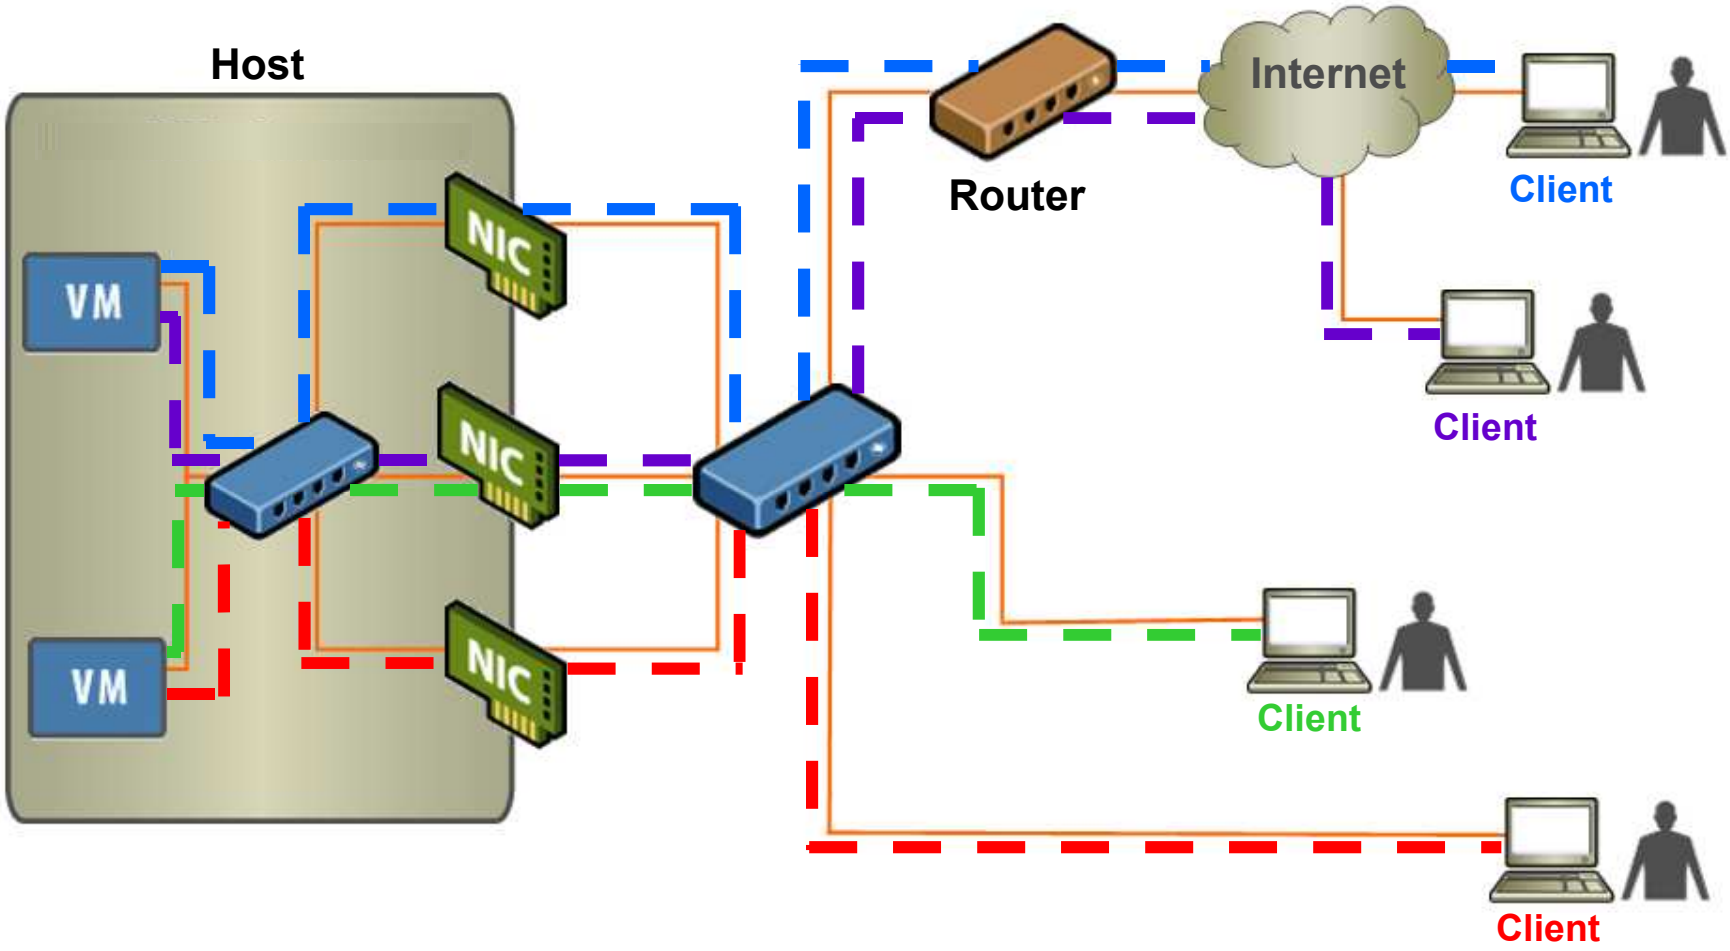

Multipathing with iSCSI Storage

SendTargets advertises multiple routes.

- It reports different IP addresses to allow different paths to the iSCSI LUNs.

Routing done via IP network.

For the software initiator:

- This counts as one network interface.
- NIC teaming and multiple SPs allow for multiple paths.

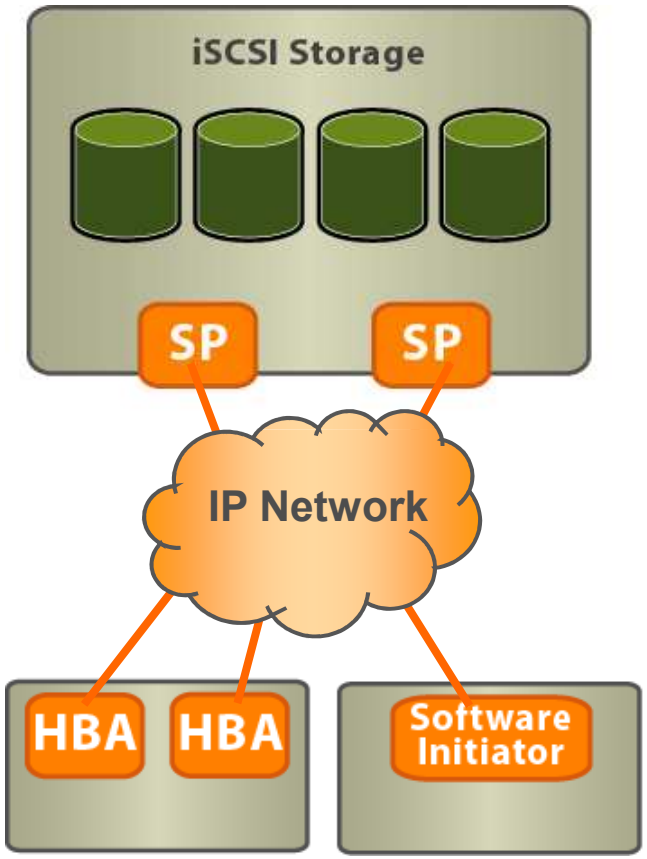

Lesson Summary

- The Round Robin multipathing policy performs load balancing across multiple paths to the datastore.
- Pluggable Storage Architecture (PSA) is a VMkernel layer responsible for managing multiple paths, load balancing across paths, and path failover.
- Three NIC teaming policies exist for load balancing: Using the virtual port ID, using the IP hash, and using the source MAC hash.

Lesson 3:
VMware VMotion Migration

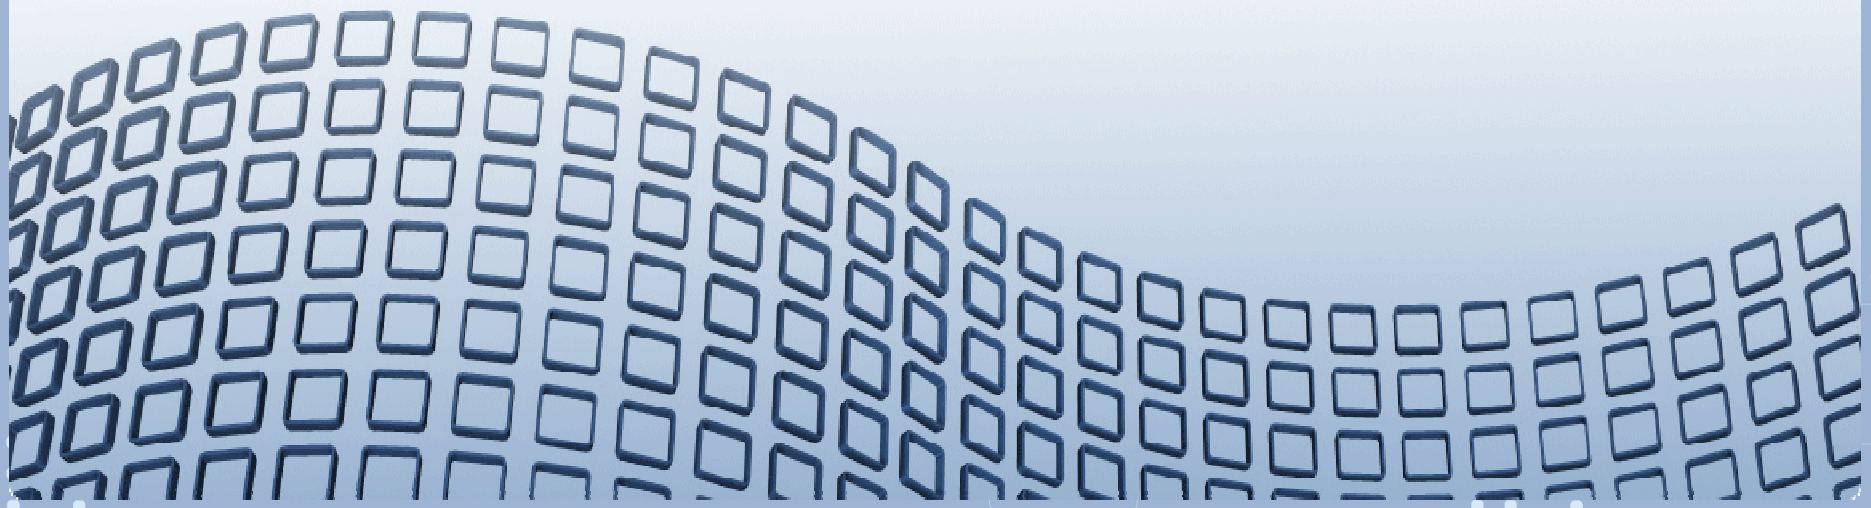

Lesson Objectives

- Understand the importance of VMware VMotion™
- Identify VMotion requirements
  - Virtual machine
  - Host
- Verify VMotion requirements
- Perform a VMotion migration

VMotion Migration

A VMotion migration moves a powered-on virtual machine from one host to another.

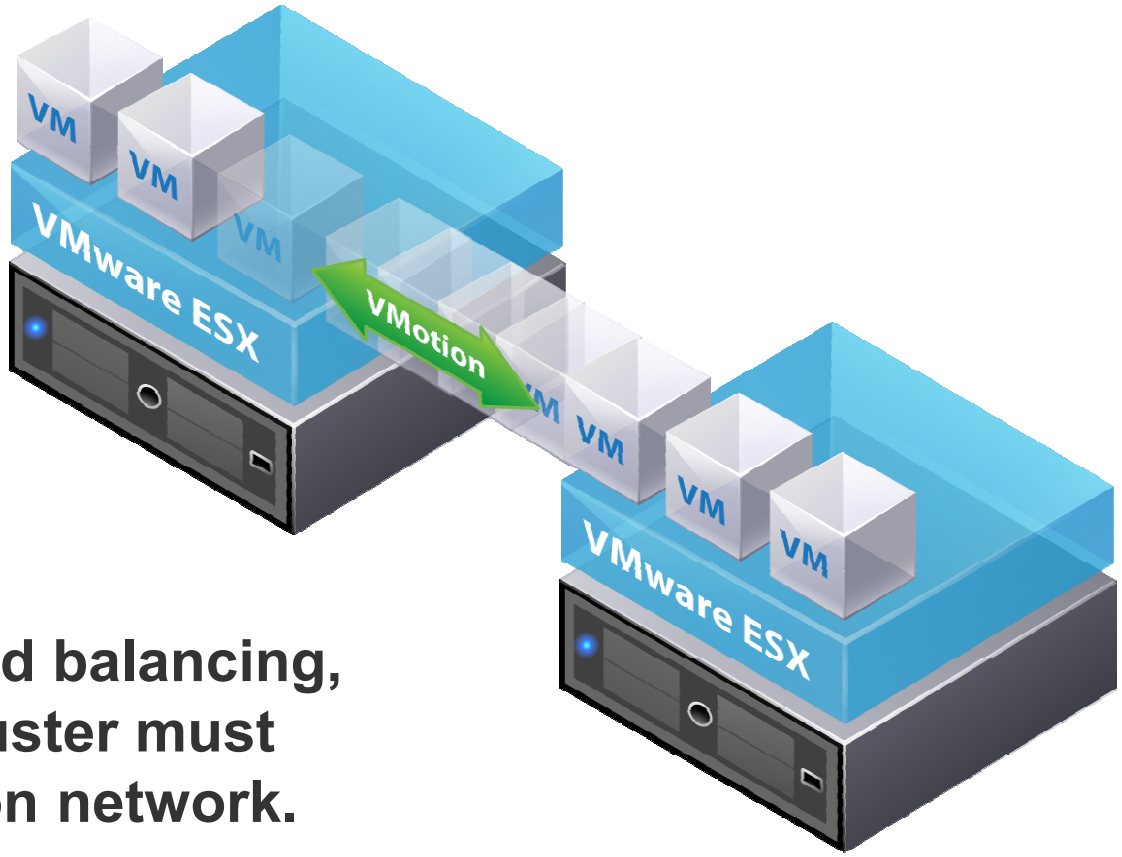

To use DRS for load balancing, the hosts in the cluster must be part of a VMotion network.

How VMotion Works

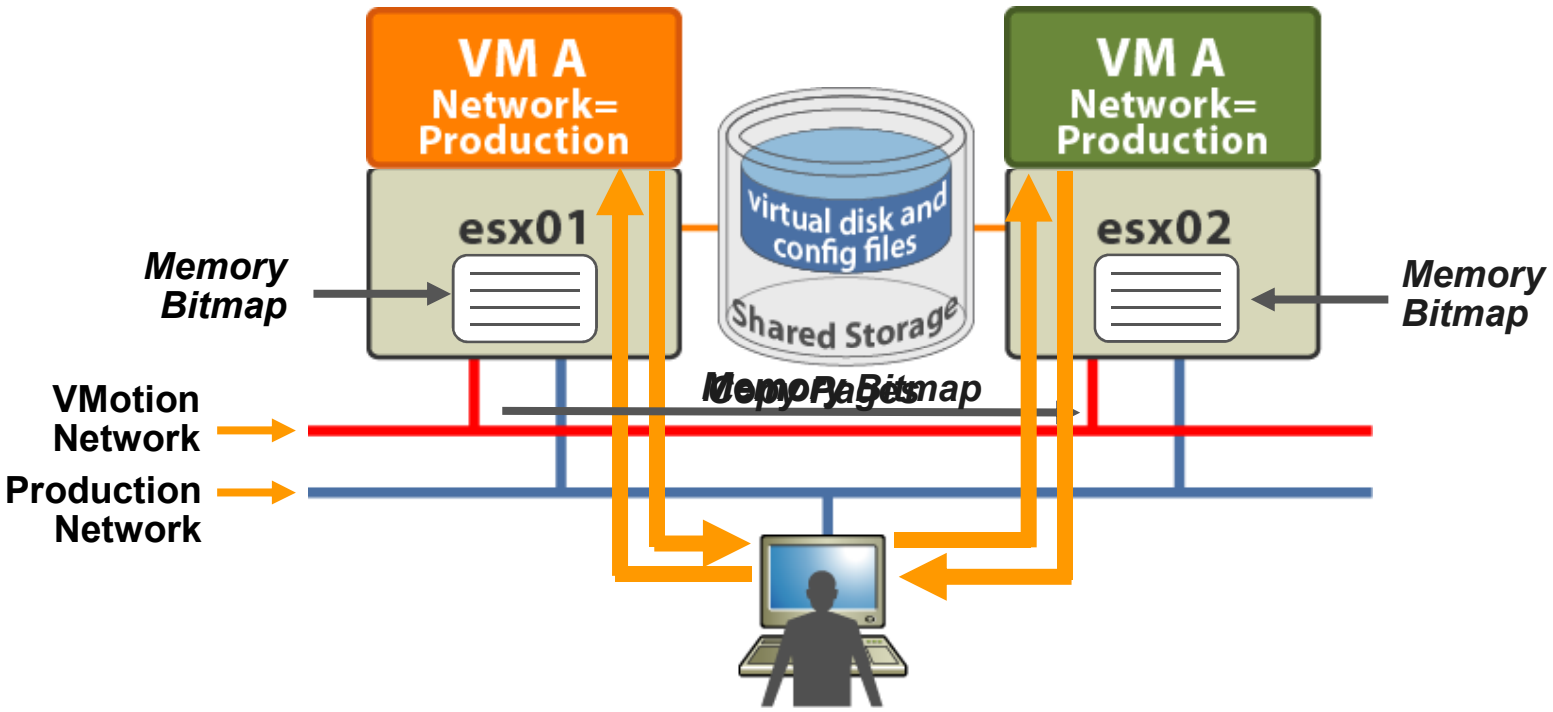

Virtual Machine Requirements for VMotion

A virtual machine must meet the following requirements:

- A virtual machine must not have a connection to an internal vSwitch (vSwitch with zero uplink adapters).
- A virtual machine must not have a connection to a virtual device (such as CD-ROM or floppy drive) with a local image mounted.
- A virtual machine must not have CPU affinity configured.
- A virtual machine must not be in an MSCS cluster relationship with another virtual machine.
- A virtual machine's swap file must be accessible by the destination host.
- If a virtual machine uses an RDM, the RDM must be accessible by the destination host.

Host Requirements for VMotion

Source and destination ESX hosts must have:

- Visibility to all storage (Fibre Channel, iSCSI, or NAS) used by the virtual machine
- A Gigabit Ethernet backplane
- Access to the same physical networks
- Compatible CPUs
  - CPU feature sets of both the source and destination host must be compatible.
  - Some features can be hidden using Enhanced VMotion Compatibility (EVC) or compatibility masks.

CPU Constraints on VMotion

CPU characteristics	Exact match required?	Why or why not?
Clock speeds, cache sizes, hyperthreading, and number of cores	No	Virtualized away by VMkernel
Manufacturer *(Intel or AMD)* Family *(P3, P4, Opteron)*	Yes	Instruction sets contain many small differences.
Presence or absence of SSE3, SSSE3, or SSE4.1 instructions	Yes	Multimedia instructions usable directly by applications
Virtualization hardware assist	For 32-bit VMs: No	Virtualized away by VMkernel
	For 64-bit VMs on Intel: Yes	VMware's Intel 64-bit implementation leverages VT.
Execution-disable (Nx/Xd bit)	Yes (but customizable)	Guest operating system relies on NX/XD bit if detected.

Exposing or Hiding NX/XD

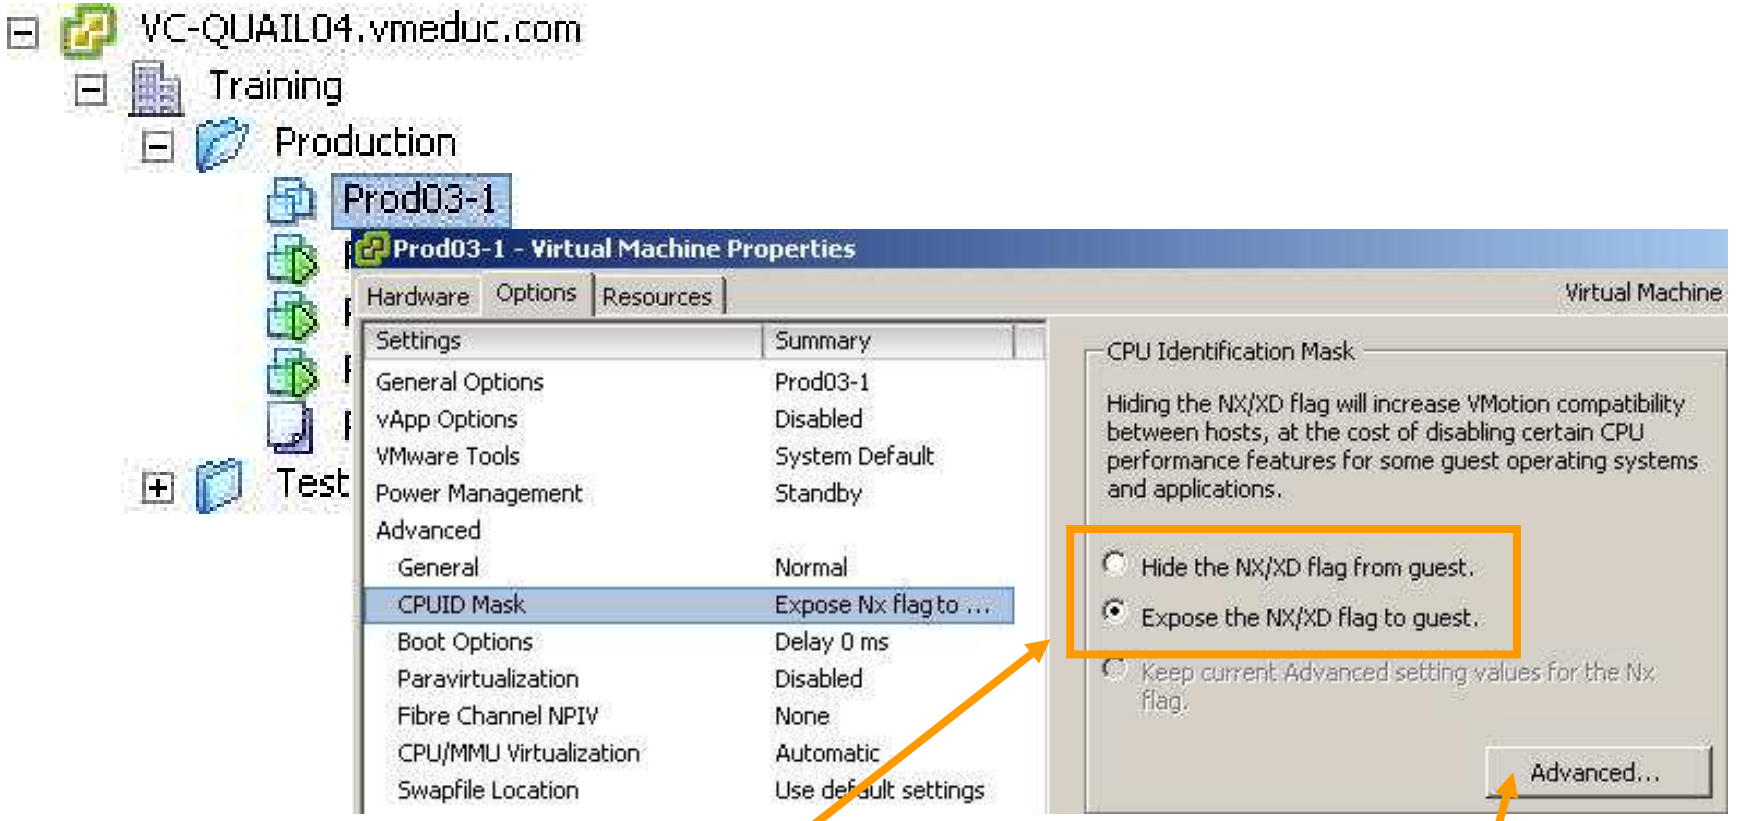

Choose between NX/XD security features or broadest VMotion compatibility.

For future CPU features, edit mask at the bit level.

Identifying CPU Characteristics

```
Random_Init: Using random seed: 2044292605 (0x79d96dfd)
Reporting CPUID for 2 logical CPUs...

All CPUs are identical

     Family: 06 Model: 17 Stepping: 6

     ID1ECX     ID1EDX     ID81ECX    ID81EDX
     0x00082201 0x0febfbff 0x00000001 0x20100000

Vendor                          : Intel
Brand String                    : "Intel(R) Xeon(R) CPU           X5482  @ 3.20GHz"
SSE Support                     : SSE1, SSE2, SSE3, SSSE3, SSE4.1
Supports NX / XD                : Yes
Supports CMPXCHG16B             : Yes
Supports RDTSCP                 : No
Hyperthreading                  : No
Supports Flex Migration         : Yes
Supports 64-bit Longmode        : Yes
Supports 64-bit VMware          : No
Supported EVC modes             : None

PASS: Test 56983: CPUID
Press any key to reboot.
```

In most cases, use server and CPU family/model specifications.

Use VMware bootable CPUID utility.

Verifying VMotion Layout: Custom Maps

Verifying VMotion Layout: Virtual Machine Map

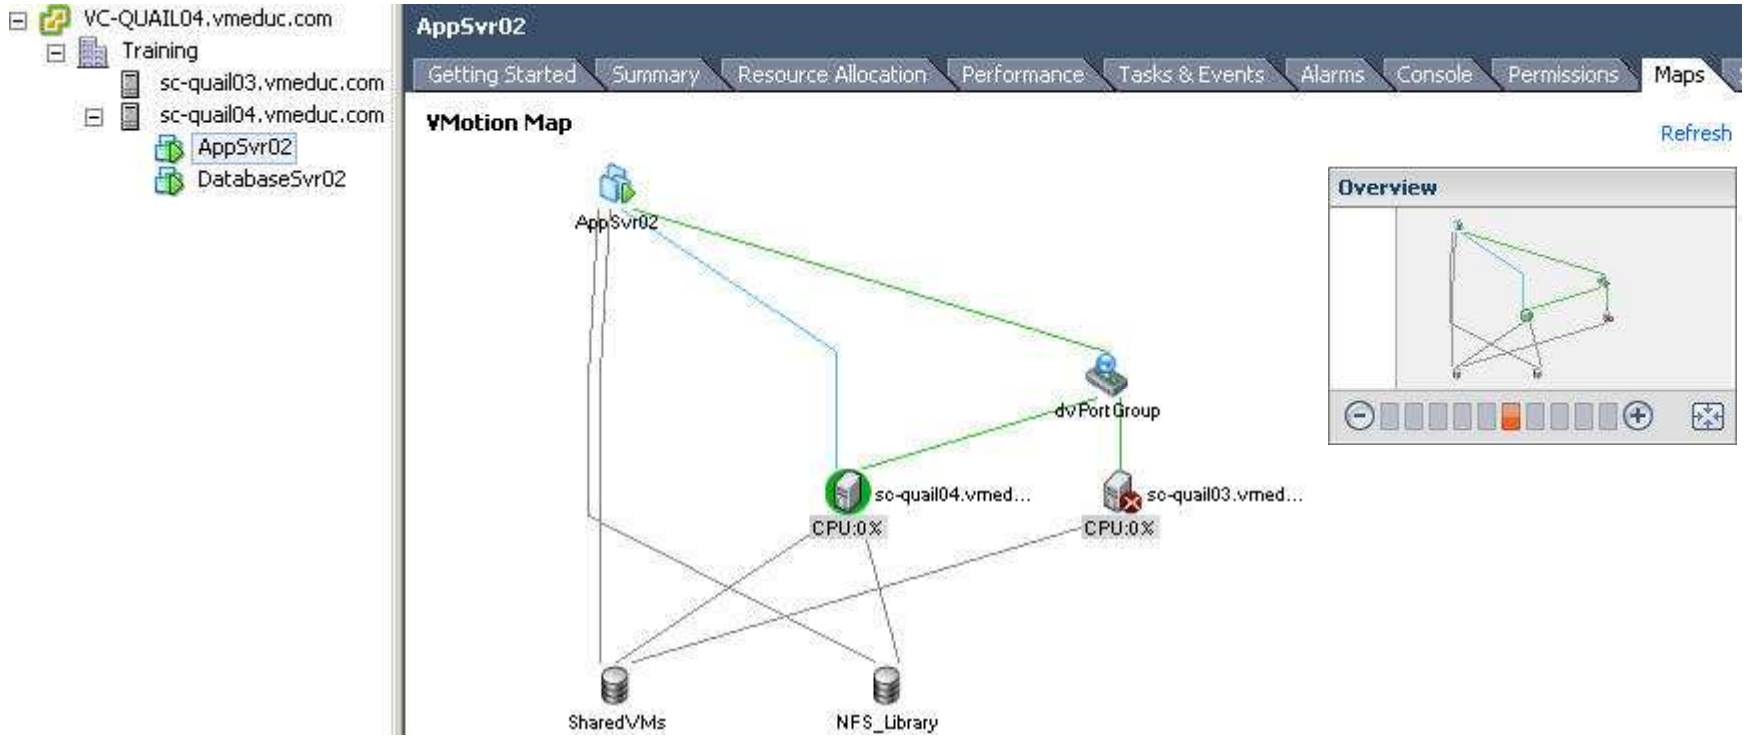

Performing a VMotion Migration

From the VMware vCenter™ Server inventory, right-click a virtual machine that is powered on, then choose Migrate.

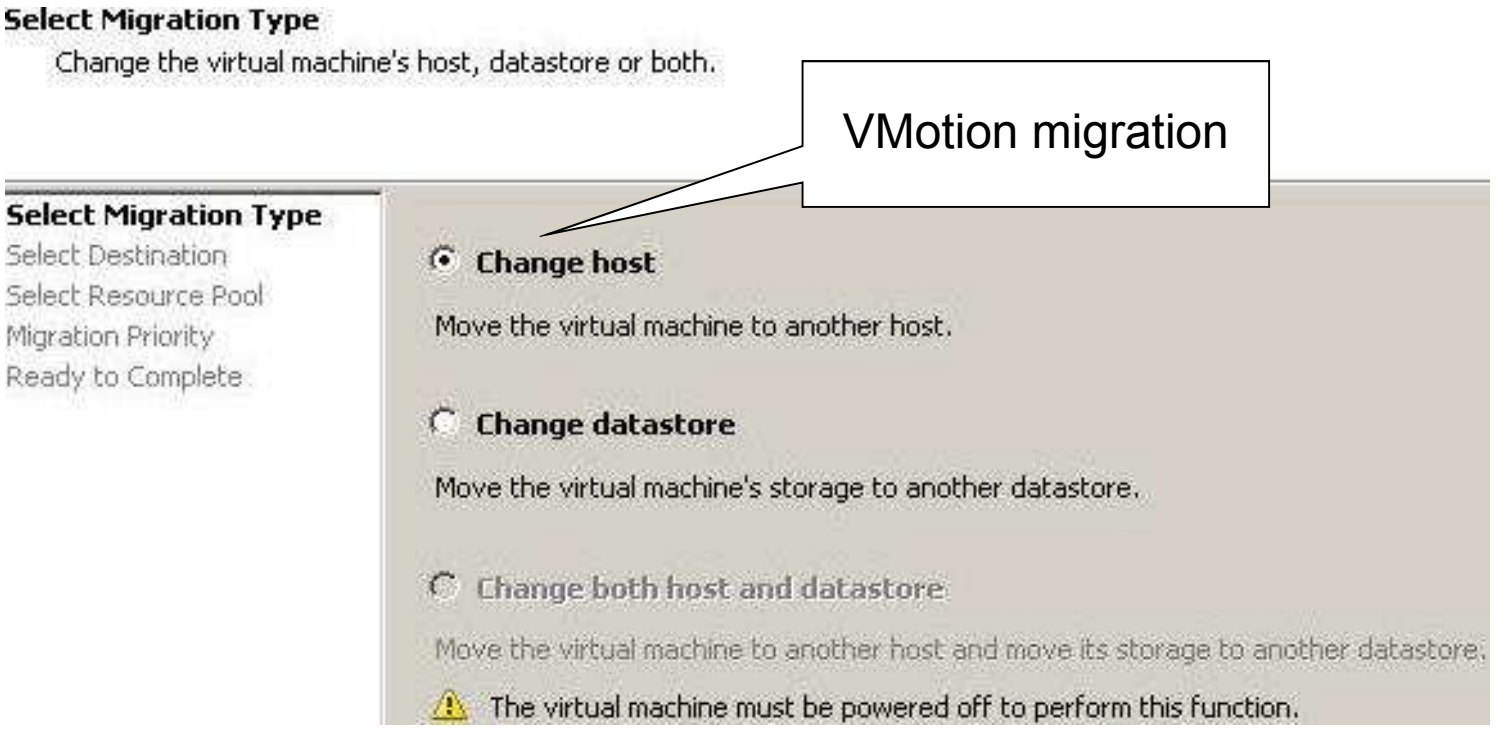

Checking VMotion Errors

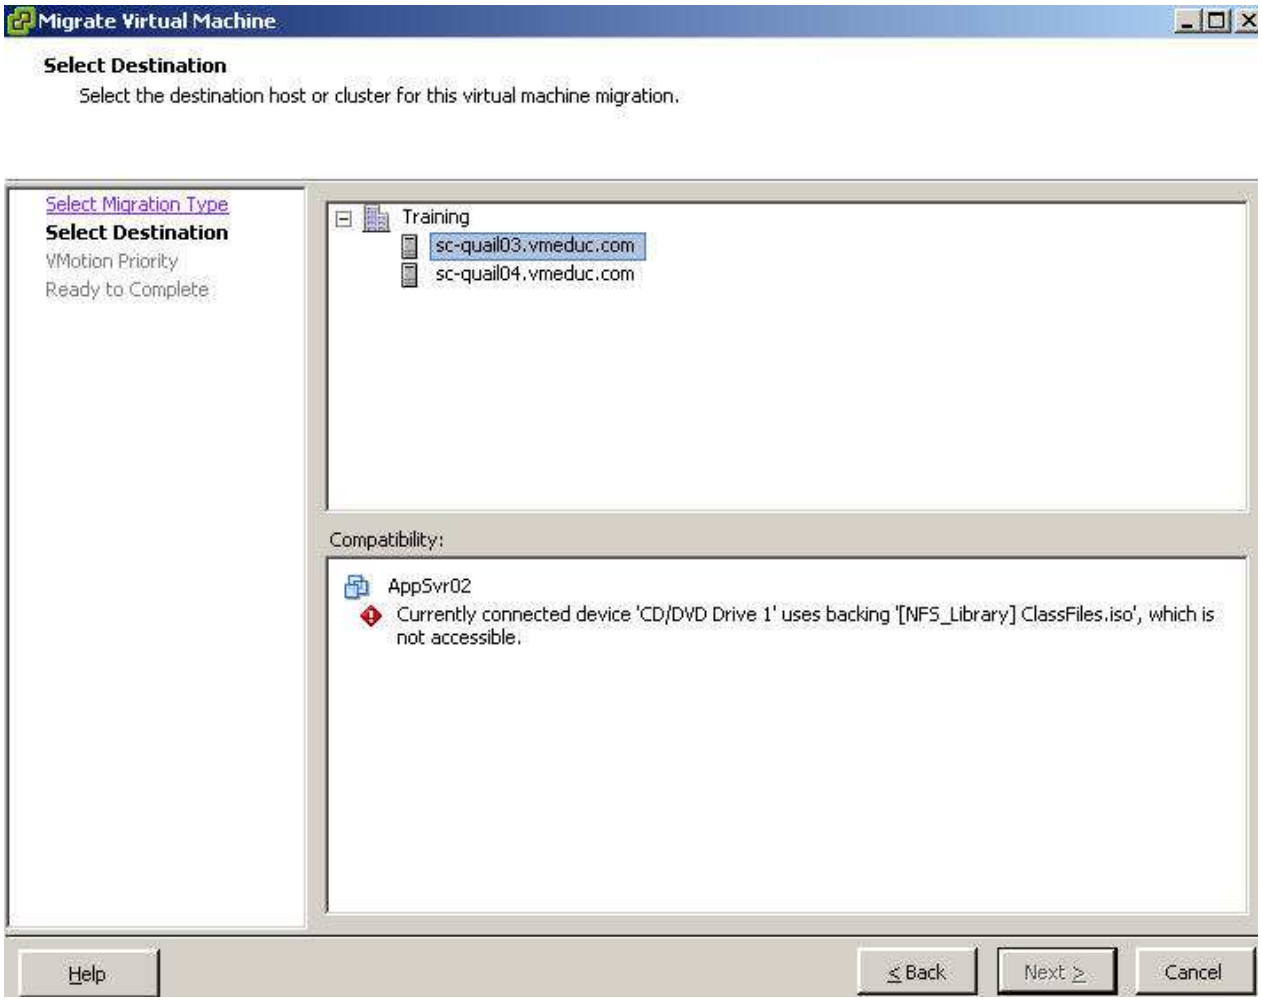

Lab 18

In this lab, you will migrate virtual machines using VMotion.

1. Add a second ESX host to the Training datacenter.
2. Add a second ESX host to the VMotion distributed switch.
3. Create a VMkernel port for your host on the VMotion distributed switch.
4. Verify that your virtual machines' settings meet VMotion requirements.
5. Verify that your ESX host meets VMotion requirements.
6. Connect virtual machines to the Production network of the lower-number ESX host.
7. Perform a VMotion migration of your virtual machine.

Lesson Summary

- VMotion is the underlying technology required for DRS to function properly.
- Use the host or virtual machine's Maps tab to help you verify your VMotion layout.
- If a virtual machine or host does not meet one or more VMotion requirements, the Migrate Virtual Machine wizard catches any inconsistencies during its validation process.

Lesson 4: VMware Distributed Resource Scheduler

Lesson Objectives

- Describe the functions of a DRS cluster
- Explain the benefits of EVC
- Create a DRS cluster
- View information about a DRS cluster
- Remove a host from a DRS cluster

What Is a DRS Cluster?

A cluster is a collection of ESX hosts and associated virtual machines.

A DRS cluster is managed by vCenter Server and has these resource management capabilities:

- Initial placement
- Load balancing
- Power management
- Virtual machine affinity rules

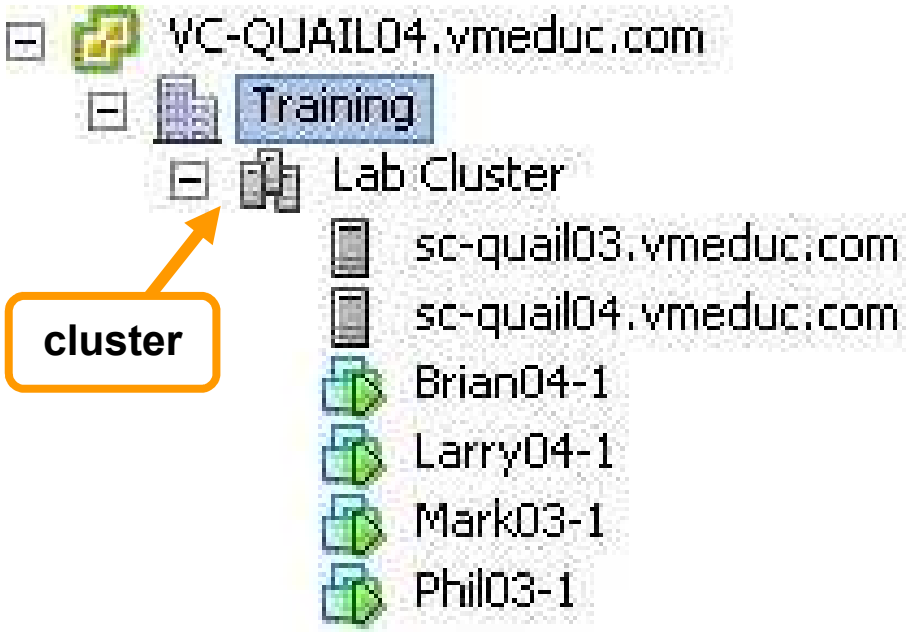

DRS Cluster Prerequisites

DRS works best if the virtual machines meet VMotion requirements.

To use DRS for load balancing, the hosts in the cluster must be part of a VMotion network.

- If not, DRS can still make initial placement recommendations.

Configure all hosts in the cluster to used shared VMware vStorage VMFS volumes.

- Volumes must be accessible by all hosts.
- Volumes must be large enough to store all virtual disks for your virtual machine.

Creating a DRS Cluster

Right-click your datacenter.

Select New Cluster**.**

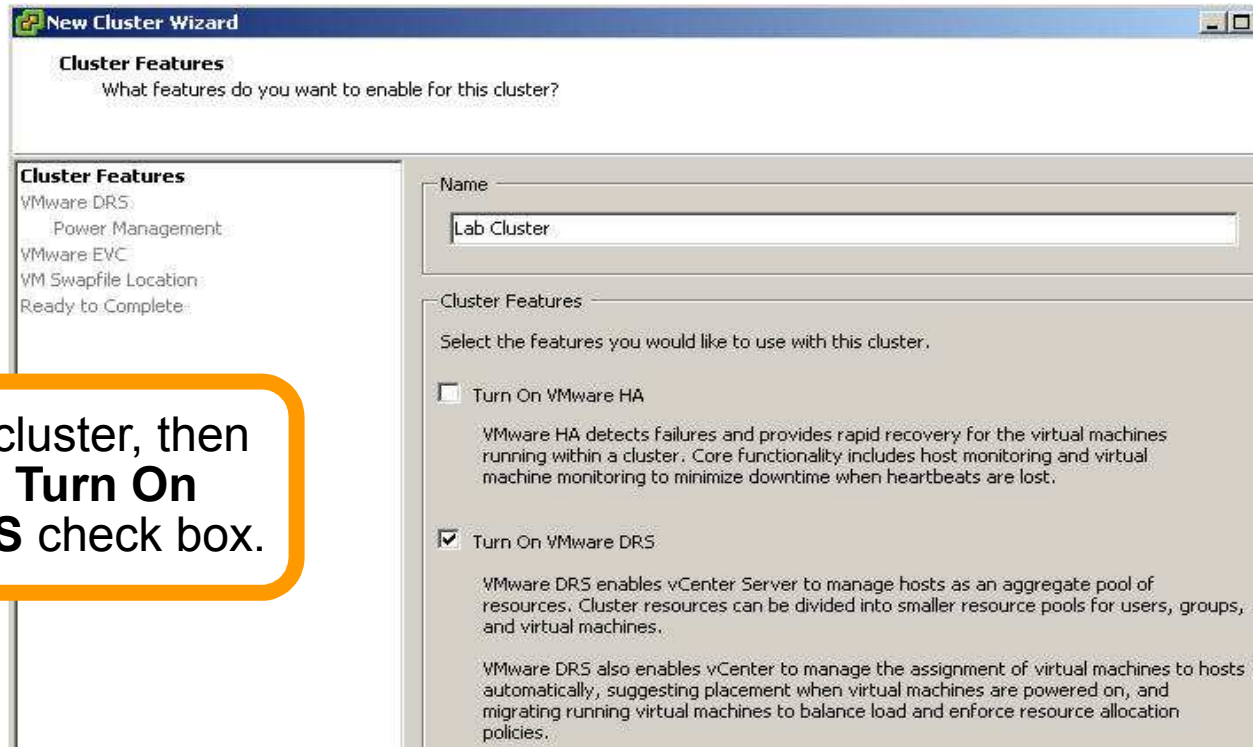

Name your cluster, then select the **Turn On VMware DRS** check box.

DRS Cluster Settings: Automation Level

Configure the automation level for initial placement of VMs and dynamic balancing while VMs are running.

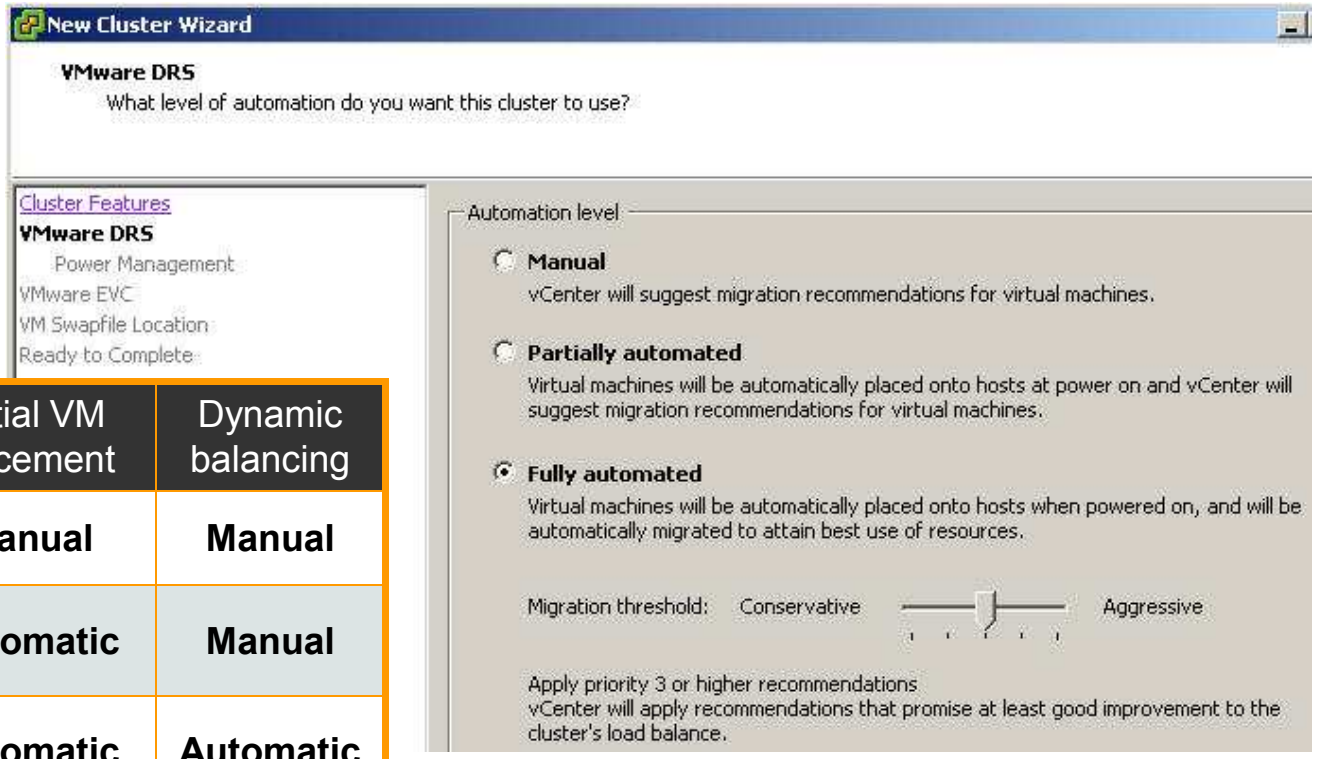

Automation level	Initial VM placement	Dynamic balancing
Manual	**Manual**	**Manual**
Partially automated	**Automatic**	**Manual**
Fully automated	**Automatic**	**Automatic**

DRS Cluster Settings: Migration Threshold

The migration threshold levels determine how quickly virtual machines are migrated.

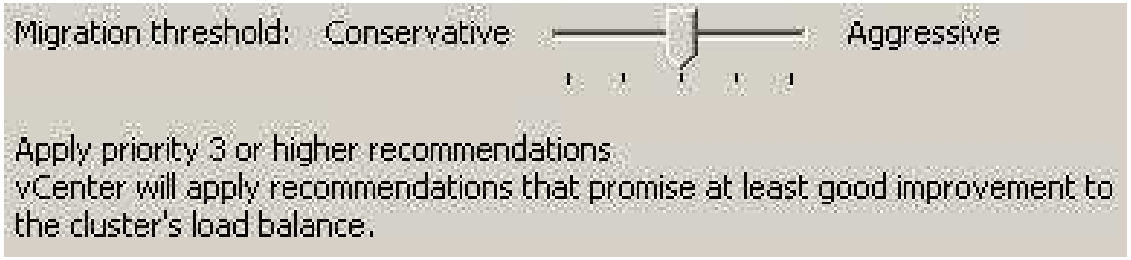

Priority level	Apply all recommendations
1 – Most conservative	with five stars only
2 – Moderately conservative	with four or more stars
3 – Midpoint (default)	with three or more stars
4 – Moderately aggressive	with two or more stars
5 – Aggressive	with one or more stars

DRS Cluster Settings: Power Management

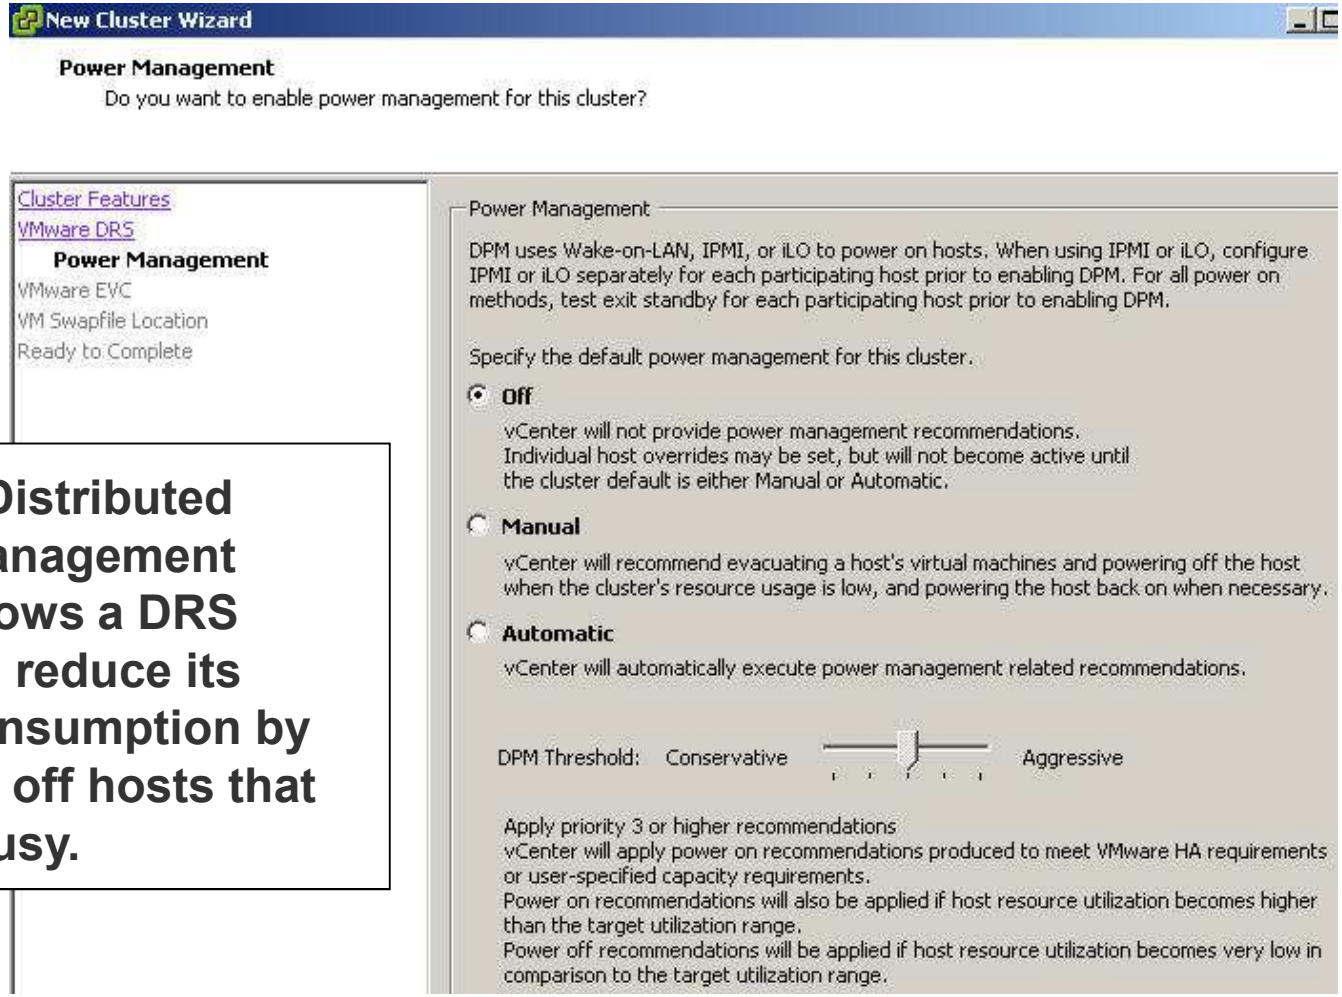

VMware Distributed Power Management (DPM) allows a DRS cluster to reduce its power consumption by powering off hosts that are not busy.

DRS Cluster Settings: VMware EVC

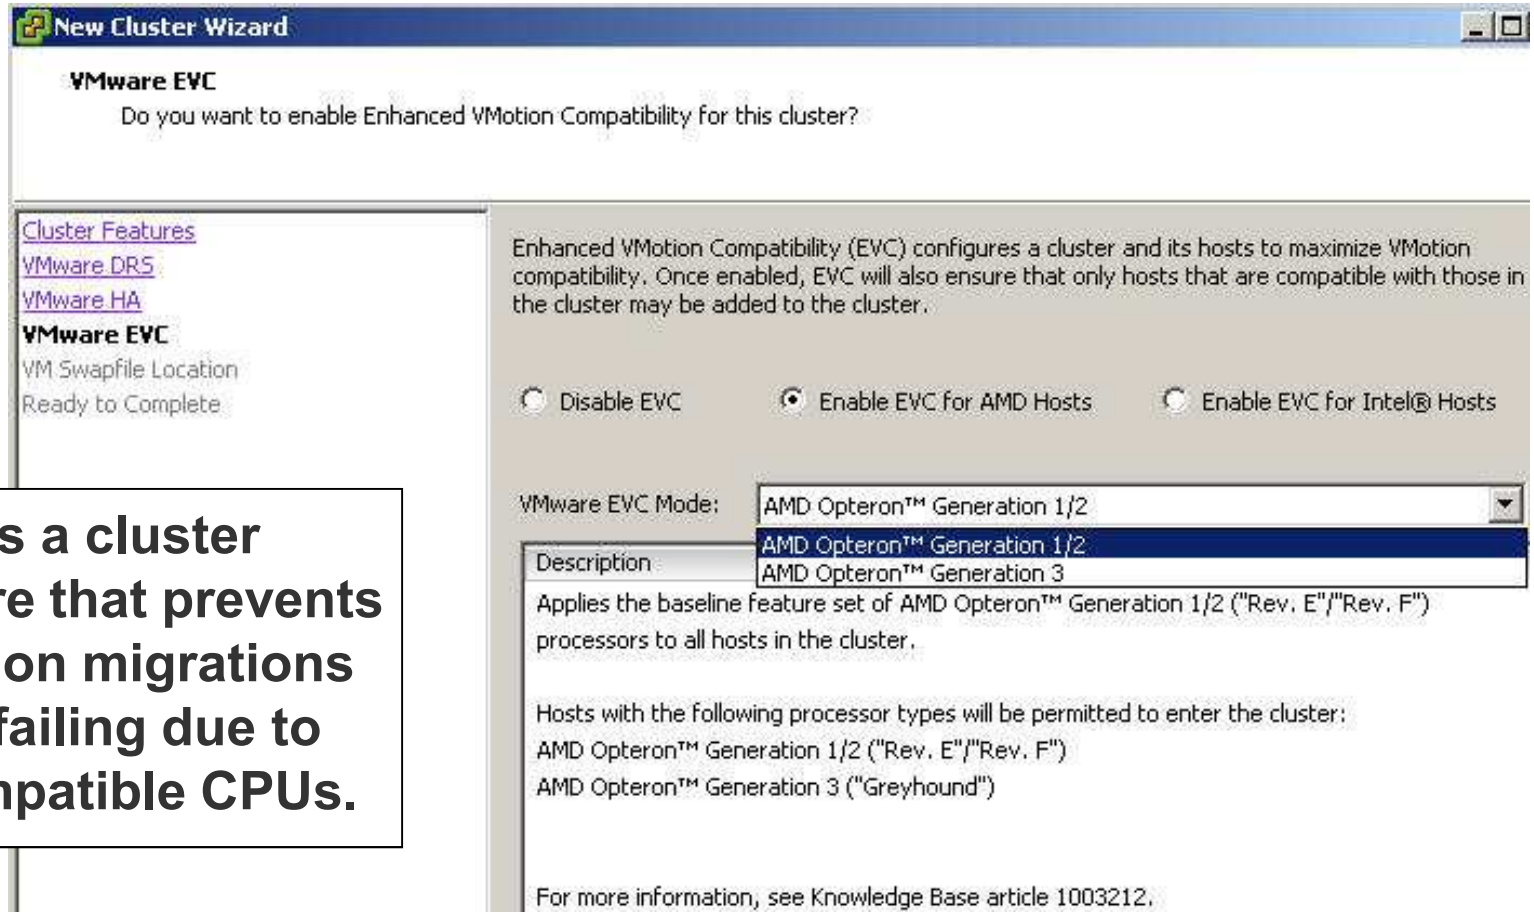

EVC is a cluster feature that prevents VMotion migrations from failing due to incompatible CPUs.

CPU Baselines for an EVC Cluster

EVC works at the cluster level using CPU baselines to configure all processors included in the EVC-enabled cluster.

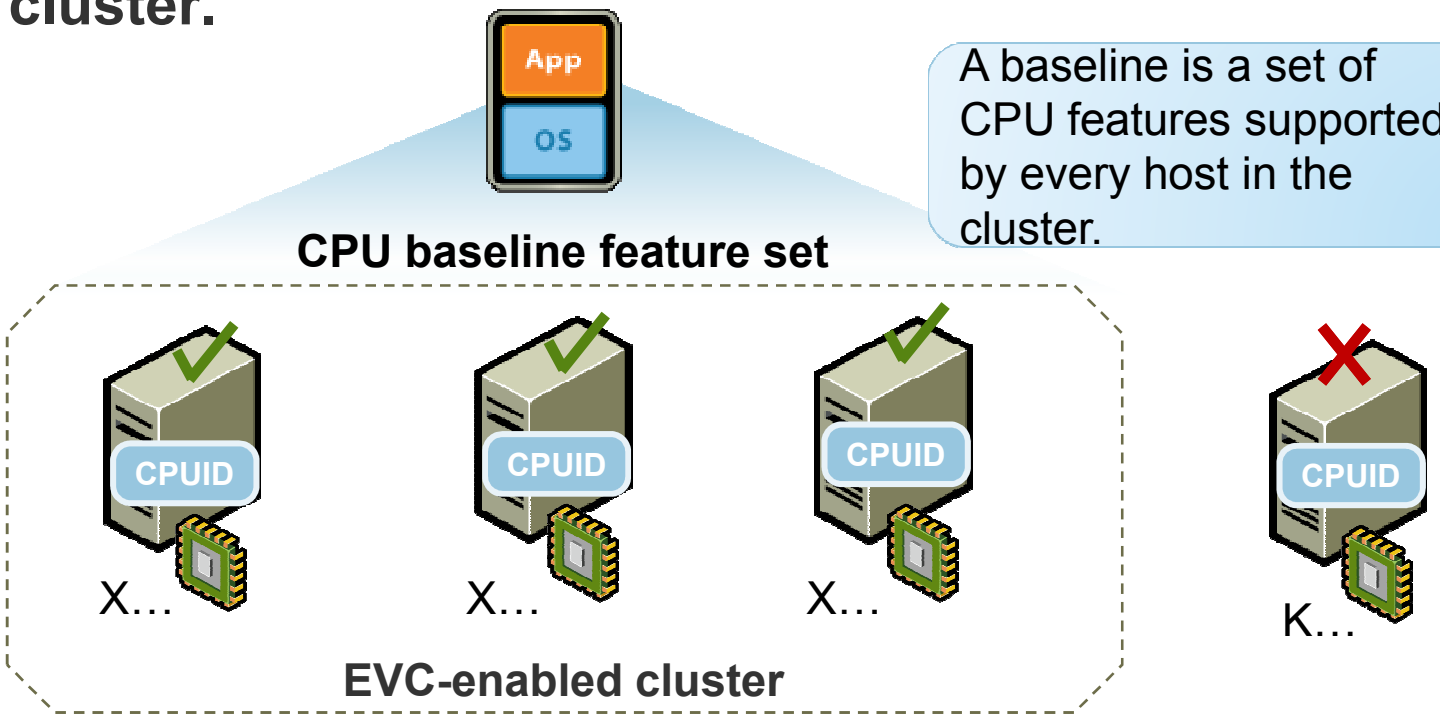

EVC Cluster Requirements

All hosts in the cluster must meet the following requirements:

- Use CPUs from a single vendor (either Intel or AMD)
  - Use Intel CPUs with Core 2 micro architecture and newer
  - Use AMD second-generation Opteron CPUs and newer
- Run ESX 3.5 Update 2 or later
- Be connected to vCenter Server
- Be enabled for hardware virtualization (AMD-V or Intel VT)
- Be enabled for execution-disable technology (AMD No eXecute (NX) or Intel eXecute Disable (XD))

Applications in virtual machines must be well-behaved.

DRS Cluster Settings: Swap File Location

Store virtual machine's swap file with virtual machine or in a specified datastore.

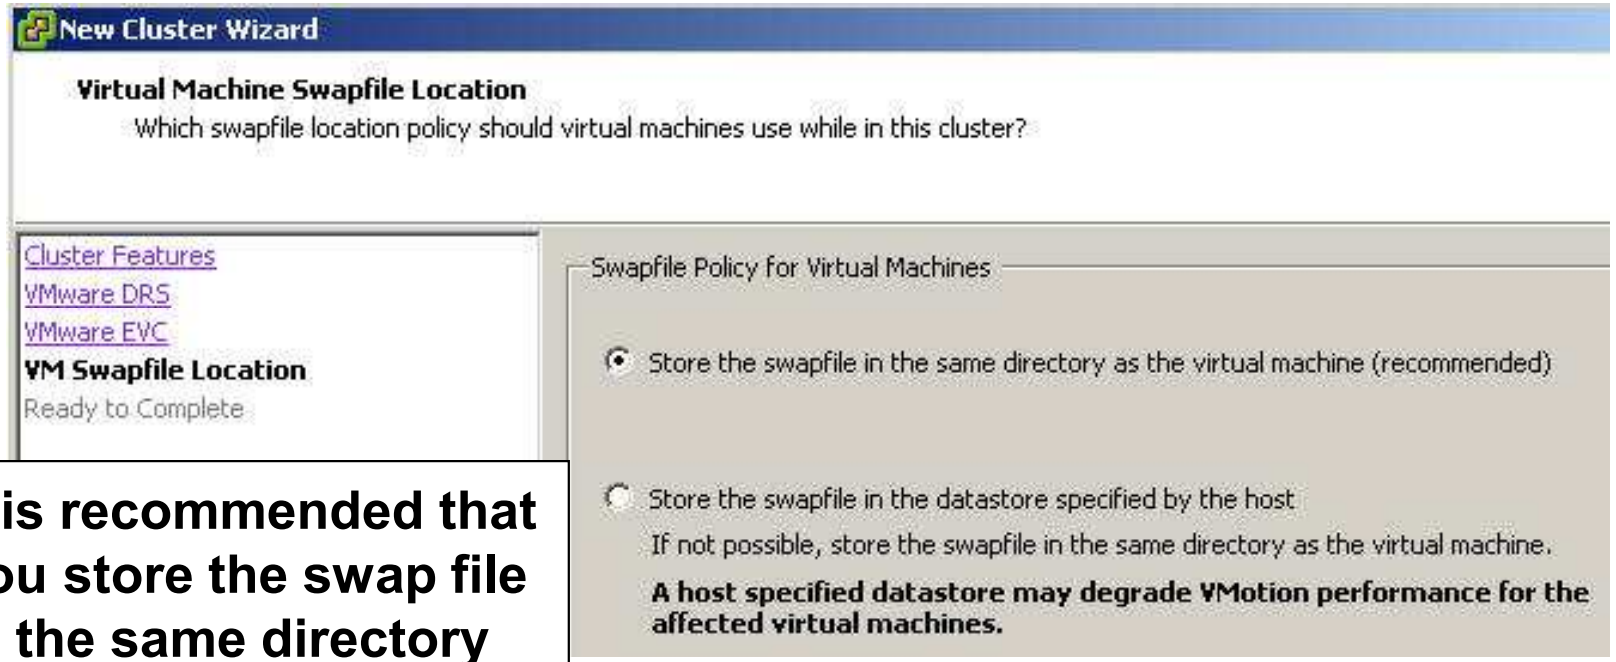

It is recommended that you store the swap file in the same directory as the virtual machine.

Adding Host to Cluster

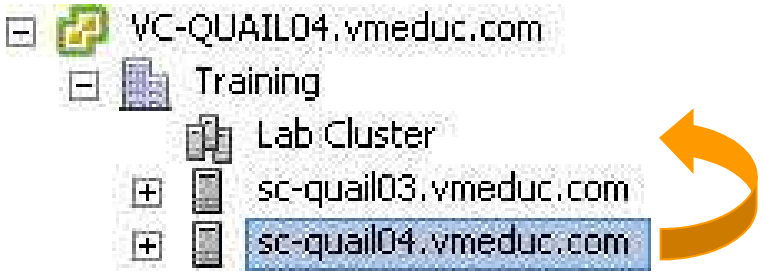

Drag ESX host into cluster.

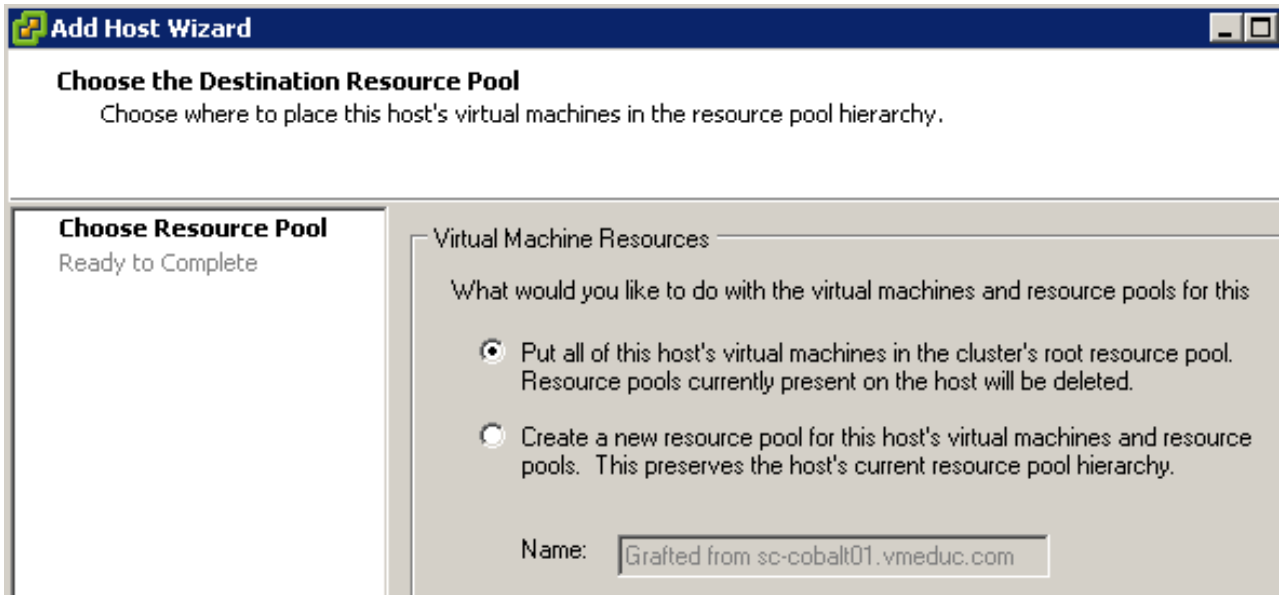

Use the Add Host wizard to complete the process.

Adding Host to Cluster: Resource Pool Hierarchy

When adding a new host or moving an existing host into the cluster, you can keep the resource pool hierarchy of the existing host.

For example, add sc-quail04 to Lab Cluster.

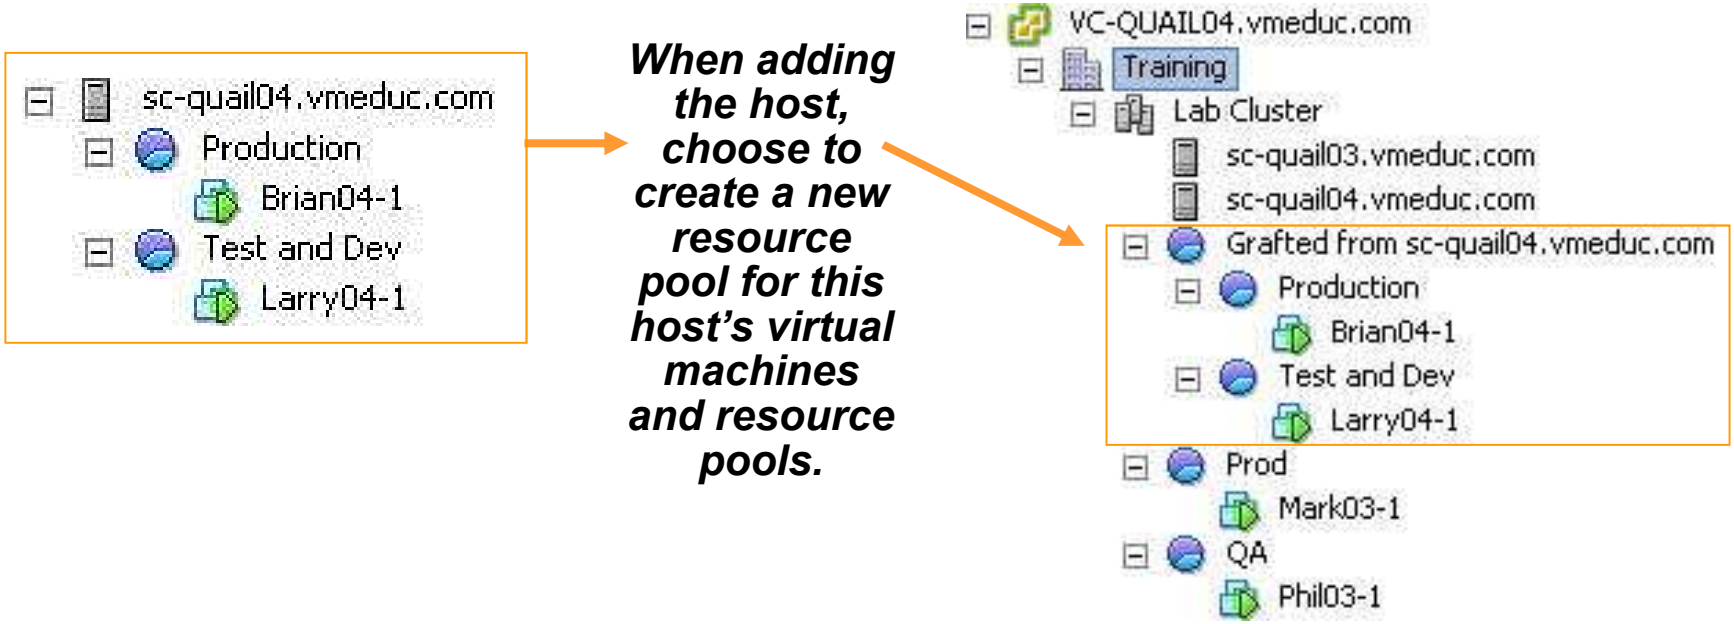

DRS Cluster Settings: Affinity Rules

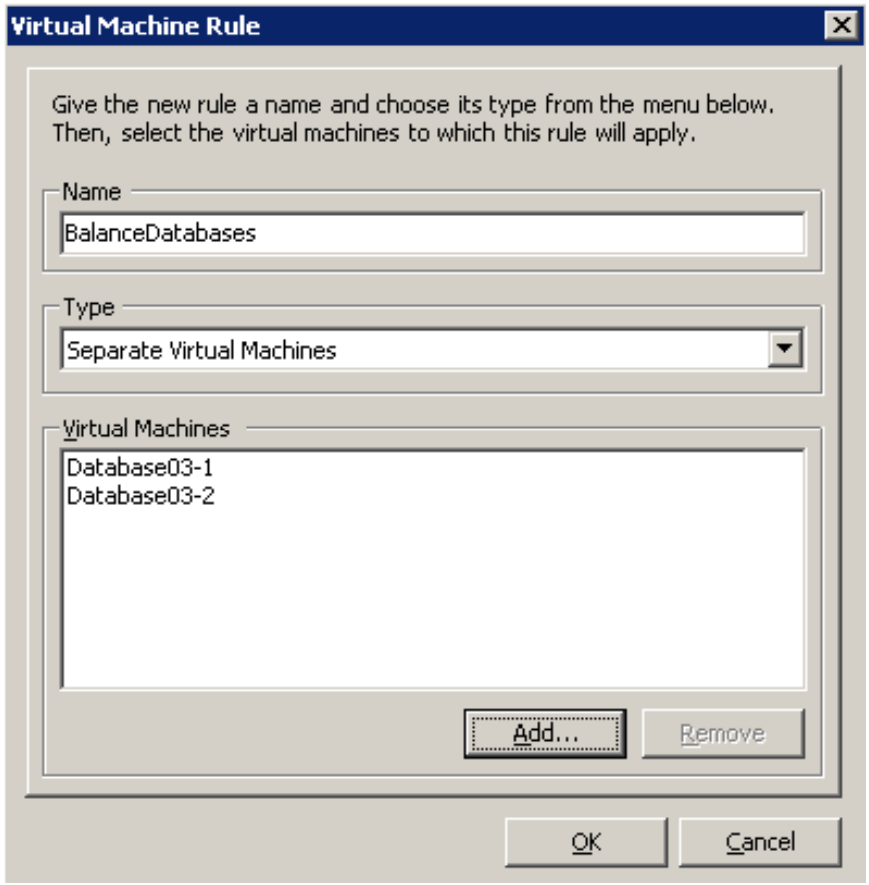

DRS affinity rules specify that either selected virtual machines should be placed on the same host or on different hosts (anti-affinity rule).

Affinity rules

- Use for multi-VM systems where performance benefits.

Anti-affinity rules

- Use for multi-VM systems that load-balance or require high availability.

DRS Cluster Settings: VM-Level Automation

Optionally, set automation level per virtual machine.

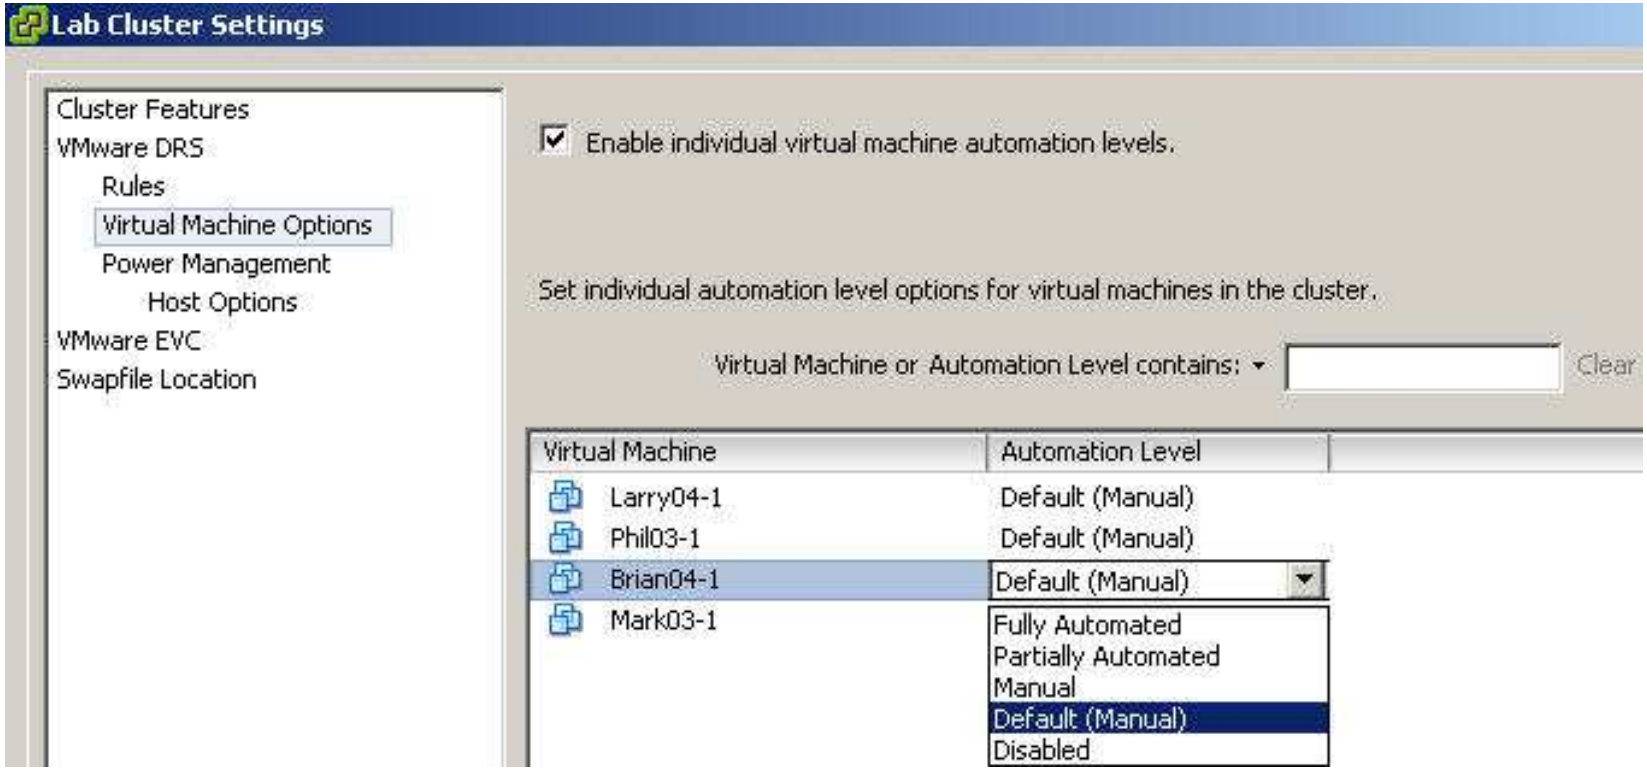

Viewing General Cluster Information

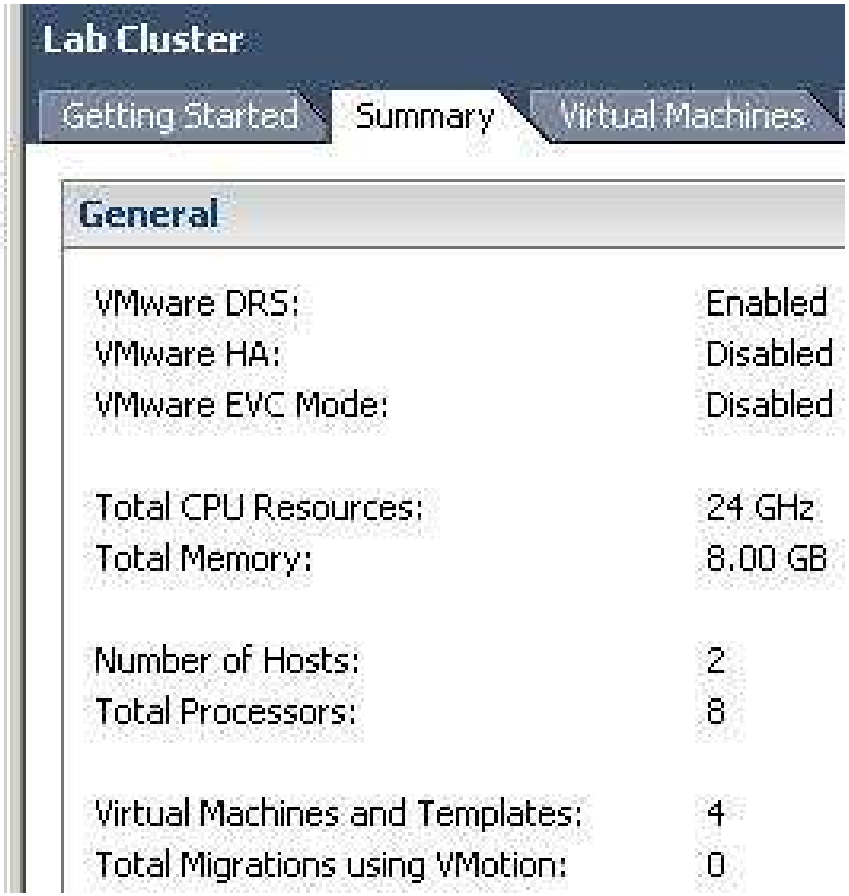

The cluster Summary **tab provides useful information about the configuration and operation of your cluster.**

Viewing DRS Cluster Information

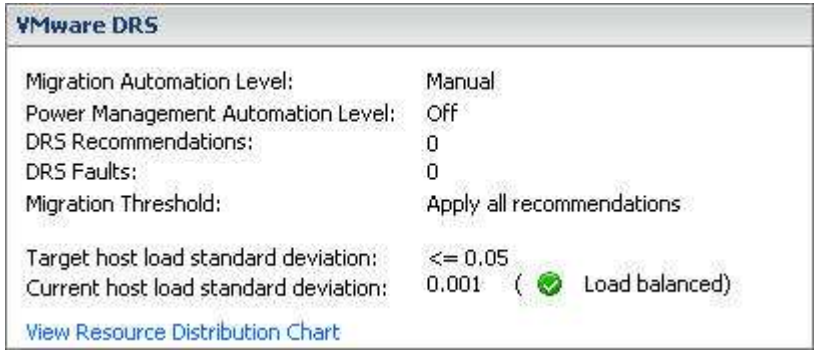

The cluster Summary **tab also provides information specific to DRS.**

The View Resource Distribution Chart **link shows CPU and memory utilization per host.**

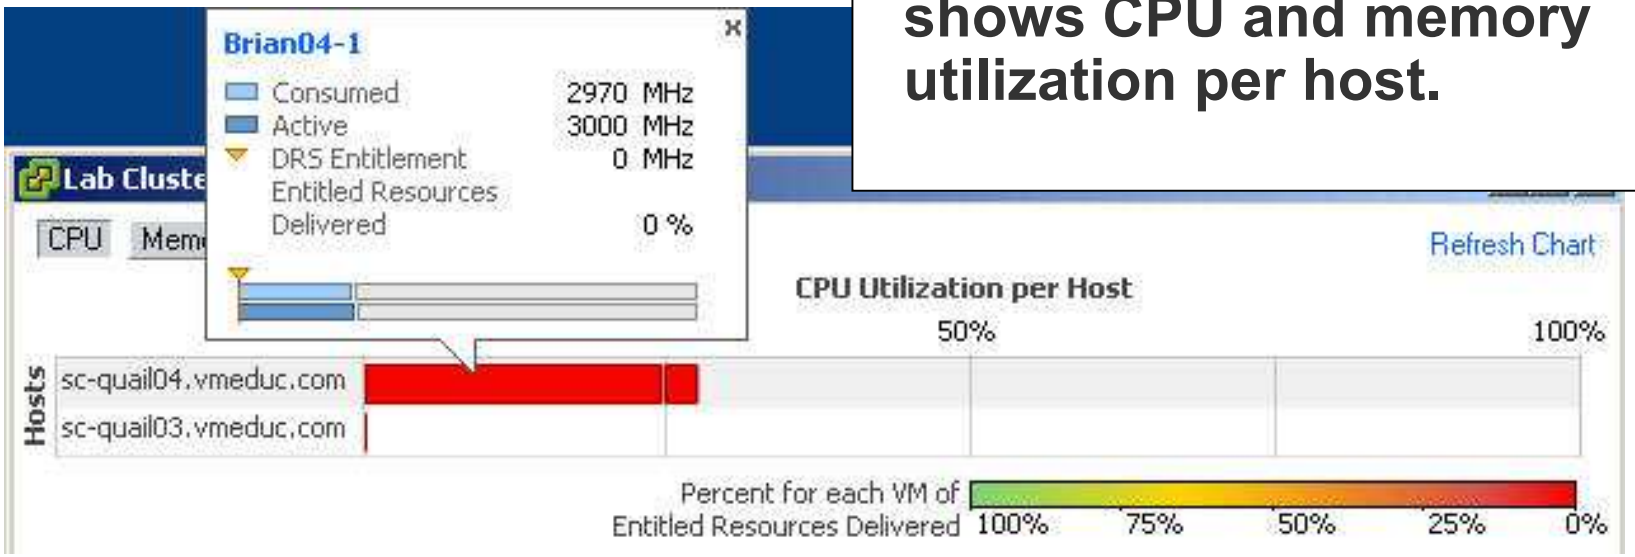

Viewing DRS Resource Allocation

The cluster Resource Allocation **tab displays information about the CPU and memory resources in the cluster.**

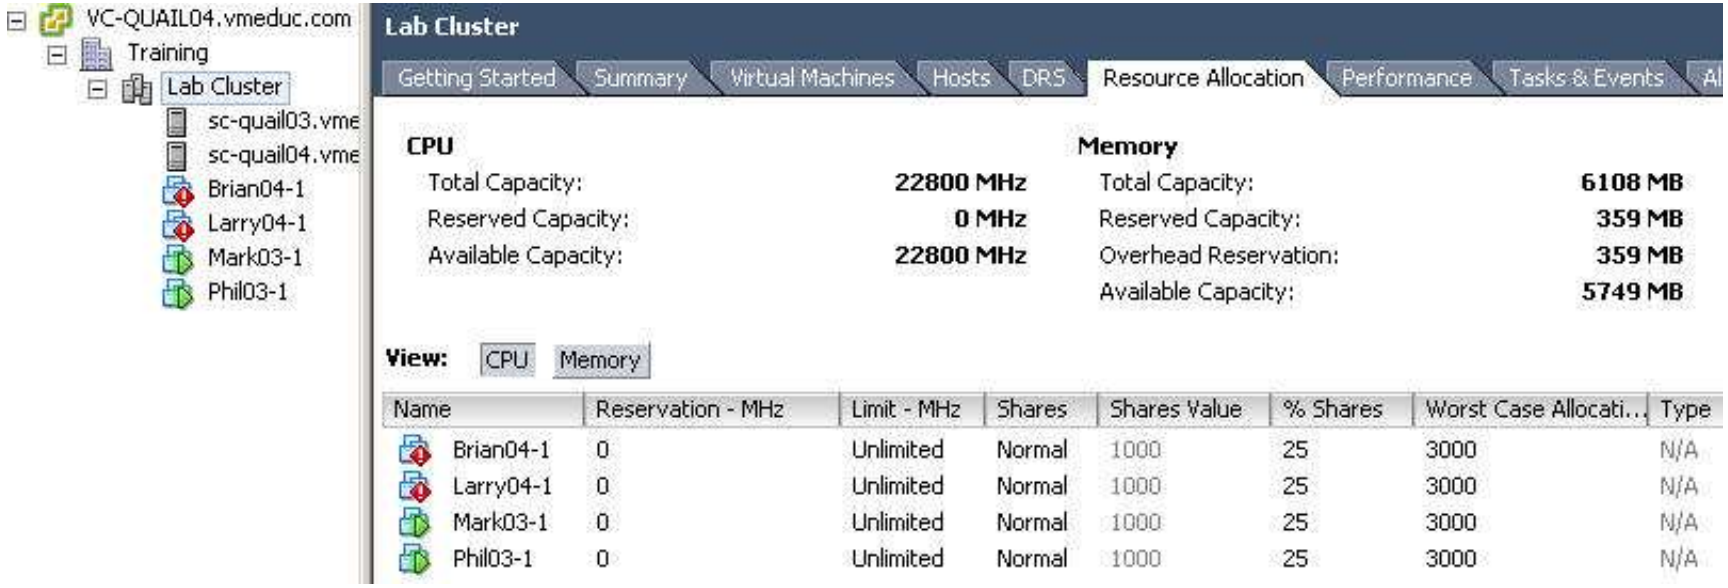

Viewing DRS Recommendations

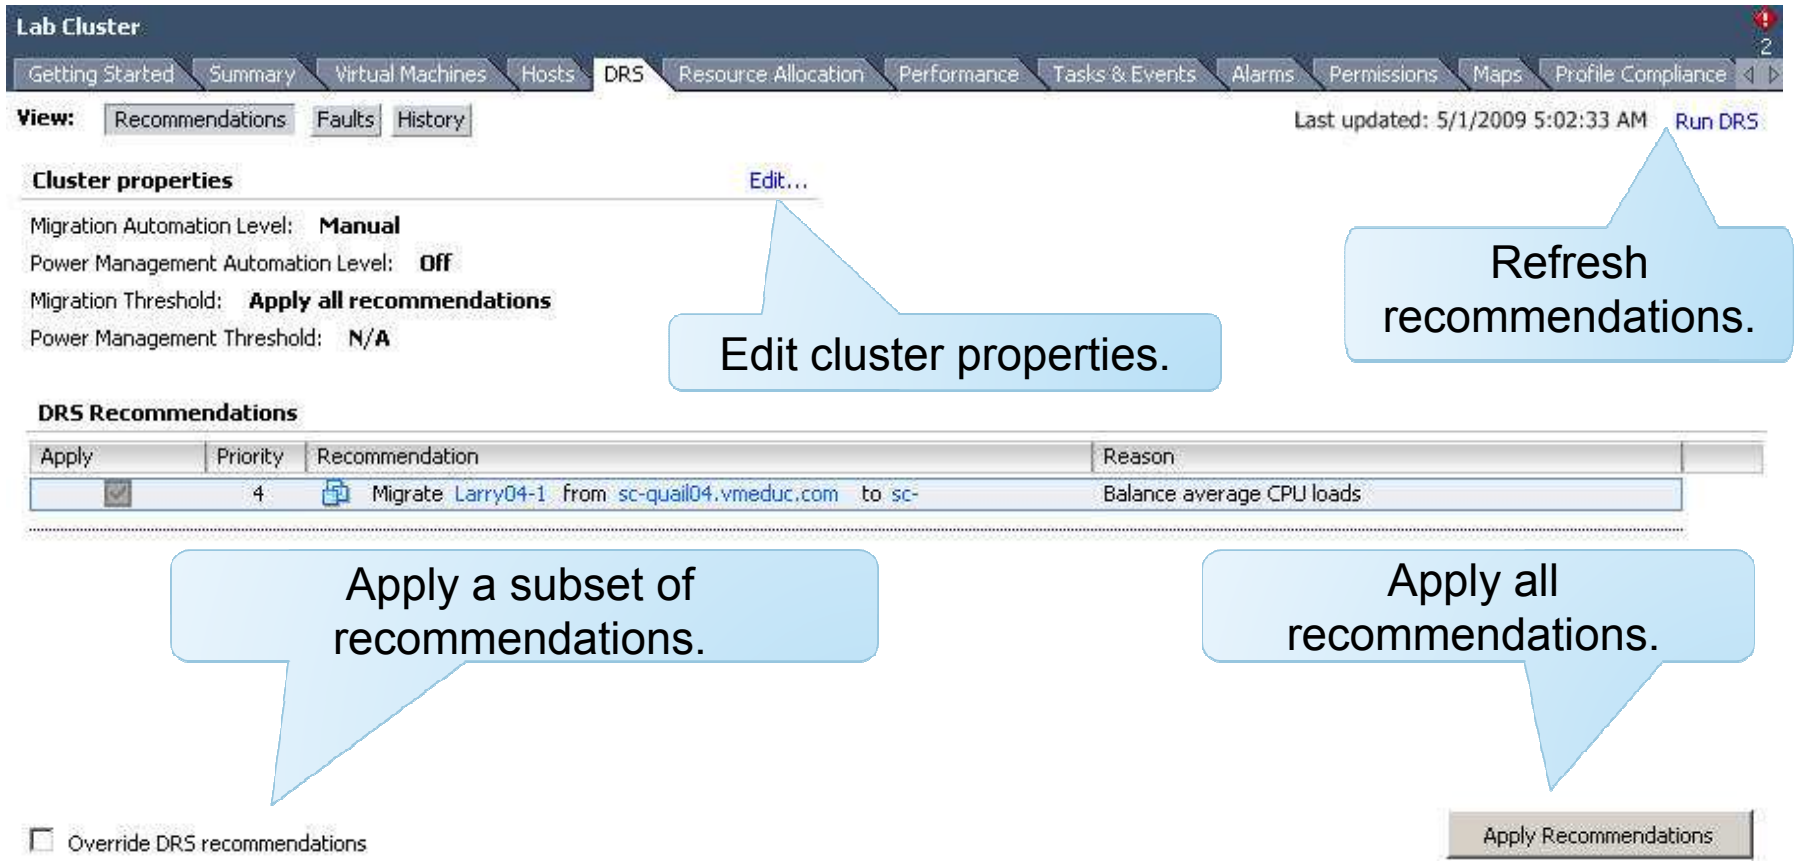

Monitoring Cluster Status

View the inventory hierarchy for the cluster state.

View the cluster's Tasks & Events **tab for further information.**

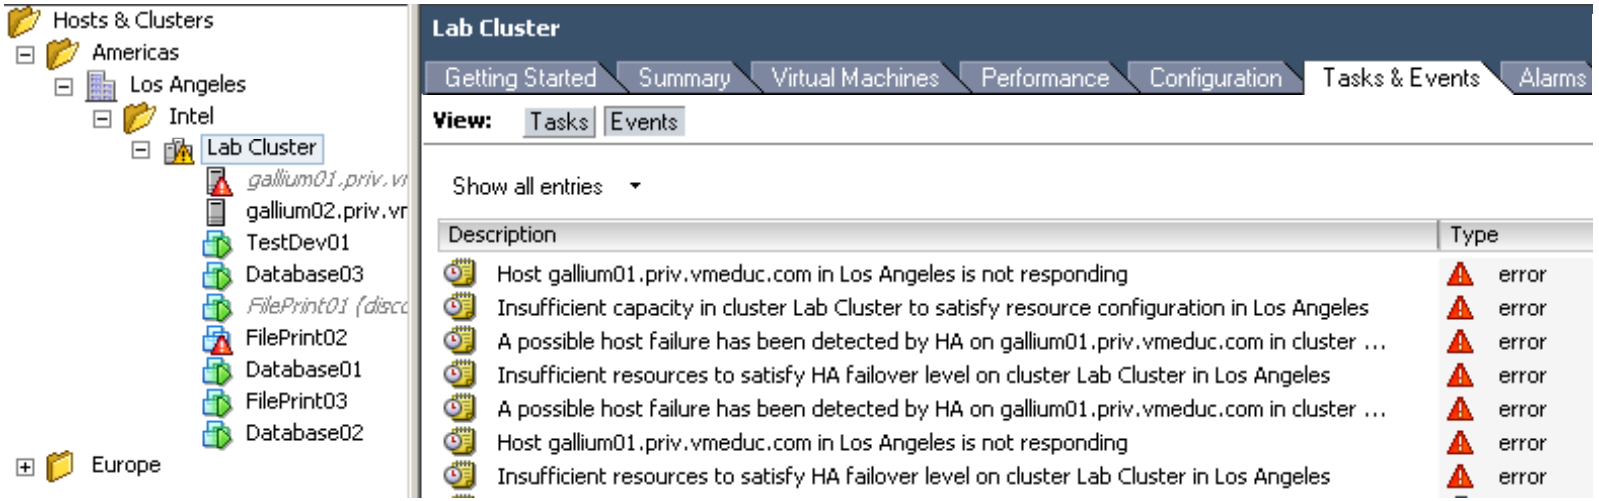

Maintenance Mode and Standby Mode

To service a host in a cluster (for example, to install more memory) or remove a host from a cluster, you must place it in maintenance mode.

- Virtual machines on the host should be migrated to another host or shut down.
- You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.
- While in maintenance mode, the host does not allow you to deploy or power on a virtual machine.

When a host is placed in standby mode, it is powered off.

- This mode is used by DPM to optimize power usage.

Removing a Host from the DRS Cluster

To remove a host from a cluster:

1. Right-click the host and choose **Enter Maintenance Mode**.
2. After the host is in maintenance mode, drag it to a different inventory location.

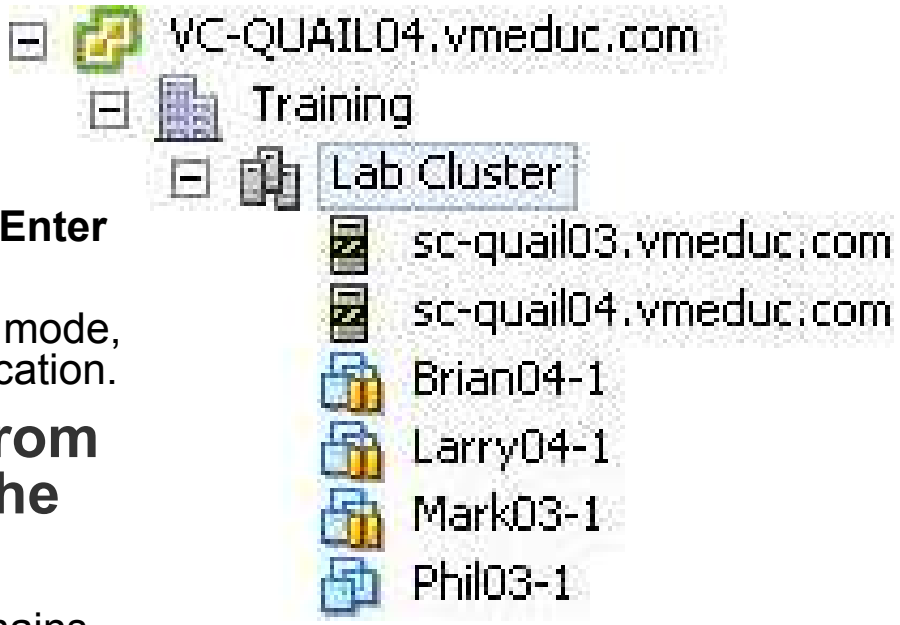

Before removing a host from a DRS cluster, consider the following issues:

- The resource pool hierarchy remains with the cluster.
- Because a host must be in maintenance mode, all virtual machines running on that host are powered off.
- The resources available for the cluster decrease.

Lab 19 and eLearning Activity

In this lab, you will implement a DRS cluster.

1. Create a DRS cluster.
2. Populate the DRS cluster.
3. Verify DRS cluster functionality.

In this eLearning activity, you will view a self-paced demonstration on how to configure and use Enhanced VMotion Compatibility.

- Ask your instructor for access to the eLearning module.

Lesson Summary

- A DRS cluster manages CPU and memory resources by initially placing virtual machines on hosts and balancing virtual machines across hosts.
- DPM allows a DRS cluster to reduce its power consumption by comparing per-host capacity versus demand and then taking, or recommending, the appropriate actions.
- EVC prevents VMotion migrations from failing due to incompatible CPUs.

Key Points

- Shares, limits, reservations, and resource pools are mechanisms for managing CPU and memory resource allocations.
- Storage multipathing and NIC teaming are mechanisms for scaling storage and network management.
- DRS clusters provide automated resource management for multiple ESX/ESXi hosts.